



Videre: Journal of Computer Vision Research

Quarterly Journal

Summer 1999, Volume 1, Number 3

The MIT Press

Article 1

Image Features from Phase Congruency

Peter Kovesi

Videre: Journal of Computer Vision Research (ISSN 1089-2788) is a quarterly journal published electronically on the Internet by The MIT Press, Cambridge, Massachusetts, 02142. Subscriptions and address changes should be addressed to MIT Press Journals, Five Cambridge Center, Cambridge, MA 02142; phone: (617) 253-2889; fax: (617) 577-1545; e-mail: journals-orders@mit.edu. Subscription rates are: Individuals \$30.00, Institutions \$125.00. Canadians add additional 7% GST. Prices subject to change without notice.

Subscribers are licensed to use journal articles in a variety of ways, limited only as required to insure fair attribution to authors and the Journal, and to prohibit use in a competing commercial product. See the Journals World Wide Web site for further details. Address inquiries to the Subsidiary Rights Manager, MIT Press Journals, Five Cambridge Center, Cambridge, MA 02142; phone: (617) 253-2864; fax: (617) 258-5028; e-mail: journals-rights@mit.edu.

Image Features from Phase Congruency

Peter Kovesi¹

Image features such as step edges, lines, and Mach bands all give rise to points where the Fourier components of the image are maximally in phase. The use of phase congruency for marking features has significant advantages over gradient-based methods. Phase congruency is a dimensionless quantity that is invariant to changes in image brightness or contrast; hence, it provides an absolute measure of the significance of feature points, thus allowing the use of universal threshold values that can be applied over wide classes of images. This paper presents a new measure of phase congruency and shows how it can be calculated through the use of wavelets. The existing theory that has been developed for 1-D signals is extended to allow the calculation of phase congruency in 2-D images. It is shown that, for good localization, it is important to consider the spread of frequencies present at a point of phase congruency. An effective method for identifying and compensating for the level of noise in an image is presented. Finally, it is argued that high-pass filtering should be used to obtain image information at different scales. With this approach, the choice of scale affects only the relative significance of features without degrading their localization.

Keywords: phase congruency, feature detection, low-level invariance, log Gabor wavelets, noise compensation

1. Department of Computer Science
The University of Western Australia
Nedlands, W.A. 6907
email: pk@cs.uwa.oz.au

Copyright © 1999
Massachusetts Institute of Technology
mitpress.mit.edu/videre.html

1 Introduction

In searching for parameters to describe the significance of image features such as edges, we should be looking for dimensionless quantities—in particular, measures that are invariant with respect to image illumination and magnification. Such quantities would provide an absolute measure of the significance of feature points that could be universally applied to any image irrespective of image illumination and magnification.

Gradient-based edge-detection methods such as those developed by Sobel [30], Marr and Hildreth [20], Canny [2, 3], and others are sensitive to variations in image illumination, blurring, and magnification. The image gradient values that correspond to significant edges are usually determined empirically, although a limited number of efforts have been made to determine threshold values automatically. In his thesis, Canny sets his thresholds on the basis of local estimates of image noise obtained via Weiner filtering. However, the details of setting thresholds on this basis and the effectiveness of this approach are not reported. Canny also introduced the idea of thresholding hysteresis which has proved to be a useful heuristic. Kundu and Pal [18] devised a method of thresholding based on human psychophysical data in which contrast sensitivity varies with overall illumination levels. However, it is hard to provide any concrete guide to the fitting of a model of contrast sensitivity relative to a digitized gray scale of 0–255. More recently, Fleck [7, 8] suggests setting thresholds at some multiple (typically 3 to 5) of the expected standard deviation of the operator output when applied to camera noise. This approach of course requires detailed a priori knowledge of the noise characteristics of any camera used to take an image.

