Efficient discrete Gabor functions for robot vision

Carl F. R. Weiman

Transitions Research Corporation, Shelter Rock Lane, Danbury, CT, 06810 USA
Voice: +1 (203) 798-8988   *Fax:* +1 (203) 791-1082   *e-mail:* weiman @ trc.com

ABSTRACT

A new, discrete Gabor function provides sub-pixel resolution of phase while overcoming many of the computational burdens of current approaches to Gabor function implementation. Applications include hyperacuity measurement of binocular disparity and optic flow for stereo vision. Convolution is avoided by exploiting band-pass to sub-sample the image plane. A general purpose front end processor for robot vision, based on a wavelet interpretation of this discrete Gabor function, can be constructed by tessellating and pyramiding the elementary filter. Computational efficiency opens the door to real-time implementation which mimics many properties of the simple and complex cells in the visual cortex.

1. INTRODUCTION

We describe a new discrete Gabor filter which serves as the kernel for many useful functions of front end vision in a computationally simple framework. These functions include edge detection, hyperacuity phase measurement, and characterization of localized, oriented spatial frequency components. Research in biological vision has discovered such functions in the simple and complex cells of the visual cortex. Artificial vision has provided theoretical and artificial models which both explain and drive biological observations. In a unifying model, Florack and Koenderink [8], [13] have derived the essential properties of a general front-end "geometrical engine" for image processing based on first principles' requirements for linearity, spatial shift invariance, isotropy, and scale invariance. The result is a pyramidal hierarchy of local directional derivatives of a Gaussian kernel, from zeroth to fourth order. These correspond intuitively to luminance, edge, gradient, curvature, and local texture elements. The collection of differential operators at a point is known as a "jet bundle". Path integrals correspond to solutions of global problems, e.g. shape-from-shading. The columnar architecture of the visual cortex corresponds to a tiling of jet bundles over the visual field.

The elementary component of Koenderink's geometrical engine is a kernel based on combinations of derivatives of Gaussian functions. This model subsumes all the well known local operators of machine vision including Sobel, differences of Guassians, LaPlacian of Gaussian and Gabor functions. When parametrized with values which optimize spatial resolution and frequency bandpass, all such kernels converge to instantiations which are nearly indistinguishable, particularly in discrete approximation. Of these, we choose Gabor functions as our point of departure because they are parametrized explicitly in spatial and frequency domain terms which are directly relevant to image representation and analysis.

2. GABOR FUNCTIONS

2.1 Definition of Gabor function

Design commences with consideration of the fundamental lower limit on positional uncertainty imposed by the discrete sampling pattern intrinsic to pixel tessellation. This, in turn, imposes an upper limit on spatial frequency resolution for image analysis. Gabor [10] defined a quadrature mirror filter which minimizes joint time/frequency domain uncertainty,

$$g(x; \mu, \sigma, \omega) = Gauss(x; \mu, \sigma)e^{i\omega x} \qquad 1$$

where

$$Gauss(x; \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(\frac{(x-\mu)^2}{2\sigma^2}\right) \qquad 2$$

Expanding the second term of equation 1 into its real and imaginary parts yields the familiar Gabor quadrature pair, illustrated in figures 1a and 1b, consisting of cosine and sine wave modulation under a Gaussian envelope[a],

$$g = g_c + i\,g_s \qquad\qquad 3$$

where

$$g_c(x;\mu,\sigma,\omega) = Gauss(x;\mu,\sigma)\cos(\omega\,x) \qquad\qquad 4$$

and

$$g_s(x;\mu,\sigma,\omega) = Gauss(x;\mu,\sigma)\sin(\omega\,x) \qquad\qquad 5$$
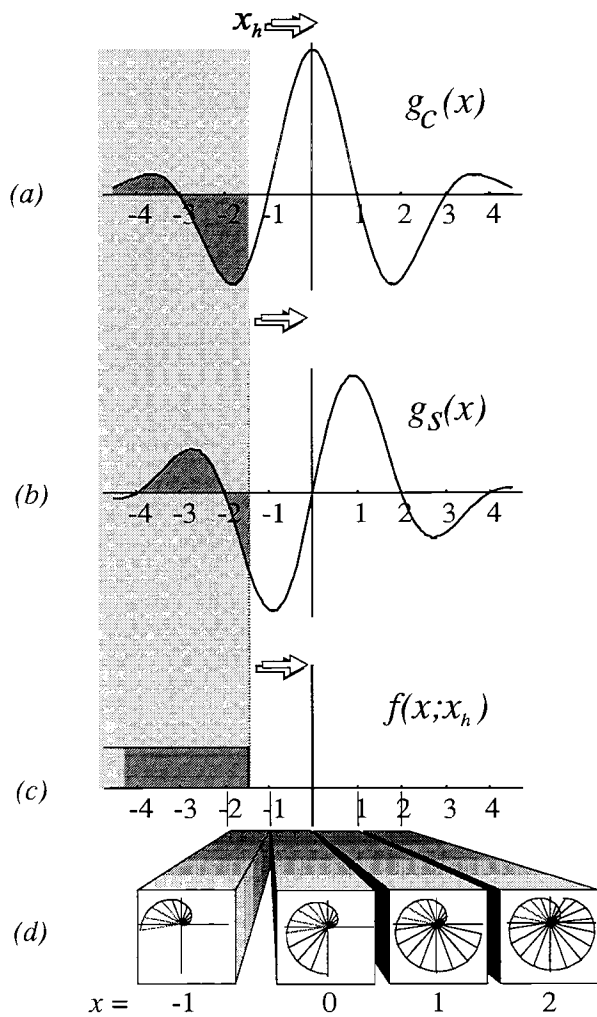
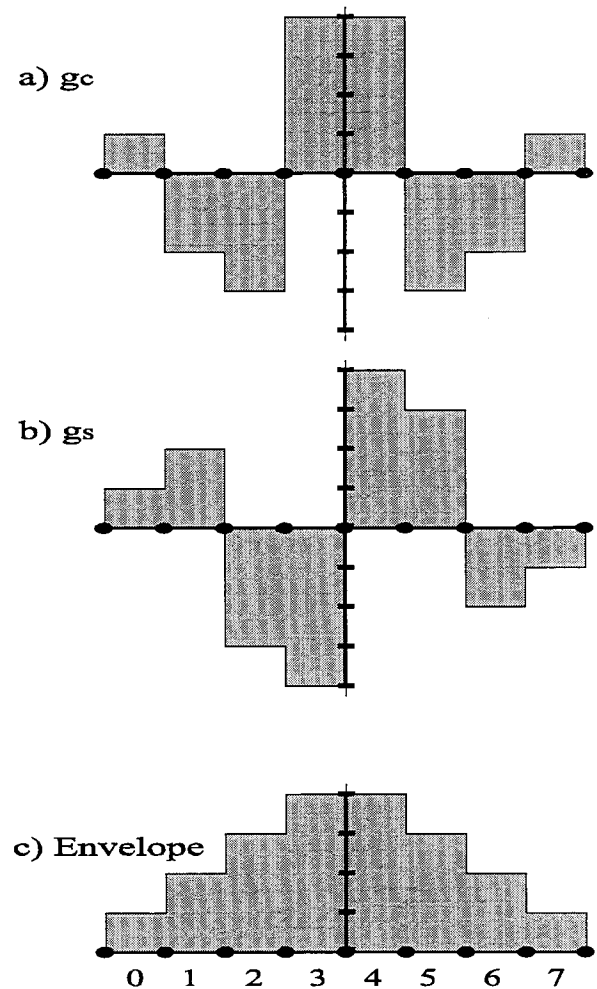

