IEEE Transactions on Image Processing, 5(9): 1377-1382, 1996

Steerable Wedge Filters for Local Orientation Analysis

Eero P. Simoncelli and Hany Farid
GRASP Laboratory, Dept. of Computer & Information Science
University of Pennsylvania
Philadelphia, PA 19104-6228

Abstract

Steerable filters have been used to analyze local orientation patterns in imagery. Such filters are typically based on directional derivatives, whose symmetry produces orientation responses that are periodic with period π, independent of image structure. We present a more general set of steerable filters that alleviate this problem.

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

1 Introduction

Oriented linear filters are used in many vision and image processing tasks, such as edge detection, segmentation, texture analysis, and motion analysis. Their basic behavior with regard to representation of orientation may be examined by computing an "orientation map": the squared filter response as a function of filter orientation (e.g., [4, 3, 8, 2]). Such filters are almost always either symmetric or anti-symmetric. The symmetry (or anti-symmetry) of the filters imposes a periodicity (of period π) on the orientation map, regardless of the underlying image structure. For example, an orientation map computed at the end of a line segment will produce a *bimodal* response. Such an ambiguity is undesirable for many applications. For this reason, some authors have recently begun to explore the use of asymmetric filters for orientation analysis [6, 7, 9].

A secondary theme in this sort of orientation analysis is that of rotation-invariance [1, 4, 5, 2, 7]. Along these lines, Freeman and Adelson developed the concept of *steerable filters*, in which an oriented filter is synthesized *exactly* from a linear combination of a fixed set of basis filters. They constructed such bases using directional derivatives of Gaussians, and used these filters to compute local orientation maps. But directional derivatives of Gaussians are either symmetric or anti-symmetric, and thus they suffer from the periodicity problem mentioned above.

In this article, we describe a new class of filters for local orientation and junction analysis. These filters are both asymmetric and steerable, and are designed to produce an optimally localized oriented energy map. A preliminary report of this work has been presented in [10].

2 Steerability and Orientation Maps

This section gives a brief overview of the principle of steerability, and describes its use in computing orientation maps. The simplest example of a steerable filter is the partial derivative of a two-dimensional Gaussian. In polar coordinates, the horizontal and vertical derivatives are written:

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

where the subscript indicates the derivative order, and the parenthesized superscript indicates the derivative direction. It is well-known (and easy to verify, using basic trigonometric identities) that $G_1(r,\theta)$ can be synthesized at an arbitrary orientation, ϕ, using the following equation:

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

This equation embodies the steerability of these functions: the directional derivative G_1 can be generated at an arbitrary orientation ϕ via a linear combination of the rotated *basis filters*, $G_1^{(0)}$ and $G_1^{(\pi/2)}$. The coefficients $\cos(\phi)$ and $\sin(\phi)$ are referred to as the *interpolation functions*. Since convolution is a linear operation, the result of convolving with an arbitrarily oriented filter may be computed via a linear combination of the results of convolving with the basis filters.

Steerability is not limited to first derivative functions. The general steerability condition, for polar-separable functions, is written as:

$$f^{(\alpha)}(r,\phi) = h(\phi-\alpha)g(r) = \sum_{n=1}^{\hat{N}} k_n(\alpha)h(\phi-\alpha_n)g(r), \quad (2)$$

where $h(\cdot)$ is the angular portion of the steerable filter, $g(\cdot)$ is the radial portion, $k_n(\cdot)$ are interpolation functions, and α_n are a fixed set of $\hat{N}$ orientations[1]. Freeman & Adelson [2] showed that this equation is satisfied by all functions with angular components that are bandlimited to contain no more than $\hat{N}/2$ harmonic terms[2]. They presented examples of steerable filter sets consisting of higher-order directional derivatives of a Gaussian, along with steerable approximations to their Hilbert transforms.

Results of applying a steerable 4th-order directional derivative of a Gaussian and an approximation to its Hilbert transform[3] to several synthetic

[1] The partial derivatives form 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 $\alpha_n \in \{0, \pi/2\}$.

