

The Design and Use of Steerable Filters

William T. Freeman

Edward H. Adelson †

The Media Laboratory

and †Dept. of Brain and Cognitive Sciences

Massachusetts Institute of Technology

Cambridge, Massachusetts 02139

Published in IEEE Trans. Patt. Anal. and Machine Intell.,

Vol. 13, No. 9, pp. 891-906, Sept., 1991.

Copyright ©1991, IEEE.

Abstract

Oriented filters are useful in many early vision and image processing tasks. One often needs to apply the same filter, rotated to different angles under adaptive control, or wishes to calculate the filter response at various orientations. We present an efficient architecture to synthesize filters of arbitrary orientations from linear combinations of basis filters, allowing one to adaptively “steer” a filter to any orientation, and to determine analytically the filter output as a function of orientation.

Steerable filters may be designed in quadrature pairs to allow adaptive control over phase as well as orientation. We show how to design and steer the filters, and present examples of their use in several tasks: the analysis of orientation and phase, angularly adaptive filtering, edge detection, and shape-from-shading. One can also build a self-similar steerable pyramid representation which may be used to implement a steerable “wavelet” decomposition. The same concepts can be generalized to the design of 3-D steerable filters, which should be useful in the analysis of image sequences and volumetric data.

1 Introduction

Oriented filters are used in many vision and image processing tasks, such as texture analysis, edge detection, image data compression, motion analysis, and image enhancement. In many of these tasks, it is useful to apply filters of arbitrary orientation under adaptive control, and to examine the filter output as a function of both orientation and phase. We will discuss techniques that allow synthesis of a filter at arbitrary orientation and phase, and will develop methods to analyze the filter outputs. We will also describe efficient architectures for such processing, develop flexible design methods for the filters in two and three dimensions, and apply the filters to several tasks in image analysis. Preliminary reports of this work appear in [12, 13].

One approach to finding the response of a filter at many orientations is to apply many versions of the same filter, each different from the others by some small rotation in angle. A more efficient approach is to apply a few filters corresponding to a few angles and interpolate between the responses. One then needs to know how many filters are required and how to properly interpolate between the responses. With the correct filter set and the correct interpolation rule, it is possible to determine the response of a filter of arbitrary orientation without explicitly applying that filter.

We use the term “steerable filter” to describe a class of filters in which a filter of arbitrary orientation is synthesized as a linear combination of a set of “basis filters.” We will show that both two and three-dimensional functions are steerable, and will show how many basis filters are needed to steer a given filter. We first discuss the two-dimensional case.

2 An Example

As an introductory example, consider the 2-dimensional, circularly symmetric Gaussian function, G , written in Cartesian coordinates, x and y :

$$G(x, y) = e^{-(x^2+y^2)}, \quad (1)$$

where scaling and normalization constants have been set to 1 for convenience. The directional derivative operator is steerable as is well-known [8, 12, 16, 18, 21, 22, 23, 24, 27, 34]. Let us write the n th derivative of a Gaussian in the x direction as G_n . Let $(\dots)^\theta$ represent the rotation operator, such that, for any function $f(x, y)$, $f^\theta(x, y)$ is $f(x, y)$ rotated through an angle θ about the origin. The first x derivative of a Gaussian, $G_1^{0^\circ}$, is

$$G_1^{0^\circ} = \frac{\partial}{\partial x} e^{-(x^2+y^2)} = -2xe^{-(x^2+y^2)}. \quad (2)$$

That same function, rotated 90 degrees, is:

$$G_1^{90^\circ} = \frac{\partial}{\partial y} e^{-(x^2+y^2)} = -2ye^{-(x^2+y^2)}. \quad (3)$$

These functions are shown in Fig. 1(a) and (b). It is straightforward to show that a G_1 filter at an arbitrary orientation θ can be synthesized by taking a linear combination of $G_1^{0^\circ}$ and $G_1^{90^\circ}$:

$$G_1^\theta = \cos(\theta)G_1^{0^\circ} + \sin(\theta)G_1^{90^\circ}. \quad (4)$$

Since $G_1^{0^\circ}$ and $G_1^{90^\circ}$ span the set of G_1^θ filters we call them *basis filters* for G_1^θ . The $\cos(\theta)$ and $\sin(\theta)$ terms are the corresponding *interpolation functions* for those basis filters.

Because convolution is a linear operation, we can synthesize an image filtered at an arbitrary orientation by taking linear combinations of the images filtered with $G_1^{0^\circ}$ and $G_1^{90^\circ}$. Letting $*$ represent convolution, if

$$R_1^{0^\circ} = G_1^{0^\circ} * I \quad (5)$$

$$R_1^{90^\circ} = G_1^{90^\circ} * I \quad (6)$$

then

$$R_1^\theta = \cos(\theta)R_1^{0^\circ} + \sin(\theta)R_1^{90^\circ}. \quad (7)$$

The derivative of Gaussian filters offer a simple illustration of steerability. In the next section, we generalize these results to encompass a wide variety of filters. (See [36, 41] for recent extensions of this approach.)

3 Steering Theorems in Two Dimensions

We want to find the conditions under which any function, $f(x, y)$, steers, i.e., when it can be written as a linear sum of rotated versions of itself.




Figure 1: Example of steerable filters. (a) $G_1^{0^\circ}$, first derivative with respect to x (horizontal) of a Gaussian. (b) $G_1^{90^\circ}$, which is $G_1^{0^\circ}$, rotated by 90° . From a linear combination of these two filters, one can create G_1^θ , an arbitrary rotation of the first derivative of a Gaussian. (c) $G_1^{30^\circ}$, formed by $\frac{1}{2}G_1^{0^\circ} + \frac{\sqrt{3}}{2}G_1^{90^\circ}$. The same linear combinations used to synthesize G_1^θ from the basis filters will also synthesize the response of an image to G_1^θ from the responses of the image to the basis filters: (d) Image of circular disk. (e) $G_1^{0^\circ}$ (at a smaller scale than pictured above) convolved with the disk, (d). (f) $G_1^{90^\circ}$ convolved with (d). (g) $G_1^{30^\circ}$ convolved with (d), obtained from $\frac{1}{2}$ [image e] + $\frac{\sqrt{3}}{2}$ [image f].

The steering constraint is

$$f^\theta(x, y) = \sum_{j=1}^M k_j(\theta) f^{\theta_j}(x, y). \quad (8)$$

We want to know what functions $f(x, y)$ can satisfy Eq. (8), how many terms, M , are required in the sum, and what the interpolation functions, $k_j(\theta)$, are.

We will work in polar coordinates $r = \sqrt{x^2 + y^2}$ and $\phi = \arg(x, y)$. Let f be any function which can be expanded in a Fourier series in polar angle, ϕ :

$$f(r, \phi) = \sum_{n=-N}^N a_n(r) e^{in\phi}. \quad (9)$$

In Appendix A, we prove the following theorem:

Theorem 1 *The steering condition, Eq. (8), holds for functions expandable in the form of Eq. (9) if and only if the interpolation functions $k_j(\theta)$ are solutions of:*

$$\begin{pmatrix} 1 \\ e^{i\theta} \\ \dots \\ e^{iN\theta} \end{pmatrix} = \begin{pmatrix} 1 & 1 & \dots & 1 \\ e^{i\theta_1} & e^{i\theta_2} & \dots & e^{i\theta_M} \\ \vdots & \vdots & \vdots & \vdots \\ e^{iN\theta_1} & e^{iN\theta_2} & \dots & e^{iN\theta_M} \end{pmatrix} \begin{pmatrix} k_1(\theta) \\ k_2(\theta) \\ \vdots \\ k_M(\theta) \end{pmatrix}. \quad (10)$$

If, for any n , $a_n(r) = 0$, then the corresponding (n th) row of the left hand side and of the matrix of the right hand side of Eq. (10) should be removed.

We are interested in the minimum number of basis functions which are required to steer a particular function, $f(r, \phi)$. Let T be the number of positive or negative frequencies $-N \leq n \leq N$ for which $f(r, \phi)$ has non-zero coefficients $a_n(r)$ in a Fourier decomposition in polar angle. For example, $\cos(\phi) = \frac{e^{i\phi} + e^{-i\phi}}{2}$ has $T = 2$ and $\cos(\phi) + 1 = \frac{e^{i\phi} + e^{-i\phi}}{2} + e^0$ has $T = 3$. In Appendix B we derive the minimum number of basis filters of any form which will steer $f(r, \phi)$, i.e., for which the following equation holds:

$$f^\theta(r, \phi) = \sum_{j=1}^M k_j(\theta) g_j(r, \phi), \quad (11)$$

where the $g_j(r, \phi)$ can be any set of functions. Theorem 2 gives the results:

Theorem 2 *Let T be the number of non-zero coefficients $a_n(r)$ for functions $f(r, \phi)$ expandable in the form of Eq. (9). Then the minimum number of basis functions which are sufficient to steer $f(r, \phi)$ by Eq. (11) is T . I.e., M in Eq. (11) must be $\geq T$.*

Using rotated versions of the function itself as the basis functions, as in Eq. (8), the T basis function orientations θ_j must be chosen so that the columns of the matrix in Eq. (10) are

linearly independent. In practice, for reasons of symmetry and robustness against noise, we choose basis functions spaced equally in angle between 0 and π . Note that the interpolation functions $k_j(\theta)$ do not depend on the values of the non-zero coefficients $a_n(r)$ in the Fourier angular decomposition of the filter $f(r, \phi)$.

A 1-D bandlimited function can be represented by a finite number of samples corresponding to the number of Fourier terms, which is the number of degrees of freedom. Theorems 1 and 2 show that angularly bandlimited functions behave the same way.

We illustrate the use of Theorem 1 by re-deriving the steering equation for G_1 . In polar coordinates, the first derivative of a Gaussian is

$$G_1^{0^\circ}(r, \phi) = -2re^{-r^2} \cos(\phi) = -re^{-r^2} (e^{i\phi} + e^{-i\phi}). \quad (12)$$

Since $G_1^{0^\circ}(r, \phi)$ has two non-zero coefficients in a Fourier decomposition in polar angle ϕ , by Theorem 1, two basis functions suffice to synthesize $G_1^{0^\circ}$. The interpolation functions are found from Eq. (10), with all entries but the second row removed:

$$\begin{pmatrix} e^{i\theta} \end{pmatrix} = \begin{pmatrix} e^{i\theta_1} & e^{i\theta_2} \end{pmatrix} \begin{pmatrix} k_1(\theta) \\ k_2(\theta) \end{pmatrix}. \quad (13)$$

If we pick one basis function to be oriented at $\theta_1 = 0^\circ$ and the other at $\theta_2 = 90^\circ$, then Eq. (13) gives $k_1(\theta) = \cos(\theta)$ and $k_2(\theta) = \sin(\theta)$. Thus, Theorem 1 tells us that $G_1^{0^\circ} = \sum_{j=1}^2 k_j(\theta) G_1^{0_j} = \cos(\theta) G_1^{0^\circ} + \sin(\theta) G_1^{90^\circ}$, in agreement with Eq. (4).

Figure 2 shows 1-D cross-sections of some steerable basis filters, plotted as a function of angle ϕ at a constant radius. An arbitrary translation of any one curve can be written as a linear combination of the basis curves shown on the graph (rotation of the filter corresponds to translation on these graphs). Figure 2 (a) shows the sinusoidal variation of 1-D slices of $G_1^{0^\circ}$ and $G_1^{90^\circ}$, plotted at a constant radius. In this case, the steering property is a re-statement of the fact that a linear combination of two sinusoids can synthesize a sinusoid of arbitrary phase. Figure 2(b) and (c) are 1-D cross-sections of steerable basis sets for functions with the azimuthal distribution $0.25 \cos(3\phi) + 0.75 \cos(\phi)$ and $0.25 \cos(3\phi) - 1.25 \cos(\phi)$, respectively. Since each function has non-zero Fourier coefficients for two frequencies, by Theorem 1, four basis functions suffice for steering. Because both functions contain sinusoids of the same frequencies (even though of different amplitudes), they use the same $k_j(\theta)$ interpolation coefficients.