A model of feature perception called the Local Energy Model has been developed by Morrone et al. [26] and Morrone and Owens [25]. This model postulates that features are perceived at points in an image where the Fourier components are maximally in phase. Other work on this model of feature perception can be found in Morrone and Burr [23], Owens et al. [28], Venkatesh and Owens [37], Kovesi [14, 15, 16], Owens [27], Morrone, et al. [24], and Robbins and Owens [31]. A wide range of feature types give rise to points of high phase congruency. These include step edges, line and roof edges, and Mach bands. Morrone and Burr [23] show that this model successfully explains a number of psychophysical effects in human feature perception.

Almost all work done so far has concentrated on finding points of maximal phase congruency by looking for maxima in local energy. However, local energy is a dimensional quantity that depends on local contrast. One is unable to specify beforehand what level of local energy corresponds to a significant feature. Here, we concentrate on the issues in calculating phase congruency itself, a dimensionless measure. Values of phase congruency vary from a maximum of 1 (indicating a very significant feature) down to 0 (indicating no significance). This allows one to specify a threshold to pick out features before an image is seen.

However, so far, phase congruency has not been used successfully for feature detection for the following reasons:

1. The calculation of phase congruency is ill conditioned if all the Fourier components of the signal are very small, or if there is only one (or nearly only one) frequency component present in the signal.
2. Being a normalized quantity, phase congruency is highly sensitive to noise.
3. The existing measure of phase congruency does not provide good localization of features.

This paper addresses these problems and is organized as follows. The existing theory behind the calculation of phase congruency in one-dimensional signals is introduced. It is then shown how phase congruency can be calculated from log Gabor wavelets. The paper then considers the effect of noise in the calculation of phase congruency and develops an effective method for identifying and compensating for these effects. It is shown that, for good localization, it is important to consider the spread of frequencies present at a point of phase congruency. Problems in the localization of blurred features are addressed by developing a new and more sensitive measure of phase congruency. This is followed by the issues involved in extending this theory to 2-D images. The issue of analysis at different scales is then considered, and it is argued that high-pass filtering should be used to obtain image information at different scales instead of the more usually applied low-pass filtering. Finally, some results and the conclusion are presented.

2 Local Energy and Phase Congruency

The local energy model of feature detection postulates that features are perceived at points of maximum phase congruency in an image. For example, when one looks at the Fourier series that makes up a square wave, all the Fourier components are sine waves that are exactly in phase at the point of the step at an angle of 0 or 180 deg. depending on whether the step is upward or downward. At all other points in the square wave, phase congruency is low. Similarly, one finds that phase congruency is a maximum at the peaks of a triangular wave (at an angle of 90 or 270 deg.).

Congruency of phase at *any* angle produces a clearly perceived feature. Figure 2 shows a grating constructed from the series

$$s(x) = \sum_0^n \frac{1}{(2n + 1)} \sin[(2n + 1)x + \phi]$$

where ϕ , the offset at which congruence of phase occurs, is varied from 0 to $\pi/2$.

Figure 1. Construction of square and triangular waveforms from their Fourier series. In both diagrams, the first few terms of the respective Fourier series are plotted with broken lines; the sum of these terms is the solid line. Notice how the Fourier components are all in phase at the point of the step in the square wave, and at the peaks and troughs of the triangular wave.





Figure 2. Interpolation of a step feature to a line feature by continuously varying the angle of congruence of phase from 0 at the top to $\pi/2$ at the bottom. Profiles of this grating corresponding to congruence of phase at 0, $\pi/6$, $\pi/3$, and $\pi/2$ are shown on the right.



Morrone and Owens define the phase congruency function in terms of the Fourier series expansion of a signal at some location x as

$$PC(x) = \max_{\bar{\phi}(x) \in [0, 2\pi]} \frac{\sum_n A_n \cos(\phi_n(x) - \bar{\phi}(x))}{\sum_n A_n} \quad (1)$$

where A_n represents the amplitude of the n th Fourier component, and $\phi_n(x)$ represents the local phase of the Fourier component at position x . The value of $\bar{\phi}(x)$ that maximizes this equation is the amplitude weighted mean local phase angle of all the Fourier terms at the point being considered. Taking the cosine of the difference between the actual phase angle of a frequency component and this weighted mean, $\bar{\phi}(x)$, generates a quantity approximately equal to one minus half this difference squared (the Taylor expansion of $\cos(x) \approx 1 - x^2/2$ for small x). Thus, finding where phase congruency is a maximum is approximately equivalent to finding where the weighted variance of local phase angles, relative to the weighted average local phase, is a minimum.