[2] This result follows from applying the Nyquist sampling theorem to the angular component.

[3] These filters are used in [2] and are notated as G_4/H_4.

images are shown in Figure 1. The response of the filters to the vertical line and cross are as expected. However, the filters respond bimodally 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 bimodal, but the peaks occur near $\phi = 3\pi/4$ and $\phi = 7\pi/4$ instead of $\phi = \pi/2$ and $\phi = \pi$. Similar problems occur with the "T-" and "Y-junctions".

These "incorrect" responses are due primarily to the inherent symmetries of directional derivatives: an even-order derivative is symmetric (when reflected through the origin) and an odd-order derivative is anti-symmetric. In either case, their squared responses (as a function of orientation) will always be symmetric. Thus, the associated orientation map will be periodic with period π.

Although the directional derivatives provide a nice example, the principle of steerability is not limited to such functions. Given that directional derivatives always exhibit symmetry (or anti-symmetry), we set out to design a more general asymmetric steerable function.

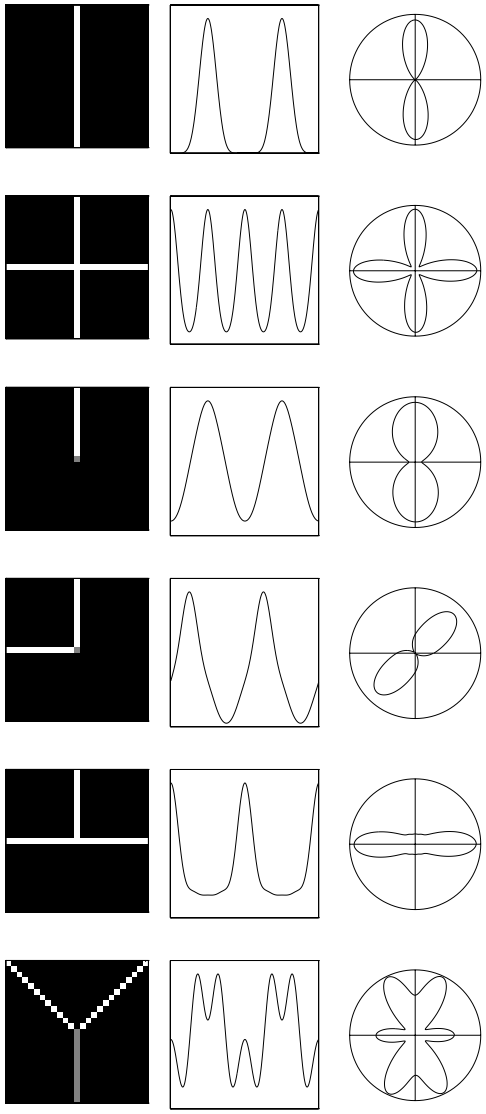

Figure 1: Orientation maps computed using G_4/H_4 symmetric steerable filters from [2] (full filter set contains eleven 13-tap filters). Left column: synthetic line images. Middle column: oriented energy as a function of angle (plotted from 0 to 2π), computed using filters centered on the corresponding image. Right column: polar plots of the same oriented energy functions. Note the symmetry in the oriented energy maps, even when the underlying image structure is asymmetric (half-line, corner, "T-junction", and "Y-junction").

3 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 independent designs for the angular and radial filter components.

3.1 Angular Function

To achieve steerability with $\hat{N} = 2N$ filters, it is sufficient that the angular portion of each filter is constrained to be a weighted sum of the first N circular harmonics[4]. In addition, we impose a Hilbert transform relationship (i.e., a 90-degree phase shift) between the *angular portions* of the two filters:

$$h_e(\phi) = \sum_{n=1}^{N} w_n \cos(n\phi), \quad h_o(\phi) = \sum_{n=1}^{N} w_n \sin(n\phi). \quad (3)$$

[4] Note that there is no $n = 0$ term, since we desire zero-mean filters.