It is convenient to have a version of Theorem 1 for functions expressed as polynomials in Cartesian coordinates x and y [12]. In Appendix C, we prove the following theorem:

Theorem 3 *Let $f(x, y) = W(r)P_N(x, y)$, where $W(r)$ is an arbitrary windowing function and $P_N(x, y)$ is an N th order polynomial in x and y , whose coefficients may depend on r . Linear combinations of $2N + 1$ basis functions are sufficient to synthesize $f(x, y) = W(r)P_N(x, y)$ rotated to any angle. Eq. (10) gives the interpolation functions, $k_j(\theta)$. If $P_N(x, y)$ contains only even [odd] order terms (terms $x^n y^m$ for $n + m$ even [odd]), then $N + 1$ basis functions are sufficient, and Eq. (10) can be modified to contain only the even [odd] numbered rows (counting from zero) of the left hand side column vector and the right hand side matrix.*




Figure 2: Three sets of steerable basis functions, plotted as a function of azimuthal angle, ϕ , at a constant radius. An arbitrary angular offset of each function (linear shift, as plotted here) can be obtained by a linear combination of the basis functions shown. (a) G_1 steerable basis set. (b) four basis functions for $0.25 \cos(3\phi) + 0.75 \cos(\phi)$; (c) four basis functions for $0.25 \cos(3\phi) - 1.25 \cos(\phi)$. The same interpolation functions apply for (b) as for (c).

Theorem 3 allows steerable filters to be designed by fitting the desired filters with polynomials times rotationally symmetric window functions, which can be simpler than using a Fourier series in polar coordinates. However, Theorem 3 is not guaranteed to find the minimum number of basis functions which can steer a filter. Representing the function in a Fourier series in angle makes explicit the minimum number of basis filters required to steer it. In a polynomial representation, the polynomial order only indicates a number of basis functions sufficient for steering. For example, consider the angularly symmetric function, $x^2 + y^2$, written in a polar representation as $r^2 e^{0\phi}$. Theorem 2 would say that only one basis function is required to steer it; Theorem 3, which uses only the polynomial order, merely says that a number of basis functions sufficient for steering is $2 + 1 = 3$.

The above theorems show that steerability is a property of a wide variety of functions, namely all functions which can be expressed as a Fourier series in angle, or in a polynomial expansion in x and y times a radially symmetric window function. Derivatives of Gaussians of all orders are steerable because each one is a polynomial (the Hermite polynomials [32]) times a radially symmetric window function.

Fig. 3 shows a general architecture for using steerable filters. (cf. Koenderink and van Doorn [22, 23, 24], who used such an architecture with derivatives of Gaussians, and Knutsson et al. [21] who used it with related filters.) The front end consists of a bank of permanent, dedicated basis filters, which always convolve the image as it comes in; their outputs are multiplied by a set of gain masks, which apply the appropriate interpolation functions at each position and time. The final summation produces the adaptively filtered image.

An alternative approach to the steerable filters presented here would be to project all rotations of a function onto a complete set of orthogonal basis functions, such as the Hermite functions, or the polynomials used in the facet model [16]. One could then steer the filter by changing the expansion coefficients. Such expansions allow flexible control over the filter, but for purposes of steering they generally require more basis functions than the minimum number given by Theorem 2. For example, $2N + 1$ basis functions are sufficient to steer any N th order polynomial, while a complete set of 2-D polynomial basis functions would require $(N + 1)(N + 2)/2$ basis functions ($n + 1$ basis functions for every order $0 \leq n \leq N$). Furthermore, a general decomposition may require extra basis functions in order to fit a rotationally symmetric component of the function, which requires no extra basis functions for steering when using rotated versions of the function itself as basis functions.

4 Designing Steerable Filters

All functions which are bandlimited in angular frequency are steerable, given enough basis filters. But in practice the most useful functions are those which require a small number of basis filters.

As an example, we will design a steerable quadrature pair based on the frequency response of the second derivative of a Gaussian, G_2 . A pair of filters is said to be in quadrature if they

Steerable Filter Architecture




Figure 3: Steerable filter system block diagram. A bank of dedicated filters process the image. Their outputs are multiplied by a set of gain maps which adaptively control the orientation of the synthesized filter.

have the same frequency response but differ in phase by 90° (i.e. are Hilbert transforms of each other [4]). Such pairs allow for analyzing spectral strength independent of phase, and allow for synthesizing filters of a given frequency response with arbitrary phase. They have application in motion, texture, and orientation analysis [1, 3, 11, 17, 19, 31, 38]. Gaussian derivatives are useful functions for image analysis [22, 23, 24, 45] and a steerable quadrature pair of them would be useful for many vision tasks.

First, we design a steerable basis set for the second derivative of a Gaussian, $f(x, y) = G_2^{0^\circ} = (4x^2 - 2)e^{-(x^2+y^2)}$. This is the product of a second order, even parity polynomial and a radially symmetric Gaussian window, so, by Theorem 3, three basis functions suffice. Equation (10) for the interpolation functions, $k_j(\theta)$, becomes

$$\begin{pmatrix} 1 \\ e^{i2\theta} \end{pmatrix} = \begin{pmatrix} 1 & 1 & 1 \\ e^{i2\theta_1} & e^{i2\theta_2} & e^{i2\theta_3} \end{pmatrix} \begin{pmatrix} k_1(\theta) \\ k_2(\theta) \\ k_3(\theta) \end{pmatrix}. \quad (14)$$

Requiring that both the real and imaginary parts of Eq. (14) agree gives a system of three equations. Solving the system, using $\theta_1 = 0^\circ$, $\theta_2 = 60^\circ$, $\theta_3 = 120^\circ$, yields

$$k_j(\theta) = \frac{1}{3}[1 + 2 \cos(2(\theta - \theta_j))], \quad (15)$$

and we have

$$G_2^\theta = k_1(\theta)G_2^{0^\circ} + k_2(\theta)G_2^{60^\circ} + k_3(\theta)G_2^{120^\circ}. \quad (16)$$

We can form an approximation to the Hilbert transform of G_2 by finding the least squares fit to a polynomial times a Gaussian. We found a satisfactory level of approximation (total error power was 1% of total signal power) using a 3rd order, odd parity polynomial, which is steerable by four basis functions. We refer to this approximation as H_2 . Its steering formula is given with that for several other polynomial orders in Appendix F.

Figures 4 (a) and (b) show 1-D slices of G_2 and H_2 . The quality of the fit of H_2 to the Hilbert transform of G_2 is fairly good, as shown by the smooth, Gaussian-like energy function $(G_2)^2 + (H_2)^2$, (c), and the closeness of the magnitudes of the Fourier spectra for each function, (d).

The seven basis functions of G_2 and H_2 are sufficient to shift G_2 arbitrarily in both phase and orientation. Those seven basis functions, and the magnitudes of their Fourier transforms, are shown in Fig. 5. Tables (1) and (2) list several quadrature pairs, based on several orders of derivatives of Gaussians and fits to their Hilbert transforms.

4.1 Designing Separable Steerable Filters

For most steerable filters, the basis filters are not all x-y separable, which can present high computational costs. For machine vision applications, we would like to have only x-y separable basis functions.

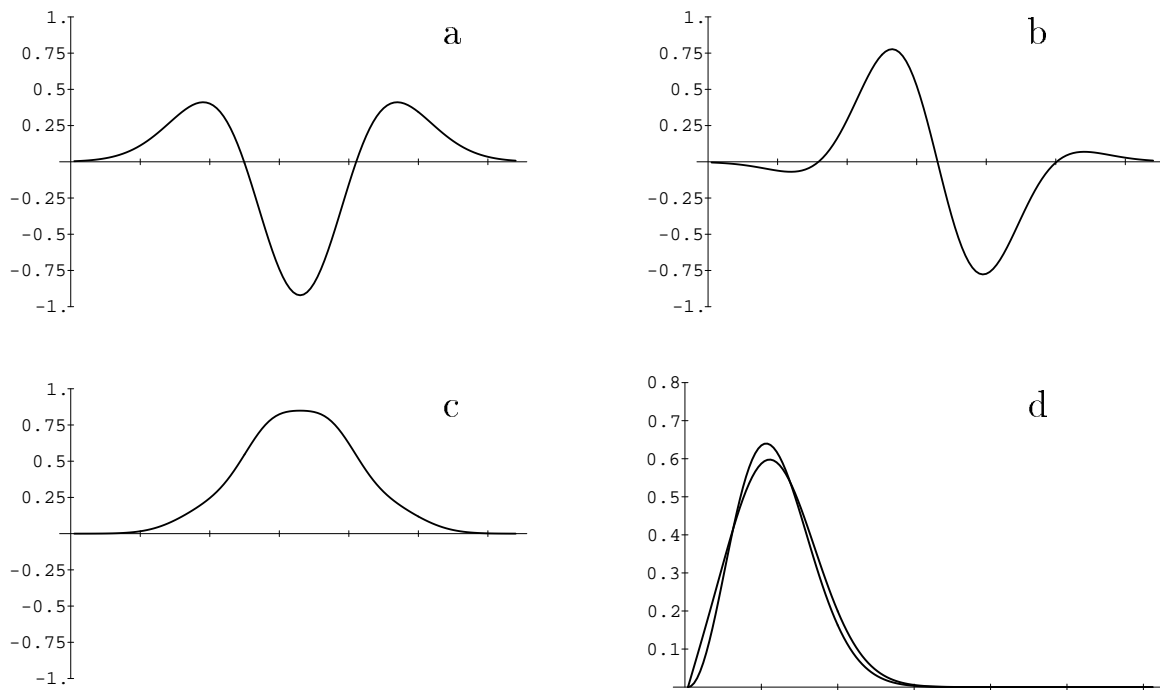




Figure 4: (a) G_2 , 2nd derivative of Gaussian (in 1 dimension). (b) H_2 , fit of 3rd order polynomial (times Gaussian) to the Hilbert transform of (a). (c) energy measure: $(G_2)^2 + (H_2)^2$. (d) magnitudes of Fourier transforms of (a) and (b).




(a) G_2 Basis Set

(b) G_2 Amplitude Spectra




(c) G_2 X-Y Separable Basis Set



(d) H_2 Basis Set

(e) H_2 Amplitude Spectra



(f) H_2 X-Y Separable Basis Set

Figure 5: G_2 and H_2 quadrature pair basis filters (rows (a) and (d)). The filters in rows (a) and (d) span the space of all rotations of their respective filters. G_2 and H_2 have the same amplitude spectra (rows (b) and (e)), but 90° shifted phase. Steerable G_2 and H_2 filters can measure local orientation direction and strength, and the phase at any orientation. Rows (c) and (f) show equivalent x-y separable basis functions which can also synthesize all rotations of G_2 and H_2 , respectively.

We first note that for all functions f which can be written as a polynomial in x and y , there is an x - y separable basis, although it may have many basis functions. Applying the rotation formula to each x and y term of the polynomial will result in a sum of products of powers of x and y , with coefficients which are functions of the rotation angle:

$$f^\theta(x, y) = \sum_l \sum_j k_{lj}(\theta) x^l y^j. \quad (17)$$

Each x and y product in the rotated polynomial can be thought of as an x - y separable basis function, with its coefficient $k_{lj}(\theta)$ the interpolation function.

In many cases, however, there exists an x - y separable basis set which contains only the minimum number of basis filters, yet spans the space of all rotations for the function of interest. Such a separable basis allows steerable filters to be applied with high computational efficiency. Rows (c) and (f) of figure 5 show x - y separable basis sets for the G_2 and H_2 filters. Tables 3, 5, 7 and 9 give the functional forms and digital filter values for x - y separable versions of the G_2 , H_2 , and G_4 and H_4 basis filters. In Appendix D we derive the steering formulas for these x - y separable functions and show how to find the separable basis functions.