As it stands, phase congruency is a rather awkward quantity to calculate. As an alternative, Venkatesh and Owens [38] show that points of maximum phase congruency can be calculated equivalently by searching for peaks in the local energy function. The local energy function is defined for a one-dimensional luminance profile, $I(x)$, as

$$E(x) = \sqrt{F^2(x) + H^2(x)},$$


where $F(x)$ is the signal $I(x)$ with its DC component removed, and $H(x)$ is the Hilbert transform of $F(x)$ (a 90 deg. phase shift of $F(x)$). Typically, approximations to the components $F(x)$ and $H(x)$ are obtained by convolving the signal with a quadrature pair of filters. Venkatesh and Owens show that energy is equal to phase congruency scaled by the sum of the Fourier amplitudes; that is,

$$E(x) = PC(x) \sum_n A_n.$$

Thus, the local energy function is directly proportional to the phase congruency function, so peaks in local energy will correspond to peaks in phase congruency.

The relationship between phase congruency, energy, and the sum of the Fourier amplitudes can be seen geometrically in Figure 3. The local Fourier components are plotted as complex vectors adding head to tail. The sum of these components projected onto the real axis represent $F(x)$, the original signal with DC component removed; the projection onto the imaginary axis represents $H(x)$, the Hilbert transform. The magnitude of the vector from the origin to the end point is the total energy, $E(x)$. One can see that $E(x)$ is equal to $\sum_n A_n \cos(\phi_n(x) - \bar{\phi}(x))$.

Figure 3. Polar diagram showing the Fourier components at a location in the signal plotted head to tail. This arrangement illustrates the construction of energy, the sum of the Fourier amplitudes, and phase congruency from the Fourier components of a signal.



Phase congruency is the ratio of $E(x)$ to the overall path length taken by the local Fourier components in reaching the end point. Thus, one can clearly see that the degree of phase congruency is independent of the overall magnitude of the signal. This provides invariance to variations in image illumination and/or contrast.

The calculation of energy from spatial filters in quadrature pairs has been central to many models of human visual perception (for example, those proposed by Heeger [11, 12, 13] and Adelson and Bergen [1]). Other researchers who have studied the use of local energy for feature detection are Perona and Malik [29], Freeman [9], and Ronse [32, 33]. Rosenthaler et al. [34] make a comprehensive study of the behavior of local energy at 2-D image feature points. Wang and Jenkin [39] use complex Gabor filters to detect edges and bars in images. They recognize that step edges and bars have specific local phase properties that can be detected using filters in quadrature; however, they do not connect the significance of high local energy with the concept of phase congruency. It should also be noted that Grossman [10] recognized that wavelets could be used for the detection of discontinuities. He recognized the fact that discontinuities have no intrinsic scale, and this is reflected in the wavelet coefficient values. However, here too the connection with the concept of phase congruency was not made.

While the use of the local energy function to find peaks in phase congruency is computationally convenient, it does not provide a dimensionless measure of feature significance as it is weighted by the sum of the Fourier component amplitudes, which have units lumens. Accordingly, this paper argues that it is phase congruency that we should be computing.

