

Steerable Filters from Erlang Functions

Anil A. Bharath

Department of Biological & Medical Systems,
Imperial College of Science, Technology & Medicine, Exhibition Road,
London SW7 2BX
a.bharath@ic.ac.uk

Abstract

Scale and orientation steerable 2D filters are constructed using a frame of Erlang functions in the Fourier domain. Erlang forms for the radial frequency characteristic are shown to provide complex quadrature filters which can be steered in a scale parameter, α . Oriented, spatial domain filters are constructed by imposing an appropriate angular selectivity in the frequency domain. A filter bank is designed, and outputs of the filters of the bank are steered to construct an oriented scale-space decomposition of an image sub-band. Some applications are discussed.

1 Introduction

The use of sub-band methods in image analysis is perhaps not as widespread as for coding. Nevertheless, quite successful and powerful analysis algorithms have been demonstrated [1]. One can make good use of sub-band processing when image features exist at several scales, so that operators of different sizes must be employed: a key issue then becomes the design of these operators. This paper, in a manner similar to recent approaches of generating linear scale spaces, shall focus on scale-steering of operators at *one* sampling rate. This has the advantage of simplifying the registration of information across multiple scales, although the computational costs and the size of a complete representation are larger. It is well known that the use of multi-scale decompositions can yield full frequency domain coverage by the Mallat [6] algorithm. However, the processing of information at intermediate levels of scale is not straightforward from the coefficient space. It is, in fact, imperative that one makes the distinction between a pyramidal decomposition used for compact representation, from that used as a vision front-end. The requirement for feature extraction can often be aided by embedding the image in a space (see, for example, [5]) more appropriate to the extraction of useful features.

The ability to steer the orientation of an operator from a number of fixed filters has been successfully demonstrated by a number of authors [2],[4], and is thought to be a function performed by parts of the human visual system [8]. By a suitable choice of radial Fourier domain filter specification, one can also steer effectively in *scale*, but this requires certain coverage and smoothness properties of the "basis", or, more correctly, "frame" filters. In what follows, a scale and orientation steerable scheme is developed based on Erlang functions.

2 The Erlang functions

The Erlang functions are one-sided density functions which are well known in statistics [7]. In these respects, they are like the lognormal functions. The family of Erlang functions can be represented by

$$f_n(x) = A_n e^{-\alpha x} x^n U(x), \quad n = 0, 1, 2, \dots \quad (1)$$

where α is some positive constant, fixed over all families, and $U(x)$ is the generalised Heaviside function. A_n is a constant for each n , generally chosen to ensure that each $f_n(x)$ is a density function in the statistical sense:

$$\int_0^{\infty} f_n(x) dx = 1, \quad n = 0, 1, 2, \dots \quad (2)$$

2.1 Properties in the frequency domain

Filters with Erlang frequency characteristics have some properties which make them attractive for analytic purposes and as practical tools for signal and image processing. In two dimensions, the form

$$H_n(\omega) = A_n \omega^n e^{-\alpha \omega} U(\omega), \quad n = 0, 1, 2, \dots \quad (3)$$

for radial frequency variation will be used.

2.1.1 Central Frequency

The radial frequency location of the peak of each family is a linear function of n . To see this, take the derivative of $H_n(\omega)$:

$$\frac{\partial}{\partial \omega} \{H_n(\omega)\} = A_n e^{-\alpha \omega} \{n\alpha \omega^{n-1} - \alpha^{n+1} \omega^n\} \quad n = 0, 1, 2, \dots \quad (4)$$

At the peak location for each member of the family in frequency, the derivative is zero. This leads to

$$\omega_{max} = n/\alpha, \quad n = 0, 1, 2, 3, \dots \quad (5)$$

Thus, the peak frequency of each family member increases linearly with the order of the family. Equivalently, for a fixed n , the modal frequency of the filters is proportional to $1/\alpha$.