4.2 Discrete Space Filters

The steering theorems have been derived for continuous functions, and one might be concerned that new difficulties would arise when one worked with discretely sampled functions. But if a continuous function is steerable, then a sampled version of it is steerable in exactly the same fashion, because the order of spatial sampling and steering are interchangeable. The weighted sum of a set of spatially sampled basis functions is equivalent to the spatial sampling of the weighted sum of continuous basis functions. So one can obtain digital steerable filters by simply sampling a continuous filter. Spatially sampled versions are given for G_2 , H_2 , G_4 and H_4 in Tables 3, 5, 7 and 9.

Filters can also be designed in the frequency domain, where one may separate the radial and angular parts of the design [19]. Conventional filter design techniques [25, 33] allow the design of a circularly symmetric 2-D filter with a desired radial response. Then, one can impose on that filter the angular variation needed to make a steerable basis set by frequency sampling [25] (if the angular response is relatively smooth). Inverse transforming the frequency sampled response gives the filter kernel.

Fig. 6 shows an example of this. The filter was designed to be part of a steerable, self-inverting pyramid image decomposition [41], described below. The constraints on the multi-scale decomposition lead to the radial frequency response shown in Fig. 6(a). We used the frequency transformation method [25] to convert the 1-D filter to a nearly angularly symmetric 2-D filter, Fig. 6 (b).

Having selected a radial frequency band, we next divided the band into four oriented subbands by imposing an angular variation of $\cos^3(\nu)$, where ν is azimuthal angle in frequency. This function has four angular frequencies (± 3 and ± 1) and so, by Theorem 1, requires four

basis functions to steer. We Fourier transformed the radially symmetric kernel, multiplied by the four desired $\cos^3(\nu - \theta_j)$ angular responses, and inverse transformed to obtain the basis filter impulse responses. Figure 6(c - f) shows the frequency amplitude responses of the resulting digital steerable filters.

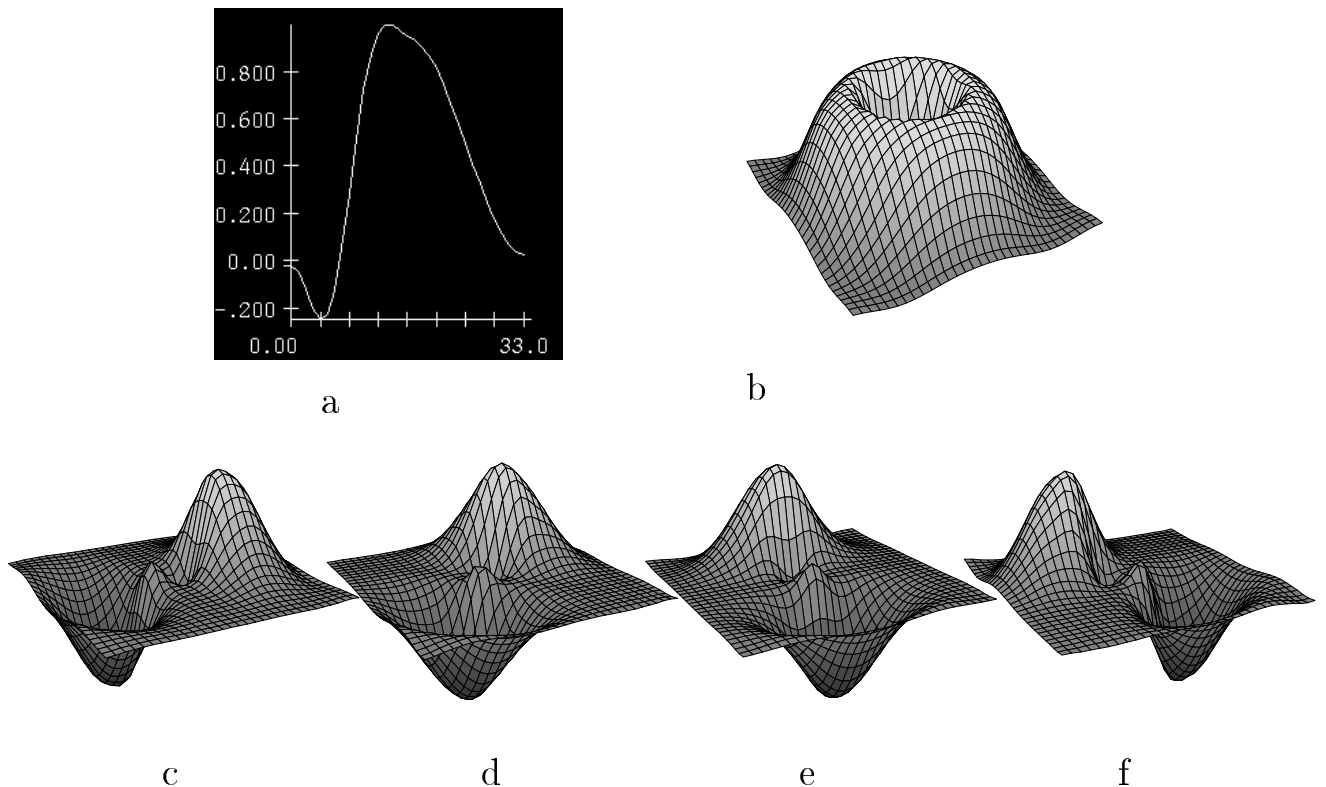


Figure 6: Frequency domain filter response plots, illustrating design procedure for steerable digital filter. (a) shows the particular desired radial frequency distribution. (b) shows the desired angularly symmetric two-dimensional frequency response, obtained through frequency transformation. (b) was multiplied by the desired $\cos^3(\nu - \theta_j)$ angular frequency responses and inverse transformed to yield the steerable basis set. The frequency responses of the resulting four steerable digital filters are shown in (c - f).

4.3 Steerable Pyramid for Multi-Scale Decomposition

We have also used the steerable filters to form a multi-scale, self-inverting pyramid decomposition [41]. Applying each filter of the decomposition once to the signal gives the transform coefficients; applying each filter a second time (with filter tap values reflected about the origin) and adding the results reconstructs a low-passed version of the image. Because all of the filters of the pyramid are bandpass, a high-pass residue image must be added back in to reconstruct the original image (as with [43]). To implement this decomposition, we designed the angular and radial components of the polar separable design so that the squares of the responses of each filter added to unity in the frequency plane.

Figure 7 shows the steerable pyramid representation. The four bandpass filters at each level of the pyramid form a steerable basis set. The pyramid basis filters were oriented at 0° , 45° , 90° , 135° , but the coefficients for any filter orientation can be found from a linear combination of the four basis filter outputs. When the basis filters are applied again at each level, the pyramid collapses back to a filtered version of the original image with near-perfect agreement. The steerable pyramid image transform allows control over orientation analysis over all scales.

The steerable pyramid is an image transform for which all of the basis functions are derived by dilation, translation, and rotation of a single function, and therefore it may be considered to be a wavelet transform [15, 26]. Most work on wavelet image decomposition has involved discrete orthogonal wavelets, in particular those known as quadrature mirror filters (QMF’s) [10, 26, 40, 42]. Pyramids made from QMF’s and other wavelets can be extremely efficient for image coding applications. Such representations are usually built with x-y separable filters on a rectangular lattice [2, 26, 44], which significantly limits the quality of orientation tuning that can be achieved. Simoncelli and Adelson [2, 39] have devised QMF pyramids based on filters placed on a hexagonal lattice; in addition to being orthogonal and self-similar, these pyramids have good orientation tuning in all bands. However, the basis functions are not steerable, so the representation is not optimal for orientation analysis. Non-orthogonal pyramids with orientation tuning have been described by [9, 14, 28, 43].

Unlike the pyramids based on QMF’s, the steerable pyramid described here is significantly overcomplete: not counting the residual image, there are $5\frac{1}{3}$ times as many coefficients in the representation as in the original image ($1\frac{1}{3}$ times over-complete, as with the Laplacian pyramid [5], but for each of 4 orientations). The overcompleteness limits its efficiency but increases its convenience for many image processing tasks. Although it is non-orthogonal, it is still self-inverting, meaning that the filters used to build the pyramid representation are the same as those used for reconstruction.

5 Applications

Steerable filters are useful for many tasks in early vision. We present four applications below—orientation and phase analysis, angularly adaptive filtering, edge detection, and shape-from-shading.

5.1 Analyzing Local Orientation

Orientation analysis is an important task in early vision [18, 19, 21, 46]. Knutsson and Granlund [19] devised an elegant method for combining the outputs of quadrature pairs to extract a measure of orientation. We describe a related method which makes optimal use of the filters designed in Section 4. We measure the orientation strength along a particular direction, θ , by the squared output of a quadrature pair of bandpass filters steered to the angle θ . We call this spectral power the “oriented energy”, $E(\theta)$.