Note that although these *angular* functions are Hilbert transforms of each other, the resulting two-dimensional filters, $f_e(r,\phi)$ and $f_o(r,\phi)$, will *not* form a Hilbert transform pair.[5] Nevertheless, the opposite symmetry of the two angular functions ensures that the orientation operator will respond equally to both line and step edges.

The angular functions in Equation (3) provide a generalization of the class of directional derivatives. An Nth order directional derivative has an angular component containing a *subset* of the harmonics between the first and the Nth. In particular, for odd N it contains only odd harmonics (up to the Nth), and for even N, only even harmonics.

The weights, w_n, are chosen to maximally localize the orientation map "impulse response": $E(\phi) = h_e^2(\phi) + h_o^2(\phi)$. We compute a total least squares solution, using a quadratic localization function: $\mathcal{P}(w_1, w_2, \ldots, w_N) = \int d\phi\ \lambda^2(\phi) E(\phi)$, where $\lambda(\phi)$ is a monotonically increasing weighting function, which we choose as $\lambda(\phi) = \phi$. To facilitate computer implementation, $\mathcal{P}(w_1, w_2, \ldots, w_N)$ is rewritten in matrix form by sampling the variable ϕ at M locations:

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

where $\vec{w}$ is a vector whose components are the weights w_n, Λ is a diagonal matrix containing the values $\lambda(\phi)$, 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. The unit vector $\vec{w}$ that minimizes $P(\vec{w})$ is simply the minimal-eigenvalue eigenvector of the matrix $[C^t \Lambda^t \Lambda C + S^t \Lambda^t \Lambda S]$.

Examples of 1D even- and odd-symmetric angular functions, with corresponding oriented energy responses, are shown in Figure 2. Note that as N increases (i.e., higher harmonics are added), the angular functions and their associated oriented energy responses become narrower. Of course, this narrowing comes at a price: the number of filters in the basis set increases by two with the addition of each harmonic term.

[5] It is not possible to construct a pair of (two-dimensional, even- and odd-symmetric) filters that are both steerable and Hilbert transforms of each other. The even-symmetric (odd-symmetric) filter will contain only even (odd) order angular components and thus the power spectra of the two filters cannot be exactly equated.

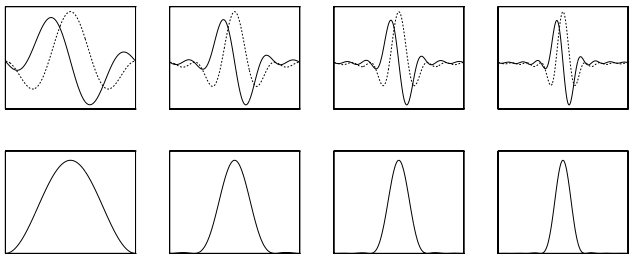

Figure 2: Top: Optimally localized even- and odd-symmetric angular components $h_e(\phi)$ (dashed) and $h_o(\phi)$ (solid), shown for $N = 2, 4, 6, 8$. Bottom: Corresponding oriented energy "impulse" response $E(\phi)$. As N increases, the width of the response decreases, but the number of filters required for a steerable basis set increases.

3.2 Interpolation Functions

We now derive the interpolation functions, $k_n(\alpha)$, for the even-symmetric filters; the derivation for the odd-symmetric filters is nearly identical. From Equation (3), the angular portion of the even-symmetric filter rotated by α is:

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

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

$$\begin{aligned} \vec{k}(\alpha) &= \begin{pmatrix} \cos(\alpha) & \sin(\alpha) & \ldots & \cos(N\alpha) & \sin(N\alpha) \end{pmatrix}, \\ \vec{f}(\phi) &= \begin{pmatrix} w_1\cos(\phi) & w_1\sin(\phi) & \ldots & w_N\cos(N\phi) & w_N\sin(N\phi) \end{pmatrix}^t. \end{aligned}$$