Figure 1 Gabor function quadrature pair

Figure 2 Quantized 1-D Gabor function

---

[a] Notation: Arguments of functions preceding the semicolon are free variables; arguments following the semicolon are fixed parameters. For example, in equation 1, x is free, followed by fixed parameters for mean, standard deviation, and frequency. For clarity, arguments are omitted later when obvious.

The modulation ω is limited by the Nyquist sampling rate of pixel tessellation and σ determines bandpass. These parameters determine the design of the discrete filter shown in figure 2. Before leaving the continuous Gabor function to derive its iscrete counterpart, we describe the concept of phase because of its importance to hyperacuity.

2.2 Gabor phase as measure of position

A primary application of Gabor filters in artificial vision is the measurement of phase to subpixel precision (hyperacuity). A major objective of our discrete filter design is to measure phase accurately and efficiently. Phase correlates with image position which must be precisely measured for binocular disparity and optic flow [6], [7], [21], [23]. Fleet and Jepson[7] have shown that phase is robust against noise, changes in illumination, and small geometric distortions arising from different perspective viewpoints. We illustrate Gabor phase with the following simple example.

Consider the image of a contrast edge which is bright to the left (intensity = 1) and dark (intensity = 0) to the right of some position x_h along the horizontal axis, as illustrated in figure 1c, where the shaded area corresponds to the bright side. That is,

$$f(x;x_h)=1 \quad for \quad x<x_h \quad ; \quad f(x;x_h)=0 \quad for \quad x \ge x_h \qquad 6$$

Now, applying the cosine and sine components of the Gabor filter (g_c and g_s in equations 4 and 5) to the edge step function of equation 6 corresponds to the functionals

$$G_c(x_h)= \int_{-\infty}^{\infty} f(x;x_h)g_c(x)dx = \int_{-\infty}^{x_h} g_c(x)dx \qquad 7$$

and

$$G_s(x_h)= \int_{-\infty}^{\infty} f(x;x_h)g_s(x)dx = \int_{-\infty}^{x_h} g_s(x)dx \qquad 8$$

illustrated by the shading in figures 1a and 1b. That is, equations 7 and 8 express the outputs of filters g_c and g_s applied to edge f(x; x_h). The curved outline of figure 3 plots the joint trajectory of G_c and G_s as x_h , the visual edge, moves across the filter window. The circular path reflects the underlying sine and cosine modulation while the spiraling to the center is a consequence of the Gaussian damping. Any given point on this trajectory corresponds to some position x_h of the edge. The spokes mark quarter-unit positions. Every fourth spoke thus represents a one complete unit of position change on the x-axis, labeled in correspondence with figure 1. Figure 1d expands the view of the x-axis above; icons are miniatures of the phase trajectory of figure 3. The progression of icons corresponds to the phase associated with the segment in which each icon resides.
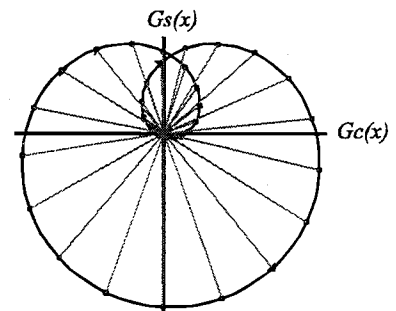

Figure 3 Gabor phase trajectory

Regarding the G_c and G_s axes in figure 3 as components of a complex number, in keeping with equation 3, the argument (phase angle) corresponds to the phase of the output of the quadrature filter pair.

$$\phi_h = Arg(G_c(x_h)+iG_s(x_h)) = \arctan(G_s(x_h)/ G_c(x_h)) \qquad 9$$

Note that phase is independent of amplitude and offset of the image step function. That is, for a stronger edge,

$$k f(x;x_h) \; ; \; k>1 \qquad 10$$

equations 7 and 8 are scaled proportionally by k, and their quotient in 9 cancels k. DC offsets integrate to zero, contributing nothing to G_c or G_s.