2.1.2 ω Coverage and Filter Synthesis

If a family of filters is constructed for $n = 0, 1, 2, \dots$, it can be shown that these filters cover all frequency space, and also that *any* other arbitrary filter magnitude specification which can be expressed as a finite polynomial in ω can be synthesized by a weighted combination of the frame filters, $H_n(\omega)$. Consider the N^{th} order polynomial filter magnitude specification, $G(\omega)$. It is proposed that one can always find a set of b_n such that

$$\sum_{n=0}^{\infty} b_n H_n(\omega) = G(\omega) \quad (6)$$

Expanding the left hand side of this equation yields

$$\sum_{n=0}^{\infty} b_n A_n \omega^n e^{-\alpha\omega} = e^{-\alpha\omega} \{b_0 A_0 + b_1 A_1 \omega + b_2 A_2 \omega^2 + \dots\} \quad (7)$$

so that

$$\begin{aligned} b_0 A_0 + b_1 A_1 \omega + b_2 A_2 \omega^2 + \dots &= e^{\alpha\omega} G(\omega) \\ &= G(\omega) \left\{ 1 + \alpha\omega + \frac{\alpha^2 \omega^2}{2!} + \dots \right\} \end{aligned} \quad (8)$$

Because $G(\omega)$ is by assumption a finite order polynomial, the effect of a multiplication with the infinite order expansion of $e^{\alpha\omega}$ simply changes the weights of each power of ω . The case above also immediately gives us an answer to the coverage provided by the infinite frame, $H_n(\omega)$. If a flat response across all ω is required, then setting $G(\omega) = c$ in Equation (8), where c is a real constant, and equating powers of ω , shows that one can obtain complete, uniform coverage over all frequencies if the sub-bands are weighted such that

$$b_n A_n = \frac{\alpha^n c}{n!}, \quad n = 0, 1, 2, \dots \quad (9)$$

In fact, the normalising weights applied to make the Erlang functions density functions are given by Equation (2) as

$$A_n = \frac{\alpha^{n+1}}{n!} \quad (10)$$

so that $b_n = \alpha c$, $n = 0, 1, 2, \dots$ is all that is required.

The properties outlined above hold for an infinite number of frame functions; it is fair to ask what happens when only a finite subset of the family is available. In this case, the approximation is close to the desired response $G(\omega)$ in the sense of an N^{th} order polynomial, if N consecutive frame members are used. Since the range of ω over which the approximation is required to hold is generally finite, the synthesis properties of these filters is quite good.

2.2 Spatial Domain Properties

2.2.1 1-D Properties

Because the frequency domain specification of $H_n(\omega)$ is such that

$$H_n(\omega) = 0, \quad n = 0, 1, 2, \dots \quad (11)$$

for $\omega < 0$, the corresponding continuous temporal or spatial domain filters given by

$$h_n(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} H_n(\omega) e^{j\omega t} d\omega, \quad n = 1, 2, 3, \dots \quad (12)$$

are complex quadrature filters.

2.3 2D properties

It is useful to understand the properties of these filters in the 2D plane. Consider a polar separable frequency domain specification of

$$H_n(\omega, \phi) = H_n^\omega(\omega)H^\phi(\phi) \quad (13)$$

where $H_n^\omega(\omega)$ is the Erlang specification of the radial frequency component, and H^ϕ . In this separable case, the following expression can be written for the inverse Fourier transform

$$f(r, \theta) = \frac{1}{2\pi} \int_0^\infty H_n^\omega(\omega) \int_{-\pi}^\pi H^\phi(\phi) e^{j\omega r \cos(\phi-\theta)} d\phi \omega d\omega \quad (14)$$

An "easy" form for $H^\phi(\phi)$ is

$$H^\phi(\phi) = \cos^2 \phi \operatorname{rect}\left(\frac{\phi}{\pi}\right) \quad (15)$$

where rect is the unit width and height "top-hat" function, given by

$$\operatorname{rect}(x) = U\left(x + \frac{1}{2}\right) U\left(\frac{1}{2} - x\right) \quad (16)$$