3 Calculating Phase Congruency via Wavelets

In this work, the wavelet transform is used to obtain frequency information local to a point in a signal. We are interested in calculating local frequency and, in particular, phase information in signals. To preserve phase information, linear-phase filters must be used; that is, we must use nonorthogonal wavelets that are in symmetric/antisymmetric quadrature pairs. Here we will follow the approach of Morlet et al. [22], but, rather than using Gabor filters, we prefer to use logarithmic Gabor

functions as suggested by Field [6]. (These are filters having a Gaussian transfer function when viewed on the logarithmic frequency scale. Log Gabor filters allow arbitrarily large bandwidth filters to be constructed while still maintaining a zero DC component in the even-symmetric filter. (A zero DC value cannot be maintained in Gabor functions for bandwidths over one octave.) On the linear frequency scale, the log Gabor function has a transfer function of the form

$$\mathcal{G}(\omega) = e^{\frac{-(\log(\omega/\omega_o))^2}{2(\log(\kappa/\omega_o))^2}},$$

where ω_o is the filter's center frequency. To obtain constant-shape ratio filters¹ the term κ/ω_o must also be held constant for varying ω_o . For example, a κ/ω_o value of 0.75 will result in a filter bandwidth of approximately one octave and a value of 0.55 will result in a two-octave bandwidth.

If we let I denote the signal and M_n^e and M_n^o denote the even-symmetric (cosine) and odd-symmetric (sine) wavelets at a scale n , we can think of the responses of each quadrature pair of filters as forming a response vector,

$$[e_n(x), o_n(x)] = [I(x) * M_n^e, I(x) * M_n^o].$$

The amplitude of the transform at a given wavelet scale is given by

$$A_n(x) = \sqrt{e_n(x)^2 + o_n(x)^2},$$

and the phase is given by

$$\phi_n(x) = \text{atan2}(e_n(x), o_n(x)).$$

At each point x in a signal, we will have an array of these response vectors, one vector for each scale of filter.² These response vectors form the basis of our localized representation of the signal, and they can be used in exactly the same way as Fourier components can be used to calculate phase congruency. This is shown in Figure 4.

The design of the wavelet filter bank needs to be such that the transfer function of each filter overlaps sufficiently with its neighbors so that the sum of all the transfer functions forms a relatively uniform coverage of the spectrum. Note that we wish to retain a broad range of frequencies in our signal, because phase congruency is of interest only if it occurs over a wide range of frequencies.

Referring to Figure 4, we can see that an estimate of $F(x)$ can be formed by summing the even filter convolutions. Similarly, $H(x)$ can be estimated from the odd filter convolutions.


$$\begin{aligned} F(x) &\simeq \sum_n e_n(x), \\ H(x) &\simeq \sum_n o_n(x), \text{ and} \\ \sum_n A_n(x) &\simeq \sum_n \sqrt{e_n(x)^2 + o_n(x)^2}. \end{aligned} \quad (2)$$

With these three components we are able to calculate phase congruency.

1. That is, filters that are all geometric scalings of some reference filter.

2. Note that, from now on, n will be used to refer to wavelet scale. (Previously n denoted frequency in the Fourier series of a signal.)

Figure 4. Calculation of phase congruency from convolution of the signal with quadrature pairs of filters. The convolution output from each quadrature pair of filters at a location in the signal can be considered to represent a response vector having length A_n and phase angle ϕ_n . When the response vectors are plotted head to tail, phase congruency can be seen to be the ratio of the length of the sum of the vectors to the total path length taken by the response vectors in reaching the end point. The noise circle represents the level of $E(x)$ one can expect just from the noise in the signal.



However, there are some problems in the calculation of phase congruency. Referring to Figure 4, one can see the following:

1. The calculation of phase congruency becomes ill conditioned if all the Fourier amplitudes are very small.
2. If the value of $E(x)$ falls within the noise circle (shown in the bottom-right diagram of Figure 4), values of phase congruency lose all significance.
3. If there is only one (or nearly only one) frequency component present in the signal, phase congruency is always one ($\sum_n A_n = E(x)$).
4. The definition of phase congruency as provided by Equation (1) does not provide good localization as this function only varies with the cosine of phase deviation, rather than, say, phase deviation itself.