For Fourier transforms, phase is linear with position shift Δx, and sensitivity is proportional to spatial frequency ω,

courtesy of the Fourier shift theorem. The circular part of the trajectory in figure 3 shows that this relationship holds fairly well at the center of the Gabor filter, but not at the ends. The Gabor function has been characterized as a local Fourier transform. The distribution of spokes in the lower half plane indicates that transit of one half wavelength in the spatial domain corresponds to approximately 180° of phase transition, and the distribution is rather uniform. For a wavelength of four units, this represents 90° phase transit per unit. Thus, even modest phase discrimination techniques yield hyperacuity measurements for binocular disparity and optic flow. We shall revisit this crucial property of phase in terms of discrete filters later, where units correspond to pixels.

3. DERIVATION OF DISCRETE GABOR FUNCTION

We now design a discrete Gabor filter whose size and computational complexity are at a minimum while maximizing performance in terms of robustness and resolution. The theoretical underpinnings of Gabor functions, jet bundles, LaPlacians of Gaussians and gradients are based on continuous mathematics. Implementations in artificial vision (robotics, image processing, image compression) are necessarily discrete. Pixel positional quantization imposes certain restrictions on choice of parameter values. Proper choice of discrete parameters and geometry, e.g. window size, has important consequences for computational efficiency, which we exploit to the fullest.

3.1 Spatial Frequency

The ultimate limit to resolution is the discrete tessellation of pixels in the image plane. Their separation dictates the highest modulation frequency possible under the Gaussian envelope of the Gabor function; courtesy of Nyquist this frequency ω and the corresponding wavelength L are

$$\omega = \pi \, (perpixel) \; ; \; L = 2 \, (pixels) \qquad\qquad 12$$

Figure 4 illustrates the quantized sine wave whose frequency is given in the preceding equation, normalized to unit amplitude. But, the Gabor quadrature pair requires a cosine wave as well, i.e., a quarter wave shift of the sine. For L = 2, this is unrealizable since the shift represents a fraction (1/2) of a pixel; subpixel weighting for such a shift would yield zero value throughout as is obvious from figure 4. Thus, the only permissible integer valued wavelengths must be multiples of four pixels. The highest possible frequency (shortest possible wavelength) satisfying these criteria are

$$\omega = \pi \, / \, 2 \; ; \; L = 4 \qquad\qquad 13$$

as illustrated in figure 5. Note that quantization and symmetry reduce the sine and cosine to square waves, which will simplify computation. Dots on the x axis in figures 4 and 5 correspond to pixel boundaries.
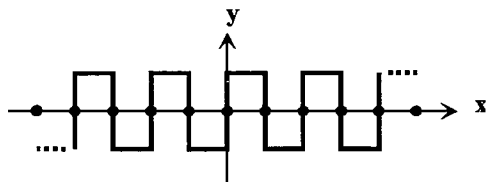

Figure 4 Highest possible frequency quantized sine wave

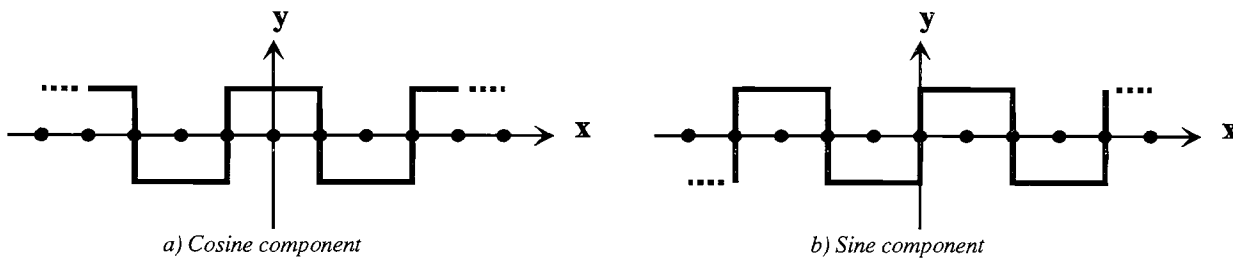a) Cosine component   b) Sine component

Figure 5 Highest possible frequency quantized quadrature pair

Having derived the frequency modulation for the discrete filter, we now derive the Gaussian envelope.

3.2 Bandpass

For the Gabor filter, reciprocal uncertainty constraints between image plane position and frequency plane bandpass are in exact analogy with the Heisenberg Uncertainty Principle, which Gabor explicitly cites [10]. Small values of sigma for the Gaussian envelope yield large bandpass which leads to undesirable aliasing; Jepson and Fleet [7] cite scale space phase detection singularities arising from such aliasing. Bandpass must however be large enough to bridge the gap between levels of an image pyramid, lest the system be blind to intermediate spatial frequencies. For dyadic up- and down-sampling (doubling/halving in each dimension), the usual case for image pyramids, each level of the pyramid corresponds to a jump of one octave. Thus one octave bandpass is necessary to span levels. Translating Sanger's equation 5 [21] from frequency to spatial domains, the relation between bandpass and σ (of equation 1) is

$$B = \log_2 \left(\frac{2\pi\sigma + L}{2\pi\sigma - L} \right)$$ 14

where B is bandpass in octaves and L is wavelength of the underlying Gabor sine and cosine functions. Thus, bandpass of one octave implies that the quantity within parentheses equals 2, whence

$$\sigma = .48L$$ 15

The functions illustrated in figure 1 exhibit precisely this relation between wavelength L and Gaussian spread σ. Note that the modulated cosine 1a) resembles a cross section of the widely used "Mexican hat function" defined as the difference of Gaussians or the 2nd derivative of a Gaussian (LaPlacian of Gaussian). The companion b) closely resembles the 1st derivative of a), in keeping with Koenderink's jet bundle model.