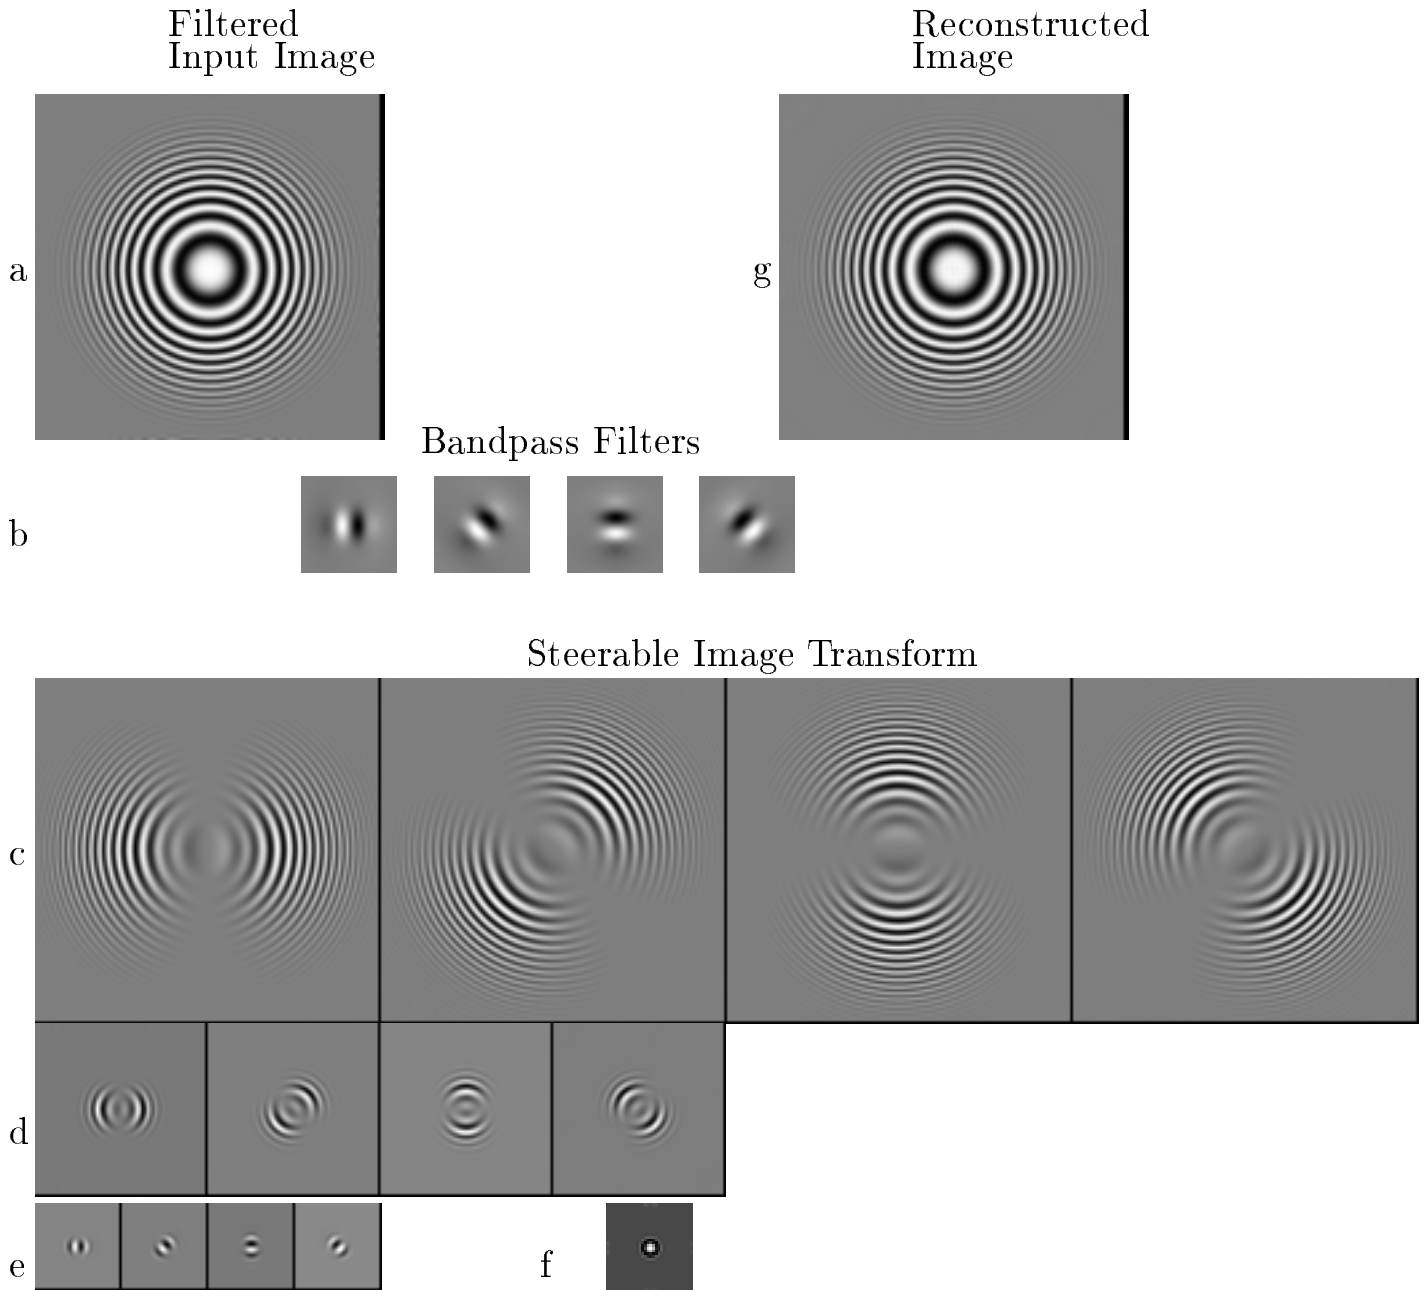


Figure 7: Steerable image transform. (a) Low-pass filtered original image. (b) Odd-phase analyzing filters, oriented at 0° , 45° , 90° , 135° . These four filters form a steerable basis set; any orientation of this filter can be written as a linear combination of the basis filters. (c) - (e) Steerable, bandpass coefficients in a multi-scale pyramid representation of (a). A linear combination of these transform coefficients will synthesize the transform coefficient for analyzing filters oriented at any angle. (f) Low-pass image. (g) Image reconstructed from the pyramid representation, showing near-perfect agreement with (a).

Using the n th derivative of a Gaussian and its Hilbert transform as our bandpass filters, we have:

$$E_n(\theta) = [G_n^\theta]^2 + [H_n^\theta]^2. \quad (18)$$

Writing G_n^θ and H_n^θ as a sum of basis filter outputs times interpolation functions, Eq. (18) simplifies to a Fourier series in angle, where only even frequencies are present, because of the squaring operation:

$$E_n(\theta) = C_1 + C_2 \cos(2\theta) + C_3 \sin(2\theta) + [\text{higher order terms } \dots]. \quad (19)$$

We use the lowest frequency term to approximate the direction, θ_d and strength, S , of the dominant orientation (the orientation which maximizes $E_n(\theta)$),

$$\theta_d = \frac{\arg[C_2, C_3]}{2} \quad (20)$$

$$S = \sqrt{C_2^2 + C_3^2}. \quad (21)$$

This approximation is exact if there is only one orientation present locally.

Figure 8 (b) shows an orientation map derived using this method, using G_2 and H_2 to measure $E_2(\theta)$. The line lengths are proportional to S , the contrast along that orientation. The measured orientations and strengths accurately reflect the oriented structures of the input image. This measurement of orientation angle was made directly from the basis filter outputs, without having to actually perform the steering operation. Table 11 lists C_2 and C_3 as functions of the basis filter outputs for x-y separable G_2 and H_2 basis filter outputs.

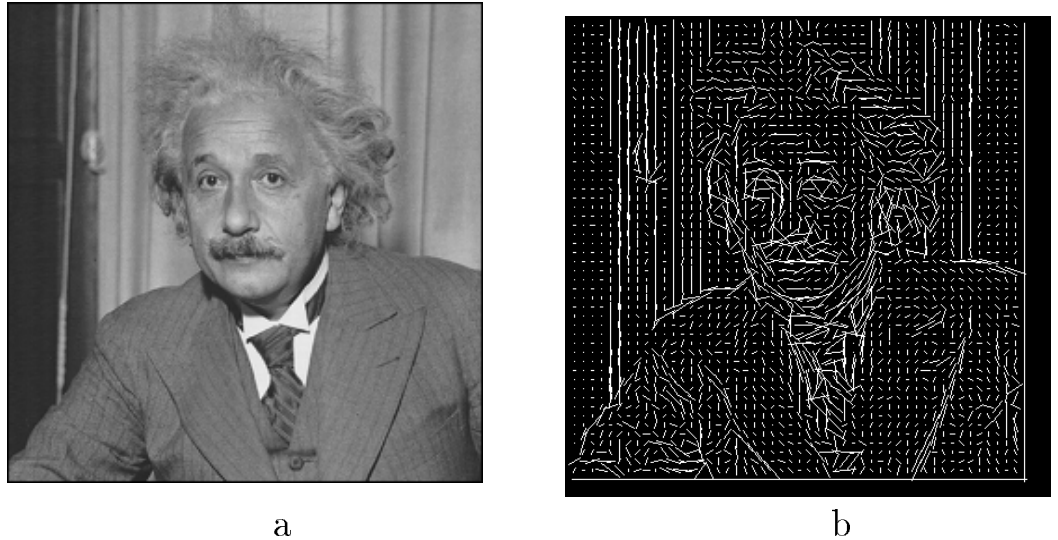


Figure 8: (a) Original image of Einstein, (b) Orientation map of (a) made using the lowest order terms in a Fourier series expansion for the oriented energy as measured with G_2 and H_2 . Table 11 gives the formulas for these terms.

5.1.1 Multiple Orientations

In regions containing corners and transparent or overlapping objects there may be more than a single orientation present at a given location. A filter such as G_2 is unable to signal the presence of two orientations at a point because of its limited angular resolution. For a higher resolution analysis of orientation, one may use a steerable filter with a narrower frequency tuning, such as the fourth derivative of a Gaussian, G_4 . This approach allows the analysis of multiple oriented structures at a single point.

The filter taps and analytical form for the steerable quadrature filter pair G_4 and H_4 are given in Appendix G. (H_4 is the least squares fit of a 5th order polynomial times a Gaussian to the Hilbert transform of G_4 .)

Figure 9 shows two test images, a vertical line, and a cross, and their oriented energy as a function of angle, measured at the center using a G_4 , H_4 quadrature pair, plotted in both Cartesian and polar coordinates. Note that the steerable filters adequately describe the multiple orientations of the cross, as seen by the floret shape.


Fig. 10 shows a test image, (a), and several measures of its oriented energy, using the G_4 , H_4 quadrature pair. Fig. 10 (b) shows the DC component of oriented energy, the angular average of Eq. (18). Because we are using a quadrature pair, the energy measure responds to both lines and edges. Fig. 10 (c) is a measure of orientation where only one orientation is allowed at each point, calculated from the lowest order Fourier terms of Eq. (18). No dominant orientation is detected at intersections of oriented structures. Fig. 10 (d) shows polar plots of the oriented energy distribution for various points in the image. Note that this measure captures the multiple orientations present at intersections and corners, shown by the florets there. These measures could all be calculated by constructing a different quadrature pair for each orientation observed; however, using the steerable filters greatly reduces the computational load.

Figure 11 shows a detail from a texture, and the corresponding polar orientation maps at every pixel in the texture image, offering a rich description of the textural details. Note that florets of one dominant orientation are separated from florets of another dominant orientation by florets where both orientations are present.


5.2 Angularly Adaptive Filtering

One can remove noise and enhance oriented structures by angularly adaptive filtering [21]. Steerable filters offer an efficient method for such processing. (Martens [27] used the steerable properties of derivatives of Gaussians for image enhancement.) Figure 12(a) shows a digital cardiac angiogram. From the outputs of the G_2 and H_2 basis filters, we found the dominant orientation direction at every point in the image, as described in Section 5.1. (In order to suppress noise, we spatially blurred the Fourier coefficients C_2 and C_3 used in Eq. (21)).


We then took the appropriate combinations of the G_2 basis filter outputs, given by Eqs. (15)




a




c




e



b



d



f

Figure 9: Test images of (a) vertical line and (b) intersecting lines. (c) and (d): Oriented energy as a function of angle at the centers of test images (a) and (b). Oriented energy was measured using the G_4 , H_4 quadrature steerable pair. (e) and (f): polar plots of (c) and (d).

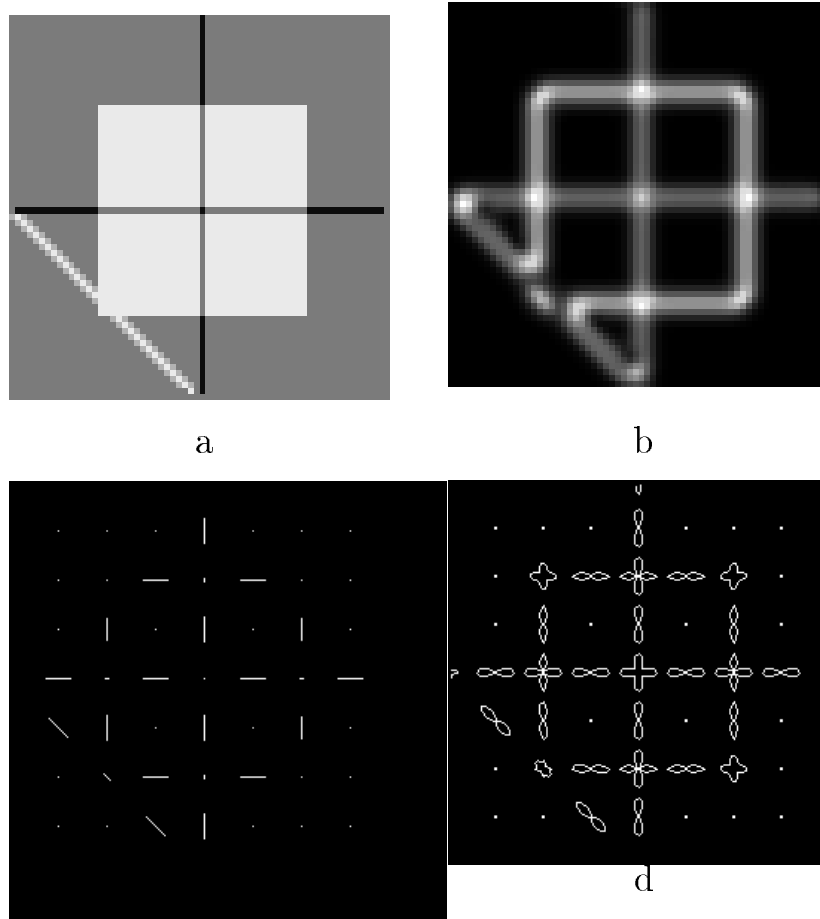





Figure 10: Measures of orientation derived from G_4 and H_4 steerable filter outputs. (a) Input image for orientation analysis (b) Angular average of oriented energy as measured by G_4 , H_4 quadrature pair. This is an oriented features detector. (c) Conventional measure of orientation: dominant orientation plotted at each point. No dominant orientation is found at the line intersection or corners. (d) Oriented energy as a function of angle, shown as a polar plot for a sampling of points in the image (a). Note the multiple orientations found at intersection points of lines or edges and at corners, shown by the florets there.