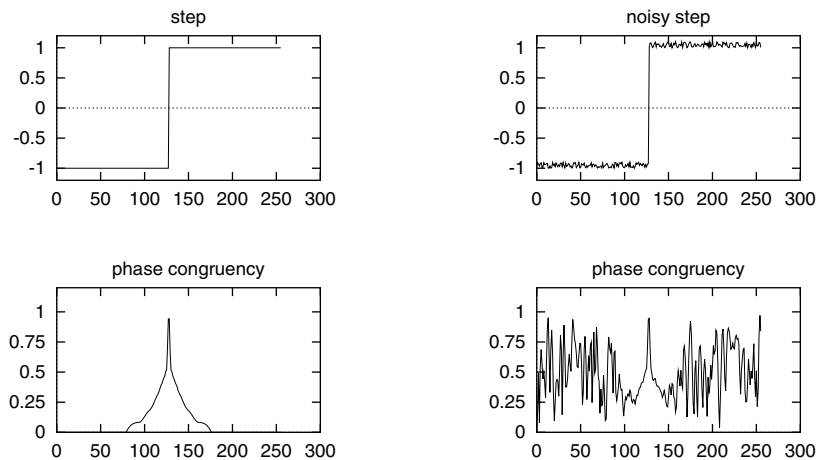
The problem of phase congruency becoming ill conditioned if all the Fourier amplitudes are very small can be addressed by adding a small positive constant, ε , to the denominator of the expression for phase congruency. Thus,

$$PC(x) = \frac{E(x)}{\sum_n A_n(x) + \varepsilon},$$

where $E(x) = \sqrt{F(x)^2 + H(x)^2}$. The appropriate value of ε depends on the precision with which we are able to perform convolutions and other operations on our signal; it does not depend on the signal itself.

The other problems in the calculation of phase congruency outlined above are addressed in the following sections.

Figure 5. Phase congruency of a step function with and without noise.



4 Noise

A difficulty with phase congruency is its response to noise. Figure 5 illustrates the phase congruency of a step function with and without noise. In the vicinity of the step, phase congruency is high only at the point of the step. However, away from the step, the fluctuations due to noise are considered to be significant relative to the surrounding signal (which is noise). This will occur no matter how small the noise level is. This is the price one pays for using a normalized measure such as phase congruency.

If one chooses to calculate local energy rather than phase congruency, one can follow the approach of Morrone, et al. [24], in which the filter scale is varied in an adaptive manner so as to keep the signal-to-noise ratio at some specified value. This approach cannot be applied here as, by definition, it results in information from only one filter scale being available at any point in the signal. The calculation of phase congruency requires the integration and normalization of information over many scales. We must identify the level of noise in the signal and remove its influence in the calculation of phase congruency.

It is possible to estimate the influence of noise in the calculation of $E(x)$ if we make the following three assumptions: image noise is additive, the noise power spectrum is constant, and features (such as edges) occur only at isolated locations in an image.³

In the following discussion, we shall use the following expression for energy

$$E = \sqrt{\left(\sum_n e_n\right)^2 + \left(\sum_n o_n\right)^2},$$

where e_n and o_n are the outputs of the even and odd symmetric filters at scale n . Energy is the magnitude of a vector sum. If our noise is Gaussian with random phase, each vector in this sum will be composed of two independent normally distributed variables. Thus, the distribution of the position of each vector will be a 2-D Gaussian centered on the origin.

3. While these assumptions may be considered simplistic, given the limited (and sometimes conflicting) data on the nature of noise in real images (see for example Fleck [7] and McIvor [21]), it can be argued that one has little basis on which to build a more formal model.