Now, using the quadrature pair of figure 5 for constructing the discrete filter, wavelength L is 4 pixels. Substituting this value into equation 15 yields

$$\sigma = .48x4 = 1.92(pixels)$$ 16

for the Gaussian modulation. Figure 6 illustrates the Gaussian for this value of σ quantized per pixel by integrating the continuous curve over unit (single pixel) intervals.

Note that contributions beyond four pixels from the center are very small. Thus an 8 pixel window is both necessary and sufficient to accommodate the derived values of L and σ (equation 16). Figure 2 illustrates the quantized values of the quadrature pair of this discrete Gabor function generated by applying the envelope of figure 6 to the waves of figure 5. Constructively, this operation can be regarded as flipping the signs of alternating pairs of coefficients of the quantized Gaussian.
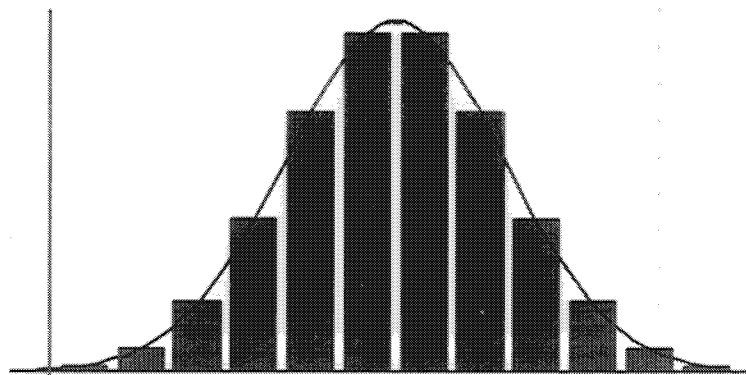
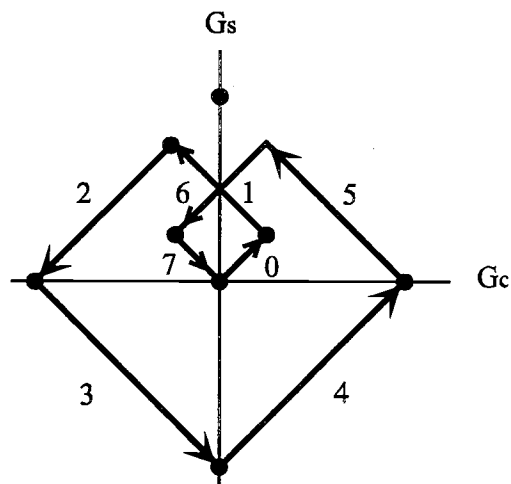Figure 6 Quantized Gaussian for $\sigma = 1.92$ pixels

Figure 7 Phase trajectory of discrete Gabor filter

3.3 Phase of discrete Gabor filter

The discrete filter offers high precision positional measurement via phase just as the continuous Gabor function. Figure 7 illustrates the phase trajectory of the discrete Gabor filter of table 2 as a vertical edge is moved through the window, analogous to figure 3. A striking difference between the continuous and discrete cases is that the latter displays straight line segments rather than a circular trajectory. Examination of figure 2 reveals why. As an edge moves across any single pixel, the contributing area under the filter grows linearly for both the sine and cosine component. Every pixel boundary is accompanied by a sign transition in either the sine or cosine filter coefficient. Thus, each linear segment in the phase trajectory corresponds to the transit of an edge across one pixel. It is evident that subpixel motions can easily be measured with high precision with such a mechanism (hyperacuity).

Two minor imperfections are manifested in figure 7 which prevent this filter from being an ideal phase detector. One is the small inscribed diamond shaped trajectory near the origin, which represents the "ramping up" of filter energy in the outer pair of entry pixels (numbered 0 and 1 in figure 2) at the beginning of the trajectory and ramping down at the pair of exit pixels (labeled 6 and 7 in figure 2) at the end of the trajectory. Aside from these four pixels, linear segments for the central four pixels are nearly segregated by quadrant. The fact that they are only *nearly* so comprises the second imperfection. Segments labeled 2 and 5 (corresponding to so numbered pixels in figure 2) are 25% shy of filling the quadrant, unlike segments 3 and 4 which fill their respective quadrants. If pixel transits were exactly segregated by quadrant, phase computation would be extremely simple. Filter signs would point to pixel, and phase would be linearly related to subpixel position.

3.4 Ideal Phase Detection Filter

We now "reverse engineer" the discrete Gabor filter of figure 2 by adjusting its coefficients in order to establish a linear relationship between phase and edge position and segregate pixel transits by phase quadrant. In so doing, we wish to preserve as much as possible the frequency and bandpass of the Gabor function. Thus, the underlying sine and cosine components and Gaussian envelope should be disturbed as little as possible.
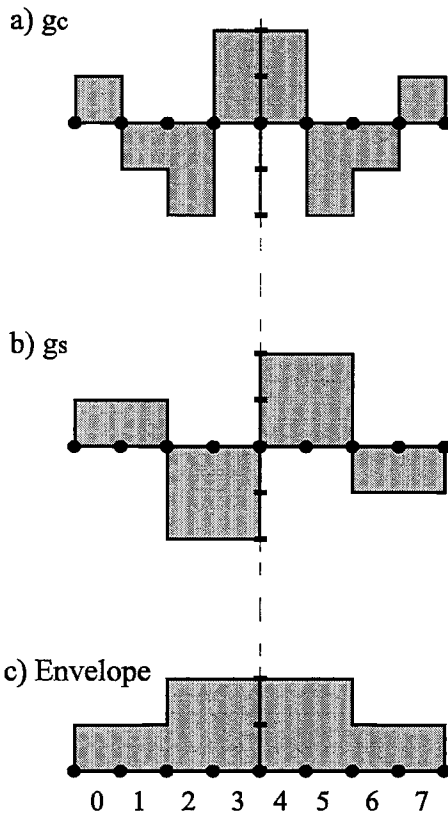
a) gc