By expanding $H^\phi(\phi)$ in a Fourier series expansion (periodicity 2π), and using

$$\frac{1}{2\pi} \int_{-\pi}^\pi e^{\pm jnu + j\beta \cos u} du = j^n J_n(\beta) \quad (17)$$

one obtains (dropping the normalising A_n for simplicity),

$$\begin{aligned} f(r, \theta) &= \int_0^\infty H_n^\omega \sum_{m=0}^\infty a_m \cos(m\theta) j^m J_m(r\omega) \omega d\omega \\ &= \sum_{m=0}^\infty a_m \cos(m\theta) \int_0^\infty H_n^\omega j^m J_m(r\omega) \omega d\omega \\ &= \sum_{m=0}^\infty j^m a_m \cos(m\theta) \int_0^\infty \omega^{n+1} e^{-\alpha\omega} J_m(r\omega) d\omega \end{aligned} \quad (18)$$

where the Fourier series coefficients converge rapidly; 98.03% of the power of $H^\phi(\phi)$ is in the first 3 terms of the series, so that the series of Bessel kernel integrals approximately terminates at $m = 2$. From a table of integrals [3] the following indefinite integral is obtained:

$$\int_0^\infty x^{m+1} e^{-\alpha x} J_\nu(\beta x) dx = (-1)^{m+1} \beta^{-\nu} \frac{\partial^{m+1}}{\partial \alpha^{m+1}} \left\{ \frac{(\sqrt{\alpha^2 + \beta^2} - \alpha)^\nu}{\sqrt{\alpha^2 + \beta^2}} \right\} \quad (19)$$

where $\nu > -(m+2)$ is real and $\beta > 0$. Thus, using 3 terms of the expansion of Equation (18), leads to the following approximate analytic expression for the filters:

$$h_n(r, \theta) \approx (-1)^{n+1} \sum_{m=0}^2 j^m a_m \cos(m\theta) \left(\frac{1}{r}\right)^m \frac{\partial^{n+1}}{\partial \alpha^{n+1}} \left\{ \frac{(\sqrt{\alpha^2 + r^2} - \alpha)^m}{\sqrt{\alpha^2 + r^2}} \right\}$$




Figure 1: Real (left) and imaginary (right) parts of complex 1st order 2D steerable wavelet at 0 degrees, $\alpha = 1$. Plotted by evaluating Equation (20).

$$\begin{aligned}
 &\approx (-1)^{n+1} \left[a_0 \frac{\partial^{n+1}}{\partial \alpha^{n+1}} \frac{1}{\sqrt{\alpha^2 + r^2}} - \frac{a_2}{r^2} \cos(2\theta) \frac{\partial^{n+1}}{\partial \alpha^{n+1}} \frac{(\sqrt{\alpha^2 + r^2} - \alpha)^2}{\sqrt{\alpha^2 + r^2}} \right. \\
 &\quad \left. + j \frac{a_1}{r} \cos(\theta) \frac{\partial^{n+1}}{\partial \alpha^{n+1}} \frac{\sqrt{\alpha^2 + r^2} - \alpha}{\sqrt{\alpha^2 + r^2}} \right] \\
 &\approx (-1)^{n+1} \frac{\partial^{n+1}}{\partial \alpha^{n+1}} \left[\frac{1}{4} \left\{ \frac{1}{\sqrt{\alpha^2 + r^2}} - \frac{(\sqrt{\alpha^2 + r^2} - \alpha)^2}{r^2 \sqrt{\alpha^2 + r^2}} \cos(2\theta) \right\} \right. \\
 &\quad \left. + \frac{4j}{3\pi} \left\{ \frac{\sqrt{\alpha^2 + r^2} - \alpha}{r \sqrt{\alpha^2 + r^2}} \cos(\theta) \right\} \right] \quad (20)
 \end{aligned}$$