The distribution of the sum of these vectors is obtained by successively convolving the position distributions for the noise vectors at each scale. As these are all 2-D Gaussians, the final distribution of the end position of the energy vector will also be Gaussian. However, what we are interested in is the distribution of the magnitude of the energy vector. This will have a Rayleigh distribution [40] of the form

$$R(x) = \frac{x}{\sigma_G^2} e^{-\frac{x^2}{2\sigma_G^2}},$$

where σ_G^2 is the variance of the Gaussian distribution describing the position of the total energy vector. The mean of the Rayleigh distribution is

$$\mu_R = \sigma_G \sqrt{\frac{\pi}{2}}, \quad (3)$$

and its variance is given by

$$\sigma_R^2 = \frac{4 - \pi}{2} \sigma_G^2. \quad (4)$$

If one can determine an expected value of energy due to noise, we can use this as an estimate of the mean of the energy's Rayleigh distribution, and, hence, determine its variance. A noise threshold can then be set in terms of a specified number of standard deviations above the mean.

Rather than construct the expected value of E , it is more convenient to estimate E^2 . Note that, while E will have a Rayleigh distribution, E^2 will have a χ^2 distribution with two degrees of freedom. The expected value of E^2 will correspond to the second moment of the Rayleigh distribution with respect to the origin:

$$\mathbb{E}(E^2) = 2\sigma_G^2, \quad (5)$$

where \mathbb{E} denotes the expected value.

The expected value for E^2 in terms of our filter responses is

$$\begin{aligned} \mathbb{E}(E^2) &= \mathbb{E}\left(\left(\sum_n e_n\right)^2 + \left(\sum_n o_n\right)^2\right) \\ &= \mathbb{E}\left(\left(\sum_n e_n\right)^2\right) + \mathbb{E}\left(\left(\sum_n o_n\right)^2\right) \\ &\quad + \mathbb{E}\left(2 \sum_{i < j} (e_i e_j + o_i o_j)\right) \\ &= 2\mathbb{E}\left(\left(\sum_n e_n\right)^2\right) + 4\mathbb{E}\left(\sum_{i < j} (e_i e_j)\right), \end{aligned} \quad (6)$$

this last step being possible because the distributions of e_n and o_n are identical, but independent.

Given that e_n is obtained by convolving the noise signal g with a filter M_n , and denoting the Fourier transform $\mathcal{F}(f) = \hat{f}$, we obtain

$$\begin{aligned}
\mathbb{E}(E^2) &= 2\mathbb{E}\left(\sum_n (M_n * g)^2\right) + 4\mathbb{E}\left(\sum_{i<j} (M_i * g) \cdot (M_j * g)\right) \\
&= 2\mathbb{E}\left(\sum_n (\hat{M}_n \cdot \hat{g})^2\right) + 4\mathbb{E}\left(\sum_{i<j} \mathcal{F}^{-1}(\hat{M}_i \cdot \hat{g}) * (\hat{M}_j \cdot \hat{g})\right) \\
&= 2|\hat{g}|^2\mathbb{E}\left(\sum_n \hat{M}_n^2\right) + 4\mathbb{E}\left(\sum_{i<j} \mathcal{F}^{-1}(|\hat{g}|^2 \cdot (\hat{M}_i * \hat{M}_j))\right) \\
&= 2|\hat{g}|^2\mathbb{E}\left(\sum_n M_n^2\right) + 4|\hat{g}|^2\mathbb{E}\left(\sum_{i<j} (M_i \cdot M_j)\right). \tag{7}
\end{aligned}$$

Note, we are assuming that g has zero mean and $|\hat{g}|$ is constant. The components of Equation (7) involving the filters M_n can be evaluated numerically, but what we do not know is the amplitude of the noise spectrum, $|\hat{g}|$. However, we can estimate $|\hat{g}|$ from the response of the smallest scale filter pair in the wavelet bank as follows.

The smallest scale filter has the largest bandwidth and as such will give the strongest noise response. Only at feature points will the response differ from the background noise response, but the regions where it will be responding to features will be small due to the small spatial extent of the wavelet. Thus, the distribution of the squared amplitude response from the smallest scale filter pair across the whole image will be primarily the noise distribution, a scaled 2 DOF χ^2 distribution, with some contamination as a result of the response of the filters to feature points in the image.