To construct the basis functions, we write $2N$ copies of Equation (5), for each of the fixed angular locations, $\{\alpha_n = \pi n/N; 1 \le n \le 2N\}$, and combine

these into a single matrix equation:

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

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

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

and K is a fixed $2N \times 2N$ matrix:

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

Equation (6) gives a linear relationship between the steerable basis set $\{h(\phi - \alpha_n); 1 \leq n \leq 2N\}$ and a weighted subset of the Fourier basis $\{w_n \cos(n\phi), w_n \sin(n\phi); 1 \leq n \leq N\}$. The matrix K is invertible (since it corresponds to a Discrete Fourier Transform), and thus we can solve for the weighted Fourier basis as a function of the steerable basis:

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

Finally, combining Equation (7) with Equation (5) gives:

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

This equation describes the construction of $h(\phi)$, rotated to arbitrary angle α, as a linear combination of the basis set $\{h(\phi - \alpha_n); 1 \leq n \leq 2N\}$. The interpolation functions are the components of 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). But this full basis is redundant: the entire set of $4N$ functions spans only a $2N$-dimensional functional subspace. As such, we can derive interpolation functions using a *combination* of N even- and N odd-symmetric filters.

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:

$$h_e(\phi - \alpha) = \vec{k}_e(\alpha) \cdot \vec{f}(\phi) \tag{8}$$
$$h_o(\phi - \alpha) = \vec{k}_o(\alpha) \cdot \vec{f}(\phi), \tag{9}$$

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

$$\vec{k}_e(\alpha) = \left(\cos(\alpha) \quad \sin(\alpha) \quad \ldots \quad \cos(N\alpha) \quad \sin(N\alpha) \right),$$
$$\vec{k}_o(\alpha) = \left(-\sin(\alpha) \quad \cos(\alpha) \quad \ldots \quad -\sin(N\alpha) \quad \cos(N\alpha) \right).$$

A collection of such equations for the reduced set of angles, $\{\alpha_n = 2\pi n/N; 1 \leq n \leq N\}$, may be combined in matrix form:

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

$$\begin{pmatrix} \cos(\alpha_1) & \sin(\alpha_1) & \ldots & \cos(N\alpha_1) & \sin(N\alpha_1) \\ -\sin(\alpha_1) & \cos(\alpha_1) & & -\sin(N\alpha_1) & \cos(N\alpha_1) \\ \vdots & & \ddots & & \vdots \\ \cos(\alpha_N) & \sin(\alpha_N) & & \cos(N\alpha_N) & -\sin(N\alpha_N) \\ -\sin(\alpha_N) & \cos(\alpha_N) & \ldots & -\sin(N\alpha_N) & \cos(N\alpha_N) \end{pmatrix} \vec{f}(\phi),$$

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:

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

where the basis set of *both even and odd* steerable filters, rotated to each of the N angles α_n, is contained in the vector $\vec{h'}(\phi)$.

3.3 Radial Function

Thus far, only the angular component of the steerable wedge filters have been specified. An arbitrary radial function may be chosen without affecting the steerability of the resulting filters. We

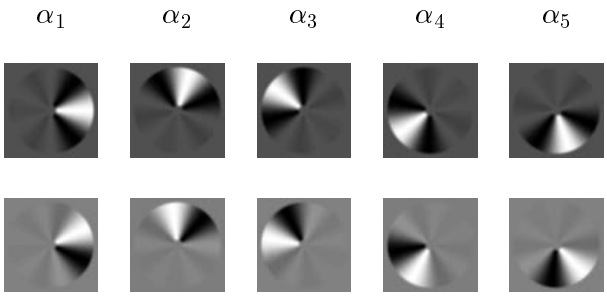

Figure 3: Example set of 10 steerable wedge functions for $N = 5$. From this basis set, a filter of either symmetry (even or odd) may be synthesized at *any* orientation.