For the case $n = 1$, this yields

$$h_1(r, \theta) = \frac{1}{4} \left\{ \frac{2\alpha^2 - r^2}{(\alpha^2 + r^2)^{5/2}} - \frac{3r^2}{(\alpha^2 + r^2)^{5/2}} \cos(2\theta) \right\} + j \frac{4}{\pi} \frac{\alpha r}{(\alpha^2 + r^2)^{5/2}} \cos(\theta) \quad (21)$$

This is a complex function. Surface plots of the real and imaginary parts of this function are given in Figure (1). Note that the axes of symmetry and antisymmetry for the real and imaginary components are co-incident. This complex quadrature filter pair represents a scale-tuned, oriented line and edge operator.

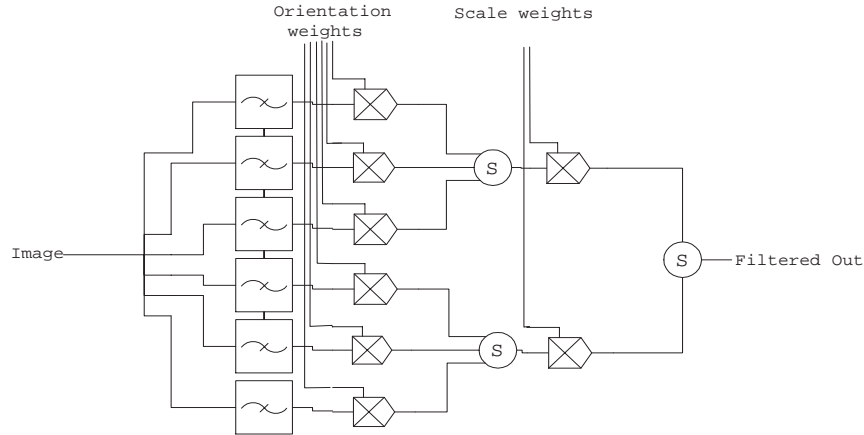


Figure 2: Simplified block architecture consisting of three filter orientations at two scales. First weight bank controls orientation steering, second bank, scale steering. **S** denotes summation.

3 Scale and Orientation Steerability

3.1 Scale Steerability

The scale steerability of the filters arises through the properties of the Erlang frequency domain specification. A particular order of filter at various scales, α , can be constructed by weighting a number of Erlang frame functions of different order, n , with quite small error. For a desired frequency characteristic, a set of weights can be computed which minimise the error in the least squares sense. In particular, for a given filter family of $h_n^\alpha(r, \theta)$, one can steer (synthesize) a response in α by weighted combination of a set of N filter outputs, all with fixed α .

3.2 Orientation Steerability

Direction steering arises immediately from the use of the $\cos^2(\phi)$ angular tapering in the frequency domain. A theorem due to Simoncelli ([10]) states that exact angular steerability over Fourier half space is provided by 3 rotated filters which span the space $(0, \pi)$. Intermediate orientations can then be synthesized by combining the filter outputs with appropriately chosen weights [2].

3.3 Scale-Orientation Steering Architecture

The software architecture for scale and orientation steering is then as illustrated in Figure (2), and its implementation is described in the following section.




Figure 3: (Left) The 6 steering weights used to sweep the scale parameter from 0.200 to 0.975. (Right) The % error in the synthesized radial frequency characteristic at each scale.