a



b



c



d

Figure 11: (a) Texture image; (b) Polar plots of oriented energy of (a) at every fourth pixel. Each plot is normalized by the average over all angles of the oriented energy. (c) Detail of (a) (zoomed and blurred); (d) Normalized polar plots showing oriented energy of (c) at every pixel.

and (16), to adaptively steer G_2 along the local direction of dominant orientation. No additional filtering was required for this step. To enhance local contrast, we divided the filtered image by a local average of its absolute value. The result, Fig. 12(b), highlights the oriented vascular structures of the angiogram. The entire process of finding the dominant orientation, steering G_2 along it, and deriving the enhanced image involved only a single pass of the image through the basis filters. The result is much less noisy than the output of an isotropic filter of the same frequency passband, Fig. 12(c), contrast enhanced in the same manner.

5.3 Contour Detection

Filters with orientation tuning are often used in the detection of lines and edges [6, 16]. One feature detector that has gained popularity is Canny’s edge operator [6], which is optimized to detect step edges; Canny’s system can also be used with different filter choices to detect features other than step edges.

A filter that is optimized for use with an edge will give spurious responses when applied to features other than edges. For example, when the Canny edge filter is applied to a line rather than an edge, it produces two extrema in its output rather than one, and each is displaced to the side of the actual line position. On the other hand, if a filter is optimized for detecting lines, it will give spurious responses with edges. Since natural images contain a mixture of lines, edges, and other contours, it is often desirable to find a contour detector that responds appropriately to the various contour types. A linear filter cannot serve this task, but a local energy measure derived from quadrature pairs can serve it quite well. Morrone et al. [31, 30] have shown that local energy measures give peak response at points of constant phase as a function of spatial frequency, and that they correspond to the points where human observers localize contours. Perona and Malik [37] have shown that energy measures are optimal with respect to a variety of edge types. We have already described the extraction of local energy measures with quadrature pairs of steerable filters. We now wish to use steerable energy measures to generate sparse image descriptions, and to compare the results with those of a system such as Canny’s.

In making this comparison we must keep in mind that Canny’s full scheme involves three stages: a filtering stage, an initial decision stage, and a complex post-processing stage which cleans up the candidate edges. The filters are merely the front end to a considerable battery of post-processing machinery. Therefore to make our comparison we removed Canny’s filtering stage and substituted the outputs of our steerable energy measures; we left the post-processing stages intact. We obtained Lisp code for the Canny edge detector from the MIT Artificial Intelligence Laboratory.

For the contour detector, we use the G_2 and H_2 quadrature steerable basis set. We first find at every position the angle of dominant orientation, θ_d , by the angle of maximum response of the steerable quadrature pair, as described in Section 5.1. We then find the squared magnitude of the quadrature pair filter response, steered everywhere in the direction of dominant orientation, $E_2(\theta_d) = [G_2^{\theta_d}]^2 + [H_2^{\theta_d}]^2$. A given point, (x_0, y_0) , is a potential contour point if $E_2(\theta_d)$ is at a




Figure 12: (a) Digital cardiac angiogram. (b) Result of filtering (a) with G_2 oriented along the local direction of dominant orientation, after local contrast enhancement (division by the image's blurred absolute value). The oriented vascular structures of (a) are enhanced. (c) Isotropic, bandpass filtering of (a), after local contrast enhancement. Note the increased noise relative to the oriented filtering results.

local maximum in the direction perpendicular to the local orientation, θ_d . The local maxima points are then thresholded with hysteresis as in the Canny method, using the values of $E_2(\theta_d)$ as the basis of thresholding, instead of the gradient magnitude.

Figure 13 (a) shows a test image consisting of a filled circle and an open square. The response of the Canny edge detector is shown in Fig. 13 (b). It correctly finds the edges of the circle, but signals double edges on either side of the lines defining the square. Figure 13 (c) shows the output using the steerable quadrature pair. The new detector responds with a single value correctly centered on both the circle and the square, giving a cleaner, sparser description of the same information.

Because the responses of G_2 and H_2 indicate the local phase, we can use them to further classify contours as edges, dark lines, or light lines. Steering G_2 and H_2 along the dominant orientation gives the phase, φ , of contour points:

$$\varphi = \arg[G_2^{\theta_d}, H_2^{\theta_d}]. \quad (22)$$

To preferentially pick-out lines or edges, we scaled the energy magnitude, $E_2(\theta_d)$ by a phase preference factor, $\Lambda(\varphi)$,

$$\Lambda(\varphi) = \begin{cases} \cos^2(\varphi - \varphi_0) & \text{if } \frac{-\pi}{2} \leq \varphi - \varphi_0 \leq \frac{\pi}{2} \\ 0 & \text{otherwise} \end{cases}, \quad (23)$$

where

$$\varphi_0 = \begin{cases} 0 & \text{for dark lines} \\ \pi & \text{for light lines} \\ \pm \frac{\pi}{2} & \text{for edges} \end{cases}. \quad (24)$$

The thresholding stage proceeds as before. Figure 13 shows the result of such processing, selecting for dark lines, (d), and edges, (e). (The blobs on the square are due to multiple orientations at a single point, and could be removed by a post-processing thinning operator.)

5.4 Shape-From-Shading Analysis

Pentland [35] has observed that in many situations, the reflectance function of a surface is approximately linear, and that under those conditions, the Fourier transform of the range image, $\hat{Z}(f, \nu)$, is related to the Fourier transform of the intensity image, $\hat{I}(f, \nu)$, by a linear transformation involving a change of phase, and scaling according to frequency,

$$\hat{Z}(f, \nu) = \frac{1}{2\pi f_{x'}} e^{-i\frac{\pi}{2}} \hat{I}(f, \nu), \quad (25)$$

where $f_{x'}$ is the x' component of frequency, and x' points toward the illuminant. Under these circumstances, shape-from-shading analysis can be performed by a filtering operation, which Pentland implemented in the Fourier domain. He also pointed out that a local approximation of the same procedure could be accomplished with Gabor-like filters.

We can describe such a shape-from-shading analysis as follows: The surface of interest, $Z(x, y)$, is considered as a sum of elementary wavelets, which we may call "bumpets," $b_j(x, y)$.

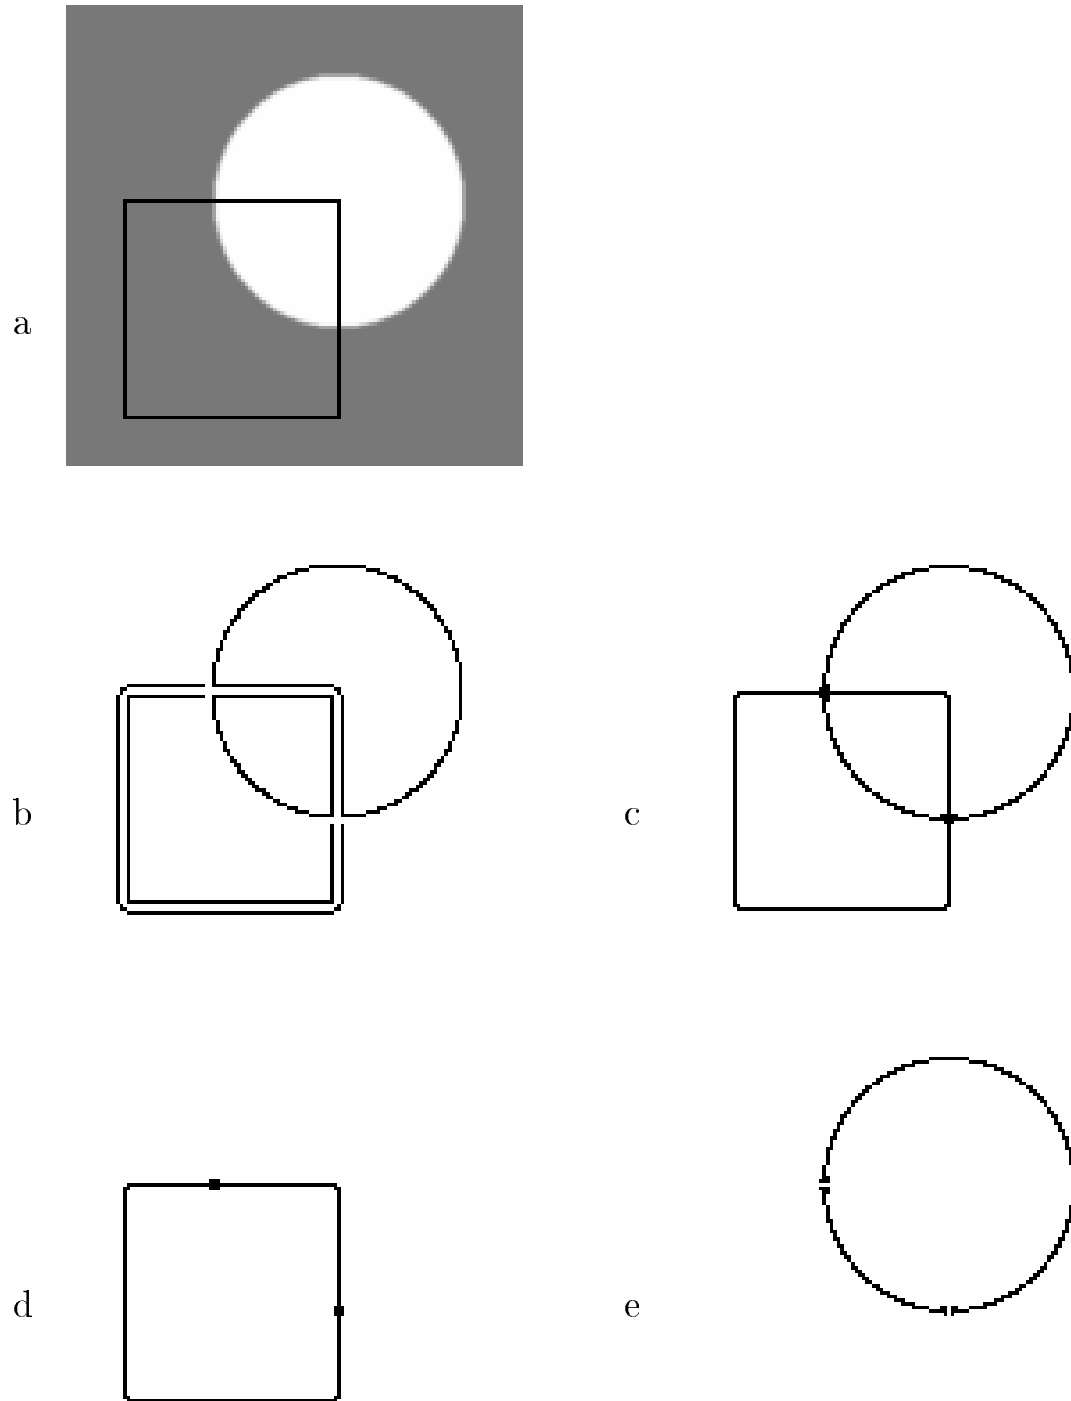


Figure 13: (a) Circle and square test image. (b) Output of Canny edge detector. The edges of the circle are accurately tracked, but the lines of the square are marked as two edges, neither at the correct position. (c) Output of steerable filter contour detector. Both edges and lines are marked as single contours, centered on the image feature. (d) Dark lines found by combining the contour detector with a phase estimator. (e) Edges found by combining the contour detector with a phase estimator.