We can obtain a robust estimate of the mean of the squared amplitude response of the smallest scale filter via the median response. The median of a 2 DOF χ^2 distribution is the value x such that

$$\begin{aligned}
\int_0^x \frac{1}{2} e^{-\frac{x}{2}} &= \frac{1}{2} \\
\Rightarrow \text{median} &= -2 \ln(1/2).
\end{aligned}$$

Noting that the mean of a 2 DOF χ^2 distribution is 2, we obtain

$$\mathbb{E}(A_N^2) = \frac{-\text{median}(A_N^2)}{\ln(1/2)},$$

where N is the index of the smallest scale filter. This allows us to form the estimate

$$|\hat{g}|^2 \simeq \frac{\mathbb{E}(A_N^2)}{\mathbb{E}(\hat{M}_N^2)}. \tag{8}$$

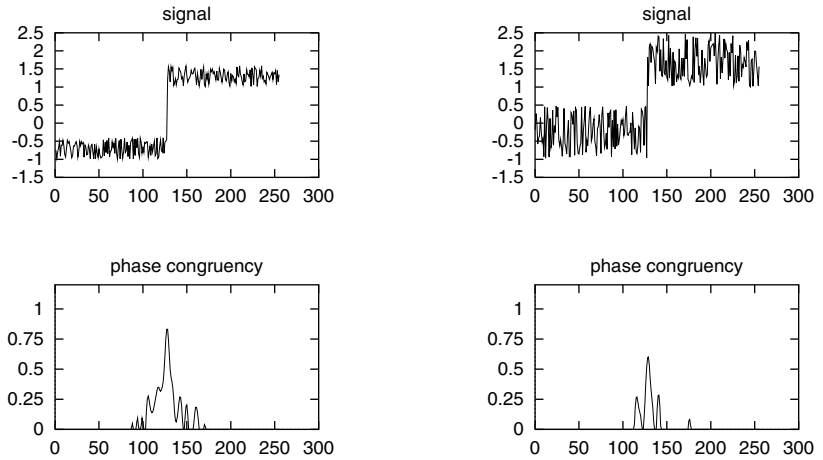
This can be substituted back into Equation (7) to obtain a value for $\mathbb{E}(E^2)$; then, using Equations (5), (4), and (3), we can obtain the mean, μ_R , and variance, σ_R^2 , of the Rayleigh distribution describing the noise energy response.

The radius, T , of the noise circle (shown in Figure 4) is taken to be the mean noise response plus some multiple, k , of σ_R ,

$$T = \mu_R + k\sigma_R, \tag{9}$$

where k is typically in the range 2 to 3. If we subtract this estimated noise effect from the local energy before normalizing it by the sum of the wavelet response amplitudes, we will eliminate spurious responses

Figure 6. Noise-compensated phase congruency of two step profiles.



to noise. Thus, we modify the expression for phase congruency to the following:

$$PC(x) = \frac{[E(x) - T]}{\sum_n A_n(x) + \varepsilon},$$

where $[\]$ denotes that the enclosed quantity is equal to itself when its value is positive, and zero otherwise. This approach to noise compensation has parallels to Donoho's techniques for de-noising via soft thresholding [5].

The phase congruency of a legitimate feature will be reduced according to the magnitude of the noise's local energy relative to the feature. Thus, we end up with a measure of confidence that the feature is significant relative to the level of noise. Figure 6 shows the results of processing two noisy step profiles. In both cases, a k value of 3 was used to estimate the maximum influence of noise on local energy.

5 The Importance of Frequency Spread

Clearly, a point of phase congruency is significant only if it occurs over a wide range of frequencies. In the degenerate case where there is only one frequency component (a pure sine wave), phase congruency will be one everywhere. A more common situation is where a feature has undergone Gaussian smoothing. The smoothing reduces the high-frequency components in the signal and reduces the frequency spread. In the extreme, the frequency spread is reduced so much that locally we approach the situation that arises with pure sine functions.