4 Steerable Filter Implementation

4.1 Filter Design

The filters were designed in the discrete Fourier domain, by evaluating the expression $H_n^\omega(\omega)H^\phi(\phi - \phi_k)$, where $\phi_k = \{0, \pi/3, 2\pi/3\}$, at uniform intervals over a 65×65 grid. A two dimensional inverse discrete Fourier transform was then applied to yield the spatial impulse responses. These were truncated to 15×15 complex filters. 6 different radial frequency specifications were used. The parameter sets for these were $\{n_i\}_{i=1..6} = [0, 3, 5, 7, 10, 15]$ and $\{\alpha_i\}_{i=1..6} = \pi/6[1/2, 1, 1, 1, 1, 1]$. The weights for synthesizing arbitrary scale filters for a 4th order filter were computed by minimising the total squared error between a desired scale value and that synthesized from the weighted frame members. The scale setting, α , was varied from 0.200 to 0.975 in steps of 0.050. The 6 weights as a function of α are plotted in Figure (3a), with the approximation error across the scale parameter α , shown in Figure (3b). These weights are directly used in the architecture of Figure (2).

4.2 Computational cost

Much of the computational cost is incurred in the filter bank. Once the "frame images" are computed, the construction of a steered image requires only 9 multiplications and 7 additions per pixel. A very small overhead is also needed for computing the weights for an arbitrary radial filter specification. In this implementation, it is given by a product between a 6×32 matrix, computed once per frame set, and a 32×1 desired radial

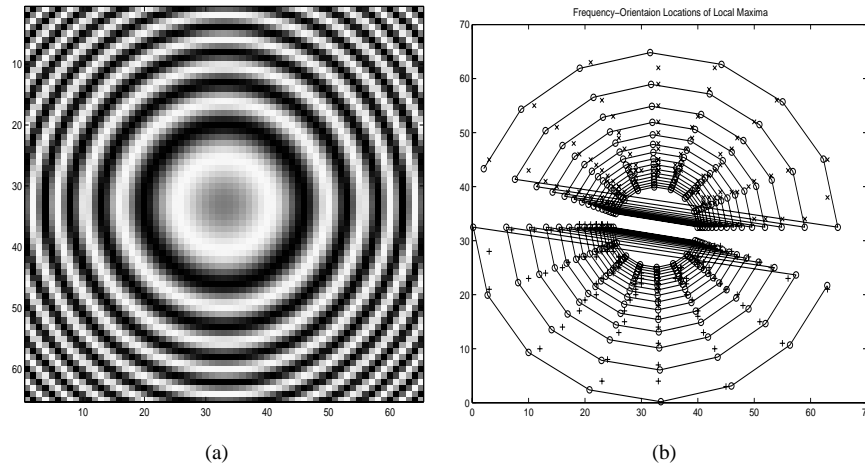


Figure 4: (Left) The test image, generated according to $f(r) = \sin((r/6)^2)$ where r is pixel distance from the origin. (Right) Positions of local maxima (\times) of filtered image and predicted locations (\circ) (desired trajectory).

frequency characteristic.

5 Validation

5.1 Scale Steering

As α is varied, one would expect the central radial frequency to scale as $1/\alpha$. In order to test this, a phantom was generated, consisting of a linear frequency modulation with radial distance from the origin. The phantom is shown in Figure (4a); it contains all possible spatial orientations (within the limits of discretisation error). A trajectory in oriented scale-space was then designed, such that at a range of 16 scales, the synthesized filter was steered through 8 orientations, spaced between 0° and 160° . To facilitate visual interpretation, the same orientations were chosen at each scale. The scale/orientation-space trajectory, using the modulation law and filter characteristics described in Section 2, is plotted in Figure (4b), along with the locations in the filtered image corresponding to the local maxima of synthesized output magnitudes. Some steering error is clearly visible at certain locations corresponding to finer scales. This error is orientation dependent. It is introduced by the truncation error in restricting Fourier space to a 65×65 grid; at angles between the grid axes, the error diminishes.