The shading process transforms these bumplets into a new set of elementary wavelets which we may call "shadelets". Each bumplet is related to its shadelet according to the transformation of Eq. (25): $b_j(x, y) \rightarrow s_j(x, y)$. Since this shading process is linear, superposition holds and it is simple to transform back from the observed intensity image to the underlying range image. One simply decomposes the intensity image, $I(x, y)$, into the shadelet coefficients, a_j :

$$I(x, y) = \sum_j a_j s_j(x, y), \quad (26)$$

then uses these coefficients to reconstruct the surface of interest with the bumplet basis set:

$$Z(x, y) = \sum_j a_j b_j(x, y). \quad (27)$$

The steerable pyramid described above offers a convenient method for implementing this. The steerable filters of Fig. 7 are the shadelets. Their steerable quadrature pair mates, scaled according to Eq. (25), approximate the corresponding bumplets. Because the steerable pyramid transform is self-inverting, applying the shadelet filters gives the coefficients a_j . Steerability allows one to easily accomodate different lighting directions, which determines what bumplet corresponds to each shadelet. Figure 14 shows the shape-from-shading algorithm applied using the pyramid decomposition illustrated in Fig. 7. The range image successfully captures the basic characteristics of the object relief.

6 Three-Dimensional Steerable Filters

Volumetric spatial data and temporal image sequences require three-dimensional processing. As with two dimensional data, the ability to adaptively orient filters has many applications (e.g., [20]). For temporal sequences of images, orientation in space-time corresponds to velocity [1], so we expect that steerable filters will be useful in motion analysis.

In three dimensions, the steering equation we wish to solve is:

$$f^{\mathbf{R}}(x, y, z) = \sum_{j=1}^M k_j(\mathbf{R}) f^{\mathbf{R}_j}(x, y, z), \quad (28)$$

where $f^{\mathbf{R}}(x, y, z)$ is $f(x, y, z)$ after application of a 3-dimensional rotation, \mathbf{R} , and each \mathbf{R}_j identifies the orientation of the j th basis function.

In two dimensions, the number of basis filters required depended on the number of different sinusoids present in an angular Fourier decomposition of the function. In three dimensions, we can make the analogous expansion in a series of spherical harmonics. The spherical harmonics, Y_l^m , form a complete, orthonormal basis set for functions on a sphere [7, 29] and are widely used in quantum mechanics (they are the eigenfunctions of the angular momentum operator). Rotation formulas for spherical harmonics [7] show that a linear combination of the $2l + 1$ spherical harmonics of order l can synthesize an arbitrary rotation of any spherical harmonic Y_l^m .

As in the two-dimensional case, it is convenient to develop steering formulas for functions which are written as polynomials times windowing functions. Of special interest as filters are functions which have an axis of rotational symmetry. These functions, rotated by a transformation \mathbf{R} to have their axis of symmetry point along the direction cosines α , β , and γ , can be written as:

$$f^{\mathbf{R}}(x, y, z) = W(r)P_N(x'), \quad (29)$$

where $W(r)$ is any spherically symmetric function, $r = \sqrt{x^2 + y^2 + z^2}$, and $P_N(x')$ is an N th order polynomial in

$$x' = \alpha x + \beta y + \gamma z. \quad (30)$$

After substituting the functional form Eq. (29) into the three-dimensional steering equation, Eq. (28), one can derive the following steering theorem for axially symmetric functions written as polynomials times spherically symmetric window functions (see Appendix E for proof):

Theorem 4 *Given a three dimensional axially symmetric function $f(x, y, z) = W(r)P_N(x)$, where $P_N(x)$ is an even or odd symmetry N th order polynomial in x . Let α , β , and γ be the direction cosines of the axis of symmetry of $f^{\mathbf{R}}(x, y, z)$ and α_j , β_j , and γ_j be the direction cosines of the axis of symmetry of $f^{\mathbf{R}^j}(x, y, z)$. Then the steering equation,*

$$f^{\mathbf{R}}(x, y, z) = \sum_{j=1}^M k_j(\alpha, \beta, \gamma) f^{\mathbf{R}^j}(x, y, z), \quad (31)$$

holds if and only if

(a) $M \geq (N + 1)(N + 2)/2$ and

(b) the $k_j(\alpha, \beta, \gamma)$ satisfy

$$\begin{pmatrix} \alpha^N \\ \alpha^{N-1}\beta \\ \alpha^{N-1}\gamma \\ \alpha^{N-2}\beta^2 \\ \vdots \\ \gamma^N \end{pmatrix} = \begin{pmatrix} \alpha_1^N & \alpha_2^N & \dots & \alpha_M^N \\ \alpha_1^{N-1}\beta_1 & \alpha_2^{N-1}\beta_2 & \dots & \alpha_M^{N-1}\beta_M \\ \alpha_1^{N-1}\gamma_1 & \alpha_2^{N-1}\gamma_2 & \dots & \alpha_M^{N-1}\gamma_M \\ \alpha_1^{N-2}\beta_1^2 & \alpha_2^{N-2}\beta_2^2 & \dots & \alpha_M^{N-2}\beta_M^2 \\ \vdots & \vdots & \vdots & \vdots \\ \gamma_1^N & \gamma_2^N & \dots & \gamma_M^N \end{pmatrix} \begin{pmatrix} k_1(\alpha, \beta, \gamma) \\ k_2(\alpha, \beta, \gamma) \\ k_3(\alpha, \beta, \gamma) \\ \vdots \\ k_M(\alpha, \beta, \gamma) \end{pmatrix}. \quad (32)$$

By adding the number of basis functions sufficient for steering even and odd symmetry polynomials, it follows from Theorem 4 that $(N + 1)^2$ basis functions are sufficient for steering functions $f(x, y, z) = W(r)P_N(x)$, where $P_N(x)$ is a general N th order polynomial. Theorem 4 permits one to design and steer arbitrary axially symmetric 3-dimensional filters. For example, one can design 3-dimensional versions of the second derivative of a Gaussian, G_2 and a third order polynomial least squares fit to its Hilbert transform, H_2 . Since G_2 can be written as a second order, even parity polynomial times a Gaussian window function, by Theorem 4, six basis functions suffice for steering it in three-dimensions. Ten basis functions will steer H_2 .

Three-dimensional filtering can be computationally intensive. For non-separable kernels, the computational cost grows as the cube of the kernel size. For separable kernels, however, the cost grows only linearly with kernel size. Thus, it is important to develop x-y-z separable steerable filters. The spherically symmetric Gaussian function can be written as a product of functions of x , y , and z . If the weighting function, $W(r)$, is a Gaussian, then functions $f^{\mathbf{R}}(x, y, z)$ of the form of Eq. (29) can be written as a sum of separable basis functions by substituting Eq. (30) for x' in Eq. (29).

7 Summary

Steerable filters can be used for a variety of operations involving oriented filters. The oriented filter, rotated to an arbitrary angle, is formed as a linear combination of basis filters. Once the basis filter responses are known, the response of the filter steered (rotated) to an arbitrary angle, can easily be found. A similar technique can be used to control the phase of the filters. We have shown that most filters can be steered in this manner, given enough basis filters, and have described how to determine the minimum number of basis functions required, and how to interpolate between them in angle.

Steerable filters may be applied to many problems in early vision and image analysis. Because the synthesis of the rotated filter is analytic and exact, steerable filters offer advantages for image analysis over *ad hoc* methods of combining oriented filters at different orientations.

We have designed steerable quadrature pair filters, and have used them to analyze orientation, adaptively filter to enhance oriented structures, and detect contours. These processing schemes require no additional convolution after the initial pass through the basis filters. The contour detector utilizes quadrature pairs to mark both lines and edges with a single response and can be used to further categorize the contours as either dark lines, light lines, or edges.

One can also build a self-similar steerable pyramid representation which may be considered to be a steerable wavelet transform, allowing the analysis and manipulation of oriented structures at all scales. The steerable pyramid can be used for local linear shape-from-shading analysis; the steering property accommodates lighting orientation.

Steering generalizes to three dimensions, and we give formulas for steering arbitrary rotationally symmetric functions. Basis functions can be separable in x-y-z, giving a tremendous computational advantage for large oriented filters. These 3-D filters should find application in motion analysis and the analysis of volumetric data.


8 Acknowledgements

The authors had helpful conversations with Jim Bergen, Dave Heeger, Jan Koenderink, Jeff Lubin, Alex Pentland, Rosalind Picard, and Eero Simoncelli. We thank Paul Granfors of G. E. Medical Systems (Milwaukee) for providing the digital cardiac angiogram of Fig. 12(a). We

thank the anonymous referees for their comments, which improved the paper. The support of DARPA/RADC grant F30602-89-C-0022 is gratefully acknowledged. The opinions expressed are those of the authors and do not necessarily represent those of the sponsor.



(a) Original Image



(b) Low-resolution plot of range data

(c) Full-resolution range map



Figure 14: (a) Image input for linear shape-from-shading analysis using steerable image transform. (Steering was used to accommodate different light directions). (b) Resulting range map, displayed as a low-resolution 3-D plot. (c) Same range map, with pixel intensity showing surface height. This simple mechanism, using biologically plausible filters, correctly derived the image surface characteristics.

References

- [1] E. H. Adelson and J. R. Bergen. Spatiotemporal energy models for the perception of motion. *J. Opt. Soc. Am. A*, 2(2):284–299, 1985.
- [2] E. H. Adelson, E. Simoncelli, and R. Hingorani. Orthogonal pyramid transforms for image coding. In *Proc. SPIE — Vis. Comm. and Image Proc. II*, pages 50–58, Cambridge, MA, 1987.
- [3] A. C. Bovik. Multichannel texture analysis using localized spatial filters. *IEEE Pat. Anal. Mach. Intell.*, 12(1):55–73, 1990.
- [4] R. N. Bracewell. *The Fourier Transform and its Applications*. McGraw-Hill, 1978.
- [5] P. J. Burt and E. H. Adelson. The Laplacian pyramid as a compact image code. *IEEE Trans. Comm.*, 31(4):532–540, 1983.
- [6] J. F. Canny. A computational approach to edge detection. *IEEE Pat. Anal. Mach. Intell.*, 8(6):679–698, 1986.
- [7] R. Courant and D. Hilbert. *Methods of Mathematical Physics*, volume 1. Interscience Publishers, Inc., 1953.
- [8] P. Danielsson and O. Seger. Rotation invariance in gradient and higher order derivative detectors. *Comp. Vis., Graphics, Image Proc.*, 49:198–221, 1990.
- [9] J. G. Daugman. Complete discrete 2-d Gabor transforms by neural networks for image analysis and compression. *IEEE Trans. Acoust., Speech, Signal Proc.*, 36(7):1169–1179, 1988.
- [10] D. Esteban and C. Galand. Application of quadrature mirror filters to split band voice coding schemes. In *Proc. ICASSP*, pages 191 – 195, 1977.
- [11] D. Fleet and A. Jepson. Computation of normal velocity from local phase information. In *Proc. IEEE CVPR*, pages 379–386, San Diego, CA, 1989.
- [12] W. T. Freeman and E. H. Adelson. Steerable filters. In *Topical Meeting on Image Understanding and Machine Vision*. Optical Society of America, June. 1989 Technical Digest Series Volume 14.
- [13] W. T. Freeman and E. H. Adelson. Steerable filters for early vision, image analysis, and wavelet decomposition. In *Proc. 3rd Intl. Conf. Computer Vision*, Osaka, Japan, 1990.
- [14] G. H. Granlund. In search of a general picture processing operator. *Comp. Graphics, Image Proc.*, 8:155–173, 1978.
- [15] A. Grossmann and J. Morlet. Decomposition of Hardy functions into square integrable wavelets of constant shape. *SIAM J. Math.*, 15:723–736, 1984.