b) gs

c) Envelope

| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |

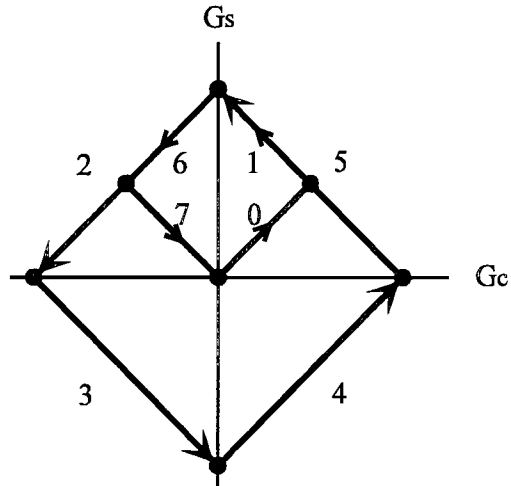*Figure 8 Ideal discrete phase filter*

Figure 9 Phase trajectory of ideal phase filter

Figure 8 illustrates an ideal phase filter which fulfills the above requirements. Figure 8c is the discrete version of its Gaussian envelope, modified so that the outer shoulders (pixels labeled 0, 1, 6, 7) are constant at half the height of the central plateau (pixels 2, 3, 4, 5). Figures 8a and 8b respectively give g_c and g_s, applying the modulation of figure 5. By weighting the central four pixels equally, the phase trajectory for the four middle pixels is a perfect square centered at the origin with vertices on the principal coordinate axes, as illustrated in figure 9. The half-weights of the outer shoulders ramp up filter energy to exactly reach the first vertex of the square. Figure 9 illustrates the perfect association between phase trajectory and position of x_h. Linear segment labels correspond to pixel numbers, and progression along a segment is linear with x_h. Arrows correspond to left-to-right motion of the edge x_h.

Since the signs of the two filter components alternate on pixel transitions, each phase quadrant corresponds to a single pixel. Thus the pixel index of an edge can be derived from the combinations of the signs of the outputs of the two filters as follows.

Labeling the pixels 0 through 7 from left to right (see figure 8) and quadrants 0 through 3 clockwise, the sign transition of quadrants is :

pixel #	0	1	2	3	4	5	6	7
quadrant	0	0	1	2	3	0	1	1
sign (g_c , g_s)	(+, +)	(+, +)	(-, +)	(-, -)	(+, -)	(+, +)	(-, +)	(-, +)

Table 1. Phase quadrant relations for ideal phase filter

The symmetry of figure 9 shows that this filter can be used as an ideal "phase only" filter if we modify the definition of phase to correspond to linear segment traverse, namely,

$$\phi \ (x_h) = \frac{G_c - G_s}{|G_c| + |G_s|}$$

17

for even quadrants and

$$\phi\ (x_h) = \frac{G_c + G_s}{\left|G_c\right| + \left|G_s\right|}$$ 18

for odd quadrants. Note that this function is invariant under brightness shifts because each filter is immune to DC changes, and invariant under brightness scaling because scale factors cancel in numerator and denominator. The values of this phase function range exactly linearly with edge subpixel position x_h from -1 to 1 within each quadrant (pixel). The phase value 0 exactly bisects the pixel. The denominator is constant over the middle four pixels of the filter window. Thus, the complexity of subpixel phase computation is simpler than the continuous case, and exactly linear with position in the image.

4. COMPUTATIONAL PROPERTIES OF DISCRETE GABOR FUNCTIONS

We now demonstrate that the ideal phase filter defined above is "well-behaved" according to criteria applied to signal and image processing filters. These criteria are quantified in terms of mathematical measures of filter energy, magnitude, symmetry, linearity, balance, isotropy, orientation, bandpass, and phase. Satisfaction of such criteria exonerates the filter from the potential onus of distortion, sanctioning broad capabilities for operations such as local derivative, bandpass filter, wavelet decomposition, phase detector, local Fourier transform and edge detector, i.e. as a qualified candidate for the elementary kernel of Koenderink's visual front-end geometric engine or any other image processing architecture.

4.1 Magnitude and Energy

A basic requirement of wavelets, bandpass filters, local derivatives, edge detectors and Gabor filters is zero response to a DC signal. For a discrete filter, this is equivalent to the requirement that the sum of coefficients must be zero. Clearly this is the case for the ideal phase filter, i.e.

$$\sum_{j=0}^{7} g_c(j) = 1-1-2+2+2-2-1+1 = 0 \ ; \ \sum_{j=0}^{7} g_s(j) = 1+1-2-2+2+2-1-1 = 0$$ 19

Note also the even and odd symmetry of g_c and g_s, respectively,

$$g_c(j) = g_c(7-j) \ ; \ g_s(j) = -g_s(7-j)$$ 20

An additional requirement of quadrature is that pointwise energy (sum of squared amplitudes) of the filter pairs sum to a constant for all values of phase; this would correspond to a circle in the phase trajectory. Clearly for discrete filters, a different definition of energy is required, namely, the so-called "metropolitan metric" which is shown in the denominator of equations 17 and 18. Under this definition, summed energy is constant for the four middle pixels.

$$\left|\ G_c\ \right| + \left|\ G_s\ \right| = 4$$ 21

The filters contain the same energy, each summing to the area of the discrete Gaussian of figure 6, and they are orthogonal by virtue of the 90° phase shift of figure 5.

The preceding discourse was restricted to one dimension for clarity. None of the essential properties above is altered by extension to two dimensions; thus notation was simplified by omitting the extra index. We now proceed to 2-D.