choose a radial function that falls smoothly to zero in the center and at the outer edges of the filter:

$$g(r) = \begin{cases} 0, & r < R_1 \\ \frac{1}{2}[1 - \cos(\pi(r - R_1)/\delta_1)], & R_1 \le r < R_1 + \delta_1 \\ 1, & R_1 + \delta_1 \le r < R_2 - \delta_2 \\ \frac{1}{2}[1 + \cos(\pi(r + \delta_2 - R_2)/\delta_2)], & R_2 - \delta_2 \le r < R_2 \\ 0, & R_2 \le r. \end{cases} \quad (12)$$

To create filters of reasonable size, the polar-separable functions, $f^{(\alpha_n)}_{\{e,o\}}(r, \phi) = g(r)h_{\{e,o\}}(\phi - \alpha_n)$, must be coarsely sampled. Sampling is a linear process, and thus does not affect the steerability of the filter set. It may, however, produce spatial aliasing artifacts that destroy the desired orientation tuning properties. In order to avoid this, we densely sampled the functions on a lattice of 63×63 points. For the radial portion of the filters, we used parameter values of $R_1 = 2$, $\delta_1 = 3$, $R_2 = 30$, and $\delta_2 = 3$. These sampled functions are shown (as grayscale images) in Figure 3. These arrays of samples were then repeatedly blurred and subsampled by a factor of two to produce a set of 7×7 filters.

4 Orientation Analysis Examples

The oriented energy, $E(\phi)$, is defined as the sum of squared outputs of a pair of even- and odd-symmetric filters "steered" to the angle ϕ. Examples of this function, computed for several synthetic images, are shown in Figure 4. These results should be compared to those of the G_4/H_4 steerable filter set, shown in Figure 1. The response to the vertical line in the two figures is similar, but the response to the half-line differs: the G_4/H_4 filters respond bimodally, while the wedge filters respond unimodally, as desired.

The response of the wedge filters to the cross and corner are "incorrect" due to the rather broad tuning of the filters. Narrower tuning can be achieved by increasing the number of harmonics, N. Results for the same set of images with a basis set of 18 filters ($N = 9$) are shown in Figure 5. With this larger and more narrowly tuned basis set, the oversmoothed responses to the cross and the corner are eliminated. Also shown in Figure 5 are responses to a "Y-junction" and a "T-junction". The wedge filters give trimodal responses to both, while the G_4/H_4 responses depicted in Figure 1 show bimodal and quadramodal responses. For tasks such as junction analysis, detection, or classification, the response of the wedge filters is clearly favorable.

Figure 6 shows orientation responses for junctions containing step edges (as opposed to lines). These edges are not pixel-centered, and thus the maximal responses do not occur at the orientations of the edges (see the next section for more examples of this situation). Nevertheless, the figure demonstrates that the filters respond to both edges and lines, as desired.

Results of applying 7-tap wedge filters ($N = 9$) to several junction regions in a real image are shown in Figure 7. Responses match the corresponding junctions in each case.

5 Discussion

Symmetric or anti-symmetric oriented filters produce a symmetric orientation response (with period π) even when the underlying image structure does not have such symmetry. This type of response is clearly not desirable for such tasks as junction analysis, detection or classification. The steerable wedge filters presented in this article are localized in their oriented energy response and there-

fore 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 derivative filters) is that they do not have a simple x-y separable implementation. However, we found that 7-tap filters are sufficient to achieve satisfactory results, and thus the computations may be performed with reasonable efficiency.

There are several issues we have left unaddressed in this article. First, we have not discussed the processing that would be necessary to characterize junctions based on the outputs of a set of steerable wedge filters. The examples we have shown indicate that the information necessary for characterizing such junctions is contained in the filter responses. But a full exploration of the problem of rotation-invariant junction classification is beyond the scope of the current article.

Second, we have avoided the question of spatial location. In each of the examples, the filters are centered over the junction. Example filter responses for an off-center junction are shown in Figure 8. These results demonstrate that the orientation response degrades gracefully with spatial displacement. They also suggest that junction detection or classification may require the analysis of orientation maps in a local neighborhood.