- [16] R. M. Haralick. The digital step edge from zero crossings of second directional derivatives. *IEEE Pat. Anal. Mach. Intell.*, 6(1):58–68, 1984.
- [17] D. J. Heeger. Optical flow using spatiotemporal filters. *Intl. J. Comp. Vis.*, 1(4):279–302, 1988.
- [18] M. Kass and A. P. Witkin. Analyzing oriented patterns. In *Proc. Ninth IJCAI*, pages 944–952, Los Angeles, CA, August 1985.
- [19] H. Knutsson and G. H. Granlund. Texture analysis using two-dimensional quadrature filters. In *IEEE Computer Society Workshop on Computer Architecture for Pattern Analysis and Image Database Management*, pages 206–213, 1983.
- [20] H. Knutsson, L. Haglund, and G. H. Granlund. Tensor field controlled image sequence enhancement. In *SSAB Symposium on Image Analysis*, Linköping, Sweden, March 1990.
- [21] H. Knutsson, R. Wilson, and G. H. Granlund. Anisotropic nonstationary image estimation and its applications: Part 1 – Restoration of noisy images. *IEEE Trans. Comm.*, 31(3):388–397, 1983.
- [22] J. J. Koenderink. Design for a sensorium. In W. von Seelen, B. Shaw, and U. M. Leinhos, editors, *Organization of Neural Networks*, pages 185–207. Verlagsgesellschaft mbH, 1988.
- [23] J. J. Koenderink. Operational significance of receptive field assemblies. *Biol. Cybern.*, 58:163–171, 1988.
- [24] J. J. Koenderink and A. J. van Doorn. Representation of local geometry in the visual system. *Biol. Cybern.*, 55:367–375, 1987.
- [25] J. Lim. *Two-Dimensional Signal and Image Processing*. Prentice Hall, Englewood Cliffs, New Jersey, 1990.
- [26] S. G. Mallat. A theory for multi-resolution signal decomposition: the wavelet representation. *IEEE Pat. Anal. Mach. Intell.*, 11(47):674–693, 1989.
- [27] J. B. Martens. Applications of polynomial transforms in image coding and computer vision. In *Proceedings of SPIE*, volume 1199, pages 1279–1290, Cambridge, MA, 1989.
- [28] J. B. Martens. The Hermite transform – theory. *IEEE Trans. Acoust., Speech, Signal Proc.*, 38(9):1595–1606, 1990.
- [29] E. Merzbacher. *Quantum Mechanics*. John Wiley and Sons, 1970.
- [30] M. C. Morrone and D. C. Burr. Feature detection in human vision: a phase-dependent energy model. *Proc. R. Soc. Lond. B*, 235:221–245, 1988.
- [31] M. C. Morrone and R. A. Owens. Feature detection from local energy. *Pattern Recognition Letters*, 6:303–313, 1987.
- [32] P. M. Morse and H. Feshbach. *Methods of Theoretical Physics*, volume 1. McGraw-Hill, 1953.

- [33] A. V. Oppenheim and R. W. Schaffer. *Digital Signal Processing*. Prentice-Hall, Englewood Cliffs, NJ, 1975.
- [34] A. P. Pentland. Local shading analysis. *IEEE Pat. Anal. Mach. Intell.*, 6(2):170–187, 1984.
- [35] A. P. Pentland. Linear shape from shading. *Intl. J. Comp. Vis.*, 1(4):153–162, 1990.
- [36] P. Perona. Finite representation of deformable functions. Technical Report 90-034, International Computer Science Institute, 1947 Center St., Suite 600, Berkeley, CA 94704-1105, 1990.
- [37] P. Perona and J. Malik. Detecting and localizing edges composed of steps, peaks and roofs. In *Proc. 3rd Intl. Conf. Computer Vision*, Osaka, Japan, 1990.
- [38] T. Sanger. Stereo disparity computation using Gabor filters. *Biol. Cybern.*, 59:405–418, 1988.
- [39] E. P. Simoncelli and E. H. Adelson. Non-separable extensions of quadrature mirror filters to multiple dimensions. *Proc. IEEE*, 78(4):652–664, 1990.
- [40] E. P. Simoncelli and E. H. Adelson. Subband transforms. In J. W. Woods, editor, *Subband Image Coding*, chapter 4. Kluwer Academic Publishers, Norwell, MA, 1990.
- [41] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger. Wavelet image transforms with continuous parameterization. Vision and Modeling Technical Report 161, The Media Lab, MIT, 20 Ames St., Cambridge, MA 02139, 1991.
- [42] M. Vetterli. Multidimensional subband coding: some theory and algorithms. *Signal Processing*, 6(2):97–112, 1984.
- [43] A. B. Watson. The cortex transform: Rapid computation of simulated neural images. *Comp. Vis., Graphics, Image Proc.*, 39:311–327, 1987.
- [44] J. W. Woods and S. D. O’Neil. Subband coding of images. *IEEE Trans. Acoust., Speech, Signal Proc.*, 34(5):1278–1288, 1986.
- [45] R. A. Young. Simulation of human retinal function with the Gaussian derivative model. In *Proc. IEEE Computer Society Conf. on Computer Vision and Pattern Recognition*, pages 564–569, 1986.
- [46] S. W. Zucker. Early orientation selection: Tangent fields and the dimensionality of their support. *Comp. Vis., Graphics, Image Proc.*, 32:74–103, 1985.

A Proof of Theorem 1

Substituting the expansion for $f(r, \phi)$, Eq. (9), into the steering constraint, Eq. (8), and projecting both sides onto the complex exponential $e^{im\phi}$, for $-N \leq m \leq N$, gives a set of simultaneous linear equations:

$$a_m(r)e^{im\theta} = \sum_{j=1}^M k_j(\theta)a_m(r)e^{im\theta_j} \quad , \quad -N \leq m \leq N. \quad (33)$$

If $a_m(r) = 0$ for some m , then we can remove that constraint from the set, otherwise, divide both sides by $a_m(r)$. The constraints above are the same for $-m$ as for m , so without loss of generality we can consider only positive frequencies $0 \leq m \leq N$ in Eq. (33). This gives Eq. (10) of Theorem 1. One can also start from Eq. (33) and derive the steering condition from it, showing that the conclusion of the theorem holds if and only if the premises hold.

B Proof of Theorem 2

We want to find the minimum number of basis filters which can span all rotations of a given filter, $f(r, \phi)$. Let $g_j(r, \phi)$, $0 \leq j \leq M$ be any set of M basis functions. We want to find the minimum number M for which Eq. (11) holds. Using the expansion for $f(r, \phi)$, Eq. (9), projecting both sides of Eq. (11) onto $e^{im\phi}$, and dividing by $a_m(r) \neq 0$ gives the following constraints:

$$e^{im\theta} = \sum_{j=1}^M k_j(\theta)c_{jm}(r) \quad , \quad 0 \leq m \leq N, \quad (34)$$

where $c_{jm}(r)$ is $a_m^{-1}(r)$ times the projection of the basis function $g_j(r, \phi)$ onto $e^{-im\phi}$.

Substituting Eq. (34) into the orthonormality relation for complex exponentials, we can write the following matrix equation,

$$\mathbf{I} = \mathbf{C}\mathbf{K}\mathbf{C}^t, \quad (35)$$

where \mathbf{I} is a T by T identity matrix; \mathbf{C} is a T by M matrix having elements $c_{lj}(r)$; \mathbf{K} is M by M , with elements $\frac{1}{2\pi} \int_{-\pi}^{\pi} k_l(\theta)k_j(\theta) d\theta$; \mathbf{C}^t is \mathbf{C} transpose; and T is the number of positive or negative frequencies m for which $a_m(r) \neq 0$. Since \mathbf{I} has rank T , then \mathbf{K} must have rank at least T , and so for steering we must have $M \geq T$, as desired.

C Proof of Theorem 3

Consider the term $x^k y^{n-k}$, where $0 \leq k \leq n$. This can be re-written in polar coordinates using $x = r \cos(\phi)$ and $y = r \sin(\phi)$:

$$x^k y^{n-k} = r^n \cos(\phi)^k \sin(\phi)^{n-k}. \quad (36)$$

It can be shown that this product of powers of sines and cosines, written as a Fourier series, can contain only the frequencies $n\phi$, $(n-2)\phi$, \dots , $-(n-2)\phi$, $-n\phi$. Thus, an N th order polynomial containing only even order terms could only have even angular frequencies m for $-N \leq m \leq N$. By Theorem 1, it would require at most $N+1$ basis functions for steering. Similarly, $N+1$ basis functions suffice for a polynomial with only odd parity terms. A general N th order polynomial could contain all angular frequencies of absolute value less than or equal to N and would need at most $2N+1$ basis functions to steer.

D Basis Functions Separable in x and y

We show how to find the steering formulas and x-y separable basis functions for some polynomial filters. We consider only the case of even or odd parity filters $f^\theta(x, y)$ which can be written as

$$f^\theta(x, y) = G(r)Q_N(x'). \quad (37)$$

where $G(r)$ is a Gaussian function (and therefore x-y separable) and $Q_N(x')$ is an N th order polynomial in

$$x' = x \cos(\theta) - y \sin(\theta). \quad (38)$$

By Theorem 3, $N+1$ functions can form a basis set for $f^\theta(x, y)$. We assume that a basis set of $N+1$ x-y separable filters exists (that is not true for all functions). Then there will be some set of separable basis functions $R_j(x)S_j(y)$ for which

$$f^\theta(x, y) = G(r) \sum_{j=0}^N k_j(\theta) R_j(x) S_j(y). \quad (39)$$

We can find the interpolation functions, $k_j(\theta)$, by equating the highest order products of x and y in Eq. (37) with those of Eq. (39), ie., equating the coefficients of $x^{(N-j)}y^j$ for $0 \leq j \leq N$. Substituting Eq. (38) into Eq. (37), the $(x')^N$ term in $f^\theta(x, y)$ will give rise to $N+1$ different products of x and y of order N , since

$$(x')^N = \sum_{j=0}^N (-1)^j \binom{N}{j} \cos^{(N-j)}(\theta) \sin^j(\theta) [x^{(N-j)}y^j]. \quad (40)$$

Each basis function $R_j(x)S_j(y)$ can contribute only one product of powers of x and y of order N (otherwise $R_j(x)S_j(y)$ would be a polynomial in x and y of order higher than N). So we must have

$$R_j(x)S_j(y) = c(x^{(N-j)} + \dots)(y^j + \dots), \quad (41)$$

where c is a constant. Therefore Eq. (39) shows that the coefficient of the highest order terms, $x^{(N-j)}y^j$, in $f^\theta(x, y)$ is $k_j(\theta)$. (The lower order terms can appear in more than one separable basis function, so their coefficients will be a sum of different $k_j(\theta)$.) Using Eq. (40) in Eq. (37) gives those same coefficients in terms of sines and cosines. Equating the two gives

$$k_j(\theta) = (-1)^j \binom{N}{j} \cos^{(N-j)}(\theta) \sin^j(\theta). \quad (42)$$

To find the separable basis functions $R_j(x)S_j(y)$ from the original filter $f(x, y)$, we note that from the steering equation for the separable basis functions, Eq. (39), we have

$$\begin{pmatrix} f^{\theta_1}(x, y) \\ f^{\theta_2}(x, y) \\ \dots \\ f^{\theta_N}(x, y) \end{pmatrix} = G(r) \begin{pmatrix} k_1(\theta_1) & k_2(\theta_1) & \dots & k_N(\theta_1) \\ k_1(\theta_2) & k_2(\theta_2) & \dots & k_N(\theta_2) \\ \vdots & \vdots & \vdots & \vdots \\ k_1(\theta_N) & k_2(\theta_N) & \dots & k_N(\theta_N) \end{pmatrix} \begin{pmatrix} R_1(x)S_1(y) \\ R_2(x)S_2(y) \\ \vdots \\ R_N(x)S_N(y) \end{pmatrix}. \quad (43)$$

The $R_j(x)S_j(y)$ can be written as a linear combination of the $f^{\theta_j}(x, y)$ by inverting the matrix of $k(\theta)$'s on the right-hand side of Eq. (43).