4.2 Generalization of Gabor functions to 2-D

We now extend the discrete 1-D filters described above to discrete 2-D filters. Discrete image processing filters are customarily defined as 2-D masks whose coefficients multiply image intensity over underlying pixels, summing the result,

$$\sum_{k=0}^{m-1} \sum_{j=0}^{n-1} M(j,k) f(x+j,y+k) \qquad\qquad 22$$

where j and k index mask row and column, x and y index image pixels, M is mask coefficient and f is pixel gray value. A number of conventions improve implementation efficiency. The double sum itself implies a rectangular domain, which simplifies indexing into the image and filter table. Small values of m and n limit the required number of multiplications and additions (m x n). Fixed point, limited precision values of $M(j,k)$ reduce the complexity of ALU's required to carry out the sums and products.

We derive such a filter by abbreviating Daugman's landmark paper [5] which generalized Gabor functions from one to two dimensions for application to visual processing. Extend the Gabor function to two dimensions by modulating the 1-D functions with a Gaussian in the orthogonal direction,

$$g_c(x,y;\mu_x,\mu_y,\sigma_x,\sigma_y,\omega_x) = g_c(x;\mu_x,\sigma_x)Gauss(y;\mu_y,\sigma_y) \qquad\qquad 23$$

$$g_s(x,y;\mu_x,\mu_y,\sigma_x,\sigma_y,\omega_x) = g_s(x;\mu_x,\sigma_x)Gauss(y;\mu_y,\sigma_y) \qquad\qquad 24$$

as illustrated in figure 10.

Cross sections parallel to the x-axis are scaled 1-D Gabor functions thereby validating, by linearity, the generalization to 2-D of 1-D magnitude, energy and symmetry properties expressed in the preceding section. Cross sections of the 2-D Gabor function parallel to the y-axis are scaled Gaussians. The 2-D generalization has several important consequences for image analysis. One is to improve signal-to-noise ratio in the detection of oriented edges. Integration along wave crests or troughs cancels out random noise along the edge, in like manner to Canny's edge detector [4]. Noise variance noise is thereby reduced proportionally to cross-wave sampling volume of the filter (classical σ/n phenomenon from statistical sampling theory). Additionally, Gaussian damping across the wave eliminates sinc function ringing in that direction which would result from rectangular windowing. And finally, Daugman shows that cross-wave damping constrains the edge-orientation tuning of the filter. Such constraint is necessary if the filter is to effectively discriminate edge orientation.

Figure 2 represents a cross section of a discrete filter along the direction of the wave. The 2-D extension is effected by Gaussian modulation in the orthogonal direction (equations 23 and 24). Daugman shows that orientation tuning is roughly proportional to the value of σ_y. In our experiments, σ_y values approximately equal to σ_x afford sufficient tuning to cover all possible orientations with a set of four filters oriented in intervals of $\pi/4$. By the argument surrounding figure 6, an 8x8 window is required. Figure 12 illustrates the 2-D extension of figure 2c, to a bivariate discrete Gaussian. Figure 11 applies this "ziggurat" function to the sine-cosine pair of figure 5 to generate the 2-D Gabor function extended from figure 2.
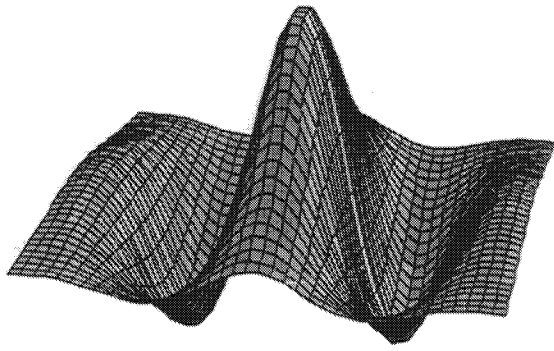
a) Cosine component, g_c
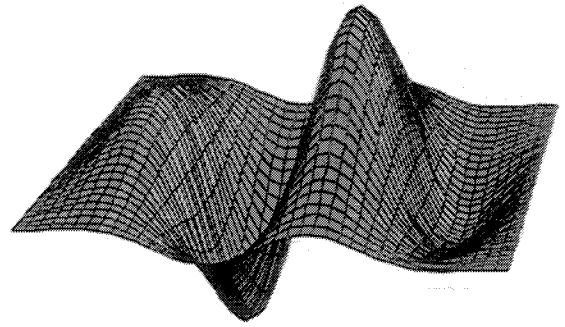
b) Sine component, g_s
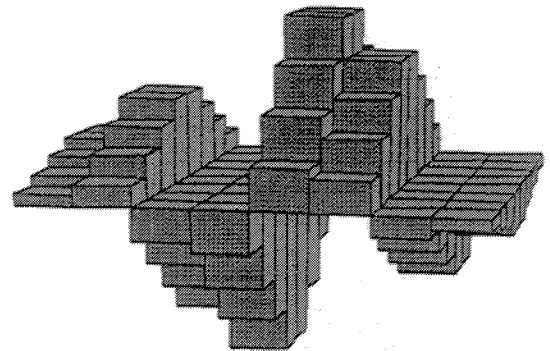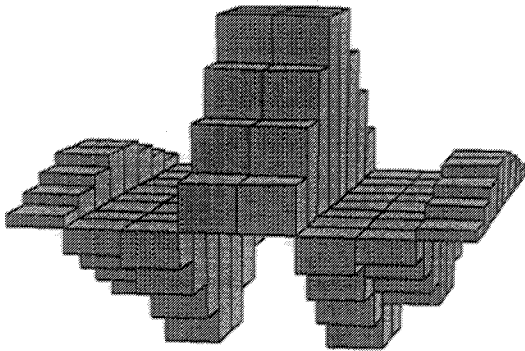
Figure 10 Continuous 2-D Gabor functions