We believe that the filters presented here will prove to be useful in analyzing the local orientation structure of images. 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 identifying motion occlusion boundaries.

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

[3] M. Kass and A. Witkin. Analyzing oriented patterns. *Computer Vision, Graphics and Image Processing*, 37:362–385, 1987.

[4] 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.

[5] J. Koenderink and A. vanDoorn. Representation of local geometry in the visual system. *Biological Cybernetics*, 55:367–375, 1987.

[6] R. Mehrotra and S. Nichani. Corner detection. *Pattern Recognition*, 23(11):1223–1233, 1990.

[7] P. Perona. Steerable-scalable kernels for edge detection and junction analysis. *Image and Vision Computing*, 10(10):663–672, 1992.

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

[9] 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.

[10] E. Simoncelli and H. Farid. Steerable wedge filters. In *International Conference on Computer Vision*, Cambridge, MA, June 1995.

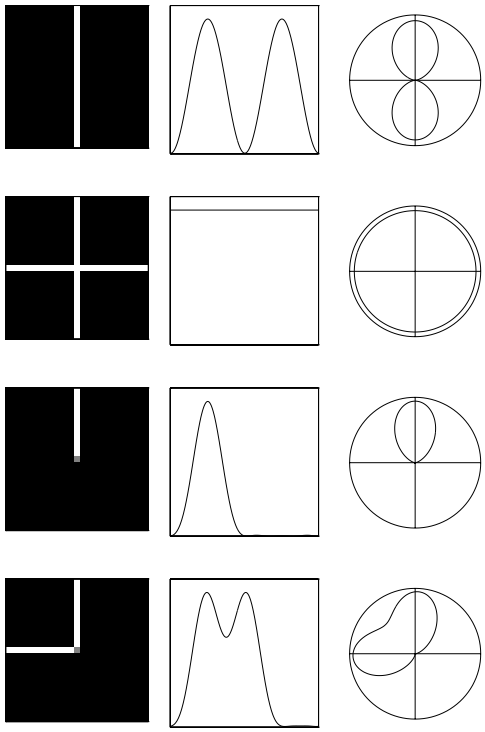

Figure 4: Orientation maps computed using a set of ten steerable wedge filters (i.e., $N = 5$) as shown in Figure 3. Left column: synthetic line images. Middle column: oriented energy as a function of angle (plotted from 0 to 2π), computed using filters centered on the corresponding image. Right column: polar plots of the same oriented energy functions. The responses match the underlying images, except that the broad orientation tuning of the filters leads to an oversmoothed response to the cross and the corner.