6 Applications

A test image for orientation and phase estimation consists of a set of 4 rectangular strips with Gaussian intensity profiles across the shorter axis. The strip width variance param-




Figure 5: (Left) (a) The test image. (Right) (b) Pixels belonging to the class of "strip axis", according to the steered magnitude and phase criteria described in the text.

ters are 25.0, 12.5, 5.0 and 2.5 pixel^2 . These are rotated to random angles using bilinear interpolation. Gaussian distributed zero-mean white noise of unit variance is added, and the resulting, heavily corrupted image, is shown in Figure (5a). Using outputs from three orientations at a particular scale, the principal direction of orientation at each pixel location is determined in a least squares sense. At all pixels, the phase relationship between the real and imaginary part of each oriented complex filter is computed. The phase relationship is indicative of a strip-like ("fattened" line) neighbourhood character when the angle is close to zero. A map of filter energy is also produced. Areas with high energy indicate that the structure size of the neighbourhood is of the same order as that of the mask. A simple classification process is applied to all pixels based on the following criteria: a pixel is classed as belonging to the axis of a strip if its local oriented absolute phase is less than 0.2 radians, and its filter energy is greater than a threshold determined heuristically from the histogram of steered output magnitude. The result of this labeling is shown in Figure (5b). A poor result is obtained for the widest Gaussian "strip", and indeed this has a width well beyond that of the filter kernels used. Such structures are best analysed by larger filter kernels, or by a multirate approach.

7 Conclusions

The filters described rely on properties of one-sided Erlang functions to yield complex quadrature filters that are scale-steerable. Orientation steerability of these filters is achieved by standard techniques. A test of the joint modal frequency and orientation steerability was performed by applying the filters to a test image. Results were good, apart from at orientation/frequency locations where mask design introduced errors in filter synthesis. This could be overcome by increasing the retina used to define the frequency-domain masks, or

by an optimisation approach such as in [4]. Orientation of structures at a particular scale was estimated using the oriented scale-space, and the steered phase was extracted from a noisy test image. This was found to contain phase information identifying the medial axes of the strips. Edges could be located using a similar procedure; one advantage of this approach is that all operations are performed in a global fashion, with no pixel-level coding required for feature generation prior to classification. This can be contrasted with, for example, zero-crossing edge-detection, where oriented searches for zero-crossings requires extensive, pixel-level coding. Additionally, the filter phase relationship, an indicator of the pixel's neighbourhood structure, is decoupled from neighbourhood energy. Whilst the scale-steerable nature of this family provides an advantage, comparative studies with well known line/edge filter pair constructions are warranted. This includes comparisons with complex Gabor and first and second derivative of Gaussian kernel pairs, and comparisons with other single-kernel operators in the literature, such as those in [9], which have similar impulse responses to either the real or imaginary part of these complex filters.

References

- [1] J. Bigün, Gösta H. Granlund, and Johan Wiklund. Multidimensional orientation: texture analysis and optical flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(8), 1991.
- [2] William T. Freeman and Edward H. Adelson. The design and use of steerable filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(9):891–906, 1991.
- [3] I.S. Gradshteyn, Iosif Moiseevich, and Alan Jeffrey. *Tables of Integrals, Series and Products: 5th Edition*. Academic Press, 1994.
- [4] Gösta Granlund and Hans Knutsson. *Signal Processing for Computer Vision*. Kluwer Academic Publishers, 1994.
- [5] J.J. Koenderink. The structures of images. *Biological Cybernetics*, 50(5):363–370, 1984.
- [6] Stephane G. Mallat. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, T-PAMI-11:674–693, July 1989.
- [7] Athanassios Papoulis. *Probability, Random Variables and Stochastic Processes*. McGraw-Hill, 1984.
- [8] G.S. Russell. Nested reentrant and recurrent computation in early vision: A bayesian neuro-morphic model applied to hyperacuity. *Biological Cybernetics*, 76(3), March 1997.
- [9] S. Sarkar and K.L. Boyer. Optimal impulse response zero crossing based edge detectors. *CVGIP-IMAGE UNDERSTANDING*, 54(2):224–243, 1991.
- [10] Eero P. Simoncelli, William T. Freeman, Edward H. Adelson, and David J. Heeger. Shiftable multi-scale transforms. *IEEE Transactions on Information Theory*, 38:587–607, March 1992.