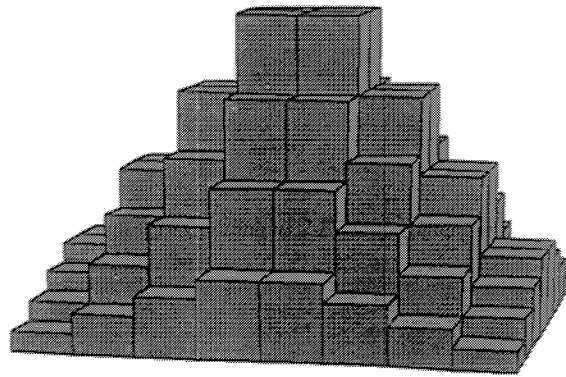*Figure 11 Discrete 2-D Gabor functions*

Figure 12 "Ziggurat" function - 2-D discrete Gaussian envelope

5. COVERING THE SPATIAL AND SPATIAL-FREQUENCY DOMAINS

The preceding derivation discussed local properties of the filter, showing that it is efficient and possesses all the required mathematical properties of quadrature mirror filters, including hyperacuity phase detection. We now discuss global coverage including the entire image plane, spatial frequency spectrum, and orientation space. Such coverage is required if the filter is to be used as the kernel of a "front end" geometrical engine satisfying Koenderink's invariance criteria. Starting with orientation, we can cover horizontal and vertical orientations by transposing indices in the discrete Gabor filters, i. e., rotating them 90°. We have found that a 45° filter is also desirable because diagonal edges do not behave well in phase detection on horizontal and vertical filters due to the coarseness of quantization.

Now consider frequency spectrum completeness. By virtue of the choice of σ_x for the (nearly) Gaussian envelope, bandpass for the filter described is one octave centered at $\pi/2$ cycles per pixel. Any dyadic pyramiding process scales this band by exactly one octave, so that all frequencies are covered, at half the frequency resolution per octave.

Square tessellation by adjacent 8x8 blocks, each covered by four Gabor filter pairs (horizontal, vertical, right diagonal, left diagonal) seems the most straightforward as illustrated in figure 13a. However, the tapering energy profiles of the filters would leave weak zones at the borders of such tiles. In addition to uneven sensitivity (a violation of Koenderink's translation invariant requirement), invalid phase measurements at the outer two-pixel zone of each filter would gravely flaw the performance of this system for optic flow or binocular disparity. Hence, we overlap the blocks by 50% in the pattern shown in figure 13b which indicates neighborhood centers as diamonds. Thus the central 4x4 block gives perfect phase readings for an edge, and the entire image plane is covered by such tiles. Mutual inhibition of nearest neighbor blocks, based on filter energy, suppresses the weaker signals from filters whose outer bands are triggered by edges which overlap the central 4x4 block of a neighboring center.
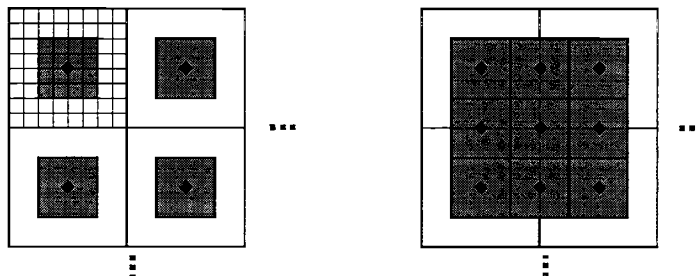

Figure 13 Tessellation of 8x8 neighborhoods

A number of dyadic up-sampling schemes are available for pyramiding. Burt [3] describes efficient hardware which can compute a Gaussian up- and down-sampling over 5x5 neighborhoods at extremely high speeds. Whatever technique is used, the discrete Gabor function described in this paper can be applied as the fundamental kernel at any level of the pyramid.

6. CONCLUSIONS

A discrete Gabor function for general image processing applications has been described. Based on 8x8 windows, it offers maximum available spatial resolution, and hyperacuity phase detection with minimal window size and computational complexity within the constraints of bandpass and intrinsic pixel resolution. Any application which uses Gabor functions, for example, wavelet [1], [16] and fractal decomposition, can benefit from its efficiencies. A striking example is the use of phase to measure binocular disparity or optic flow. The tessellation of centers in 4-fold decimation in each dimension of the pixel grid represents a 16-fold reduction in computation count compared to convolution, which centers filters at every pixel. A similar reduction applies in comparison to cross-correlation methods for binocular disparity measurement.

Other applications which can benefit from the efficiency of the discrete Gabor function include steerable filters [9], image compression [28], shape from shading [19] and pattern recognition.

In general terms, the Gabor function serves as wavelet basis for representation and analysis of imagery. Higher order jet bundles can be derived from Gabor kernels by differencing nearest neighbors in the filter tessellation, or by taking advantage of the Fourier-like local properties of the Gabor function, e.g. scaling filter outputs proportional to spatial frequency to get the Fourier transform of the derivative:

$$G(\frac{df}{dx};\omega) \Leftrightarrow \omega\, G(\,f\,;\omega) \qquad\qquad 25$$

As a basis for representation and analysis, the underlying image coordinate system is of no consequence to the downstream computational mechanics, but its transformations may prove useful in the semantics of the application. For example, log-

polar coordinates render zoom and rotation separable [26], [27]. Hence image looming induced by robot motion through an environment becomes a one-dimensional transformation. If the underlying coordinate transformation is conformal (which log-polar is [24]), local orientation measures are rotated orthogonally, i.e. preserving the angles of vertices and the shapes of texture components. This is most important for perspective invariance; cross-ratios based on lengths may be distorted by the new metric but angle-based cross-ratios are unchanged [17].

In the spirit of Koenderink, the discrete system presented here can be regarded as a model for certain aspects of the visual cortex, with Gabor functions as simple cells. This suggests searching for biological analogues of phase detection circuitry expressing equations 17 and 18. Pollen [20] has demonstrated the existence of Gabor-like quadrature pair neurons in the visual cortex. Ohzawa's model [18] of binocular disparity detection is based on the energy of such pairs. The model presented here suggests a search for cells which respond to phase downstream from these quadrature pairs, i.e. scaling the difference in amplitude by the sum of energies of pair members as in equations 17 and 18.