E Proof of Theorem 4

First, equating only the highest order terms of Eq. (31) (after dividing both sides by $W(r)$), we have

$$(\alpha x + \beta y + \gamma z)^N = \sum_{j=1}^M k_j (\alpha_j x + \beta_j y + \gamma_j z)^N. \quad (44)$$

Expanding the N th power of the sums on both sides, and equating like powers of x , y , and z gives the constraints of Eq. (32).

The constraint equations resulting from any lower order polynomial terms of $f^{\mathbf{R}}$ and $f^{\mathbf{R}_j}$ in Eq. (31) will turn out to be linearly dependent on the constraints of Eq. (32). This can be seen as follows. Consider the coefficients of $x^p y^q z^r$ in Eq. (31), for $p + q + r < N$. Dividing out common factors, we have

$$\alpha^p \beta^q \gamma^r = \sum_{j=1}^M k_j \alpha_j^p \beta_j^q \gamma_j^r. \quad (45)$$

Because $P_N(x')$ is assumed to have even or odd symmetry, then powers of x' can differ only by even integers. Consider coefficients resulting from terms in Eq. (31) of order $p + q + r + 2$. There will be at least the following three equations:

$$\alpha^{p+2} \beta^q \gamma^r = \sum_{j=1}^M k_j \alpha_j^{p+2} \beta_j^q \gamma_j^r \quad (46)$$

$$\alpha^p \beta^{q+2} \gamma^r = \sum_{j=1}^M k_j \alpha_j^p \beta_j^{q+2} \gamma_j^r \quad (47)$$

$$\alpha^p \beta^q \gamma^{r+2} = \sum_{j=1}^M k_j \alpha_j^p \beta_j^q \gamma_j^{r+2}. \quad (48)$$

Now utilize the fact that the sum of the squares of direction cosines is one: substituting $\alpha^2 = 1 - \beta^2 - \gamma^2$ and $\alpha_j^2 = 1 - \beta_j^2 - \gamma_j^2$ into Eq. (46), and adding Eqs. (47) and (48) to it gives Eq. (45). Thus, every constraint equation resulting from terms of polynomial order n is linearly dependent on the constraint equations from the polynomial order $n + 2$. So if the constraints

of the highest order terms, Eq. (32), are satisfied, and the polynomial $P_N(x')$ contains terms of only even or odd order, then Eq. (31) holds. Because there are $(N+1)(N+2)/2$ constraint equations in Eq. (32), we must have $M \geq (N+1)(N+2)/2$. One can proceed from Eq. (32) back to Eq. (31), and so the theorem conclusions hold if and only if the premises hold.

F Formulas for Steering Even or Odd Parity Polynomials

Polynomial Order	Steering Equation		
1	$k_j(\theta)$	=	$\frac{1}{2}[2 \cos(\theta - \theta_j)]$
2	$k_j(\theta)$	=	$\frac{1}{3}[1 + 2 \cos(2(\theta - \theta_j))]$
3	$k_j(\theta)$	=	$\frac{1}{4}[2 \cos(\theta - \theta_j) + 2 \cos(3(\theta - \theta_j))]$
4	$k_j(\theta)$	=	$\frac{1}{5}[1 + 2 \cos(2(\theta - \theta_j)) + 2 \cos(4(\theta - \theta_j))]$
5	$k_j(\theta)$	=	$\frac{1}{6}[2 \cos(\theta - \theta_j) + 2 \cos(3(\theta - \theta_j)) + 2 \cos(5(\theta - \theta_j))]$

Table 1: Interpolation functions $k_j(\theta)$ in Eq. (8) needed to synthesize $f^\theta(x, y)$ from the basis functions $f^{\theta_j}(x, y)$, where $f(x, y)$ is a polynomial in x and y (times any window function $W(r)$) with only even or odd parity terms. The orientations of the $n + 1$ basis functions were assumed to be evenly spaced between 0 and π , i.e. $\theta_j = j\pi/(n + 1)$, where $j = 0, 1, \dots, n$. Under those conditions, the pattern apparent in the terms above continues to all polynomial orders.

G Steerable Quadrature Filter Pairs




Figure 15: Three steerable quadrature filter pairs, listed in Table 2.

G_2	=	$0.9213(2x^2 - 1)e^{-(x^2+y^2)}$
H_2	=	$(-2.205x + 0.9780x^3)e^{-(x^2+y^2)}$
G_3	=	$(2.472x - 1.648x^3)e^{-(x^2+y^2)}$
H_3	=	$(-0.9454 + 2.959x^2 - 0.6582x^4)e^{-(x^2+y^2)}$
G_4	=	$(0.9344 - 3.738x^2 + 1.246x^4)e^{-(x^2+y^2)}$
H_4	=	$(2.858x - 2.982x^3 + 0.3975x^5)e^{-(x^2+y^2)}$

Table 2: Several Gaussian derivatives and polynomial fits to their Hilbert transforms (transforms and derivatives taken along x axis). The listed functions are normalized so that the integral over all space of their square equals one. To steer each of these, use Eq. (8) and the appropriate $k_j(\theta)$ from Table 1.

H X-Y Separable, Steerable Quadrature Pair Basis Filters




Figure 16: X-Y separable basis filters for G_2 , listed in Tables 3 and 4.

$G_{2a} = 0.9213(2x^2 - 1)e^{-(x^2+y^2)}$	$k_a(\theta) = \cos^2(\theta)$
$G_{2b} = 1.843xye^{-(x^2+y^2)}$	$k_b(\theta) = -2 \cos(\theta) \sin(\theta)$
$G_{2c} = 0.9213(2y^2 - 1)e^{-(x^2+y^2)}$	$k_c(\theta) = \sin^2(\theta)$

Table 3: X-Y separable basis set and interpolation functions for second derivative of Gaussian. To create a second derivative of a Gaussian rotated along to an angle θ , use: $G_2^\theta = (k_a(\theta) G_{2a} + k_b(\theta) G_{2b} + k_c(\theta) G_{2c})$. The minus sign in $k_b(\theta)$ selects the direction of positive θ to be counter-clockwise.

tap #	f1	f2	f3	G_2 basis filter	filter in x	filter in y
0	-0.9213	1.0	0.0			
1	-0.0601	0.6383	0.5806	G_{2a}	f1	f2
2	0.3964	0.1660	0.3020	G_{2b}	f3	f3
3	0.1148	0.0176	0.0480	G_{2c}	f2	f1
4	0.0094	0.0008	0.0028			

Table 4: 9-tap filters for x - y separable basis set for G_2 . Filters f1 and f2 have even symmetry; f3 has odd symmetry. These filters were taken from Table 3, with a sample spacing of 0.67. Use the interpolation functions of Table 3.




Figure 17: X-Y separable basis filters for H_2 , listed in Tables 5 and 6.

H_{2a}	$= 0.9780(-2.254x + x^3)e^{-(x^2+y^2)}$	$k_a(\theta) = \cos^3(\theta)$
H_{2b}	$= 0.9780(-.7515 + x^2)(y)e^{-(x^2+y^2)}$	$k_b(\theta) = -3 \cos^2(\theta) \sin(\theta)$
H_{2c}	$= 0.9780(-.7515 + y^2)(x)e^{-(x^2+y^2)}$	$k_c(\theta) = 3 \cos(\theta) \sin^2(\theta)$
H_{2d}	$= 0.9780(-2.254y + y^3)e^{-(x^2+y^2)}$	$k_d(\theta) = -\sin^3(\theta)$

Table 5: H_2 basis set: x - y separable basis set and interpolation functions for fit to Hilbert transform of second derivative of Gaussian. To synthesize a filter oriented along direction θ , use: $H_2^\theta = (k_a(\theta)H_{2a} + k_b(\theta)H_{2b} + k_c(\theta)H_{2c} + k_d(\theta)H_{2d})$. The distance between filter taps should be the same as that used with the quadrature pair derivative of Gaussian filter.

tap #	f1	f2	f3	f4	H_2 basis filter	filter in x	filter in y
0	0.0	1.0	0.0	-0.7349	H_{2a}	f1	f2
1	-0.7551	0.6383	0.4277	-0.1889	H_{2b}	f4	f3
2	-0.0998	0.1660	0.2225	0.1695	H_{2c}	f3	f4
3	0.0618	0.0176	0.0354	0.0566	H_{2d}	f2	f1
4	0.0098	0.0008	0.0020	0.0048			

Table 6: 9-tap filters for x - y separable basis set for H_2 . Filters for which tap 0 is 0.0 have odd symmetry about tap 0; the others have even symmetry. These filters were taken from Table 5, with a sample spacing of 0.67. Use the interpolation functions of Table 5.




Figure 18: X-Y separable basis filters for G_4 , listed in Tables 7 and 8.

G_{4a}	$= 1.246(0.75 - 3x^2 + x^4)e^{-(x^2+y^2)}$	$k_a(\theta)$	$= \cos^4(\theta)$
G_{4b}	$= 1.246(-1.5x + x^3)(y)e^{-(x^2+y^2)}$	$k_b(\theta)$	$= -4 \cos^3(\theta) \sin(\theta)$
G_{4c}	$= 1.246(x^2 - 0.5)(y^2 - 0.5)e^{-(x^2+y^2)}$	$k_c(\theta)$	$= 6 \cos^2 \sin^2(\theta)$
G_{4d}	$= 1.246(-1.5y + y^3)(x)e^{-(x^2+y^2)}$	$k_d(\theta)$	$= -4 \cos(\theta) \sin^3(\theta)$
G_{4e}	$= 1.246(0.75 - 3y^2 + y^4)e^{-(x^2+y^2)}$	$k_e(\theta)$	$= \sin^4(\theta)$

Table 7: X-Y separable basis set and interpolation functions for fourth derivative of Gaussian. To create a fourth derivative of a Gaussian rotated through an angle θ , use: $G_4^\theta = (k_a(\theta)G_{4a} + k_b(\theta)G_{4b} + k_c(\theta)G_{4c} + k_d(\theta)G_{4d} + k_e(\theta)G_{4e})$.

tap #	f1	f2	f3	f4	f5	G_4 basis filter	filter in x	filter in y
0	0.9344	1.0	0.0	0.0	-0.5581			
1	0.0606	0.7788	-0.4867	0.4851	-0.2173	G_{4a}	f1	f2
2	-0.5729	0.3679	-0.1839	0.4583	0.2053	G_{4b}	f3	f4
3	-0.1231	0.1054	0.1186	0.1970	0.2059	G_{4c}	f5	f5
4	0.1084	0.0183	0.0916	0.0456	0.0715	G_{4d}	f4	f3
5	0.0507	0.0019	0.0229	0.0060	0.0124	G_{4e}	f2	f1
6	0.0084	0.0001	0.0028	0.0005	0.0012			

Table 8: 13-tap filters for x - y separable basis set for G_4 . Filters for which tap 0 is 0.0 have odd symmetry about tap 0; the others have even symmetry. These filters were taken from Table 7, with a sample spacing of 0.5. Use the interpolation functions of Table 7.