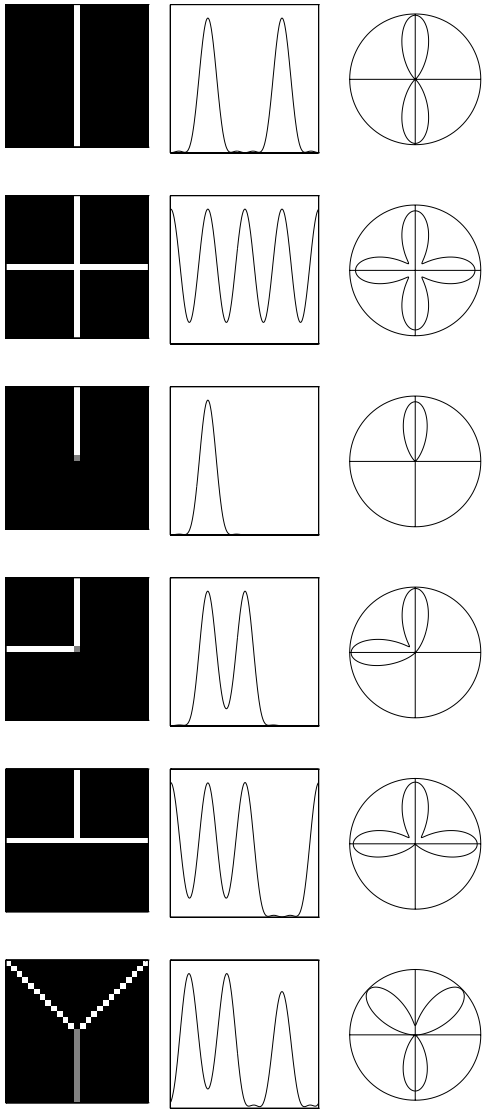

Figure 5: Orientation maps computed using a set of 18 steerable wedge filters (i.e., $N = 9$). Left column: synthetic line images. Middle column: oriented energy as a function of angle (plotted from 0 to 2π), computed using filters centered on the corresponding image. Right column: polar plots of the same oriented energy functions. Note that the narrower tuning (relative to the filters of Figure 4) produces a better representation of the image structure.

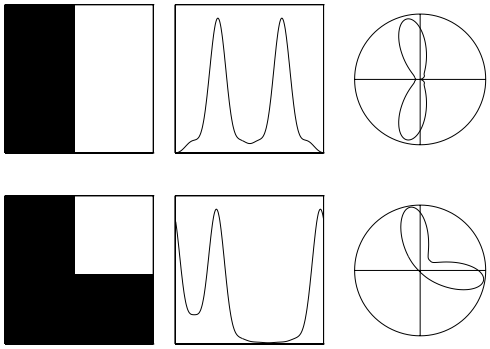

Figure 6: Orientation maps computed using a set of 18 steerable wedge filters (i.e., $N = 9$). Left column: synthetic images containing step edges. Middle column: oriented energy as a function of angle (plotted from 0 to 2π), computed using filters centered on the corresponding image. Right column: polar plots of the same oriented energy functions. Note that the step edges are not pixel-centered, and thus the responses do not lie along the expected axes.

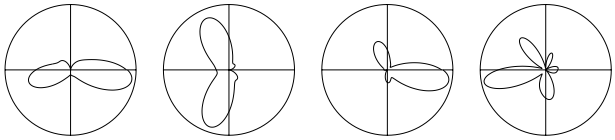

Figure 7: Parkbench image, and orientation maps computed at selected locations using the set of 7-tap steerable wedge filters with $N = 9$ (i.e., 18 basis filters). Polar plots of oriented energy maps are given for: (a) a horizontal edge; (b) a vertical edge; (c) a corner; and (d) a "T-junction".

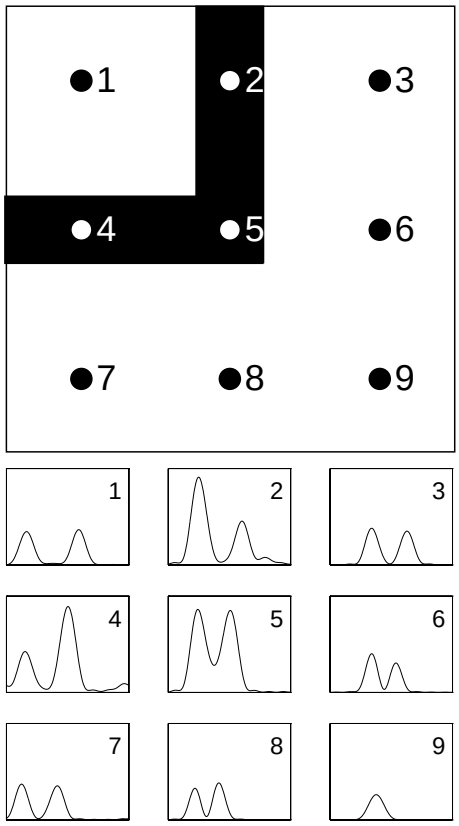

Figure 8: Left: synthetic image. Right: responses at each of nine positions. Off-center orientation maps were computed with a set of 18 steerable wedge filters (i.e., $N = 9$). Maps were computed with filters centered at each of 9 locations, separated by intervals of 2 pixels.