7. ACKNOWLEDGMENTS

8. REFERENCES

[1] Benedetto, J. J. and Frazier, M. W., Wavelets: Mathematics and Applications, CRC Press, Boca Raton, 1994.

[2] Burt, P. J., and Adelson, E. H., "The LaPlacian pyramid as a compact image code", IEEE Trans. Comm., Vol. COM-31, April 1983.

[3] Burt, P. J., "Smart sensing with a pyramid vision machine", Proc IEEE, Vol 76, pp1006-1014, August, 1988.

[4] Canny, J., "A computational approach to edge detection", IEEE Trans PAMI, 8, pp. 679-698, 1987.

[5] Daugman, J. C., "Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters", Journal of the Optical Society of America, Vol 2, No. 7, July 1985, pp. 1160-1169.

[6] Fleet, D. J., Measurement of Image Velocity, Kluwer Academic Publishers, Boston, 1992.

[7] Fleet, D. J., and Jepson, A. D., "Stability of phase information", IEEE Trans PAMI, Vol 15, pp. 1253-1268, Dec, 1993.

[8] Florack, L.M.J., Romeny, B. M. tH., Koenderink, J.J., and Viergiver, M. A., "Scale and the differential structure of images", Image and Vision Computing, Vol 10, pp 376-387, 1992.

[9] Freeman, W. T., and Adelson, E. H., "The design and use of steerable filters", IEEE Trans. PAMI, Vol 13, No. 9, pp. 891-906, September, 1991.

[10] Gabor, D., "Theory of Communication", Journal of the Inst. of Elec. Engrs, v. 93, pp 429-459, 1946.

[11] Gennert, M. A. and Malin, J. A., "Stereo vision using Gabor receptive fields", Proceedings of of SPIE Conference on Intelligent Robots and Computer Vision XI: Biological, Neural Net and 3-D Methods, Casasent, Vol 1826, pp 64-75, Boston, Nov. 1992.

[12] Huebel, D. G. and Wiesel, T. N. "Receptive fields, binocular interaction, and functional architecture in the cat's visual cortex", Journal of Physiology Vol 160, pp. 106-154, 1962.

[13] Koenderink, J. J., and van Doorn, "Representation of Local Geometry in the Visual System", Biological Cybernetics, Vol 55, pp. 367-375, 1987.

[14] Laine, A. and Fan, J., "Texture classification by Wavelet Packet Signatures", *IEEE Transactions on PAMI*, Vol. 15, No. 11, pp. 1186-1191, November, 1993.

[15] Laine, A., Schuler, S., and Girish, V., "Orthonormal wavelet representations for recognizing complex notations", *Machine Vision and Applications*, Vol 6, pp. 110-193, 1993.

[16] Meyer, Y. (translated by R. D. Ryan), Wavelets: Algorithms and Applications, Society for Industrial and Applied Mathematics, Philadelphia, 1993.

[17] Mundi, J. L., and Zisserman, A., Geometric Invariance in Computer Vision , MIT Press, Cambridge, 1992.

[18] Ohzawa, I., DeAngelis, G. C., and Freeman, R. D., "Stereoscopic depth discrimination in the visual cortex: neurons ideally suited as disparity detectors", *Science*, Vol 249, pp. 1037-1041, August, 1990.

[19] Pentland, A. P., "Linear shape from shading", *Int'l Jnl of Computer Vis.*, Vol 4, pp. 153-162, 1990.

[20] Pollen, D. and Ronner, S, "Phase relationships between adjacent simple cells in the visual cortex", *Science*, Vol 212, pp 1409-1411, 1981.

[21] Sanger, T. D., "Stereo disparity computation using Gabor filters", *Biological Cybernetics*, Vol 59, pp. 405-418, 1988.

[22] Theimer, W. M., Mallot, H. A., and Tölg, S., "Phase Method for binocular vergence control and depth reconstruction", *Proceedings of of SPIE Conference on Intelligent Robots and Computer Vision XI: Biological, Neural Net and 3-D Methods,* Casasent, Vol 1826, pp 76-87, Boston, Nov. 1992.

[23] Toelg, S., "Gaze control for an active camera system by modeling human pursuit eye movements", *Proceedings of of SPIE Conference on Intelligent Robots and Computer Vision XI: Algorithms, Techniques and Active Vision,* Casasent, Vol 1825, pp 585-598, Boston, Nov. 1992.

[24] Weiman, C. F. R., and Chaikin, G. M., "Logarithmic spiral grids for image processing and display", *Computer Graphics and Image Processing*, Vol 11, pp 197-226, 1979.

[25] Weiman, C. F. R., "3-D sensing with exponential sensor arrays", *Proc. SPIE Conf. on Pattern Recognition and Signal Processing*, Vol. 938, *Digital and Optical Shape Representation and Pattern Recognition*, Orlando, 1988.

[26] Weiman, C. F. R. and Juday, R. D., "Tracking Algorithms Using Log-Polar Mapped Image Coordinates", *Proc. SPIE Conf. on Intelligent Robots and Computer Vision VIII: Algorithms and Techniques*, Vol. 1192, Phila., November 1989.

[27] Weiman, C. F. R., "Log Polar Vision for Mobile Robot Navigation", *Electronics Imaging: International Electronic Imaging Exposition and Conference*, pp.382-386, Oct 1990a.

[28] Weiman, C. F. R., "Video Compression Via Log Polar Mapping", *Proc. SPIE Conf. on Real-Time Image Processing II*, Vol. 1295, Orlando, April 1990b.

[29] Weiman, C. F. R., and Evans, J.M., Digital Image Compression Employing a Resolution Gradient, U.S. Patent No. 5,103,306, 1992.