To counter this problem, one can incorporate low-frequency components in the calculation of phase congruency. These low-frequency components are the least affected by any smoothing of the signal. Even so, the values of energy and $\sum_n A_n$ can still be nearly equal over an extended region about the feature, producing a poorly localized phase congruency response. Thus, as a measure of feature significance, phase congruency should be weighted by some measure of the spread of frequencies present. What then is a significant distribution of frequencies? If we consider some common feature types such as the square waveform (step edge), the triangular waveform (roof edge), and the delta function (line feature) as some of the "edgiest" waveforms imaginable, we can use their frequency distributions as a guide to the ideal.

The power spectrum of a square wave is of the form $1/\omega^2$. Each of the wavelets that we use to analyze the signal gathers power from

geometrically increasing bands of the spectrum. The net result is that filter responses are constant, independent of filter center frequency. Hence, the expected distribution of amplitude responses to a step feature will be a uniform one. Field [6] points out that, in many cases, images of natural scenes have overall spectral distributions that fall off inversely proportional to frequency, and for this reason he also advocates the use of geometrically scaled filter banks. Under these conditions, filters at all scales will, on average, be responding with equal magnitudes, which is likely to maximize the precision of any computation (numerical or neural) that we make with the filter outputs.

The other important feature types we must consider are the delta function (corresponding to line features) and roof edges. The power spectrum of a delta function is uniform. One can show that, for a delta feature, the amplitude of the wavelet filter responses will be proportional to their bandwidths, and hence their center frequencies. This will give a distribution of filter responses strongly skewed to the high-frequency end. On the other hand, for a triangular waveform where all the features are roof edges, the power spectrum falls off at $1/\omega^4$, resulting in a distribution of filter amplitude responses that is strongly skewed to the low-frequency end.

Thus, the difficulty we face here is that there is no one ideal distribution of filter responses. All we can say is that the distribution of filter responses should not be too narrow in some general sense. We can also say that a uniform distribution is of particular significance as step discontinuities are common in images.

Accordingly, we can construct a weighting function that devalues phase congruency at locations where the spread of filter responses is narrow. A measure of filter response spread can be generated by taking the sum of the amplitudes of the responses and dividing by the highest individual response to obtain some notional “width” of the distribution. If this is then normalized by the number of scales being used, we obtain a fractional measure of spread that varies between 0 and 1. This spread is given by

$$s(x) = \frac{1}{N} \left(\frac{\sum_n A_n(x)}{\varepsilon + A_{max}(x)} \right),$$

where N is the total number of scales being considered, $A_{max}(x)$ is the amplitude of the filter pair having maximum response at x , and ε is used to avoid division by zero and to discount the result should both $\sum A_n(x)$ and $A_{max}(x)$ be very small.

A phase congruency weighting function can then be constructed by applying a sigmoid function to the filter response spread value, namely


$$W(x) = \frac{1}{1 + e^{\gamma(c-s(x))}},$$

where c is the cut-off value of filter response spread below which phase congruency values become penalized, and γ is a gain factor that controls the sharpness of the cutoff. Note that the sigmoid function has been merely chosen for its simplicity and ease of manipulation.

Thus,

$$PC(x) = \frac{W(x)[E(x) - T]}{\sum_n A_n(x) + \varepsilon}.$$

Figure 7. Frequency spread weighting function with a cut-off value of 0.4 and γ value of 10.



Weighting by frequency spread—as well as reducing spurious responses where the frequency spread is low—has the additional benefit of sharpening the localization of features, especially those that have been smoothed.

6 A New Measure of Phase Congruency

Even with the addition of frequency spread weighting, one finds that the localization of phase congruency remains poor on blurred features. The reason for this is evident when one studies the expression for energy. Energy is proportional to the cosine of the deviation of phase angle, $\phi_n(x)$ from the overall mean phase angle, $\bar{\phi}(x)$. While the cosine function is maximized when $\phi_n(x) = \bar{\phi}(x)$, it requires a significant difference between $\phi_n(x)$ and $\bar{\phi}(x)$ before its value falls appreciably. For example, the filter outputs could be such that all phase angles were $\bar{\phi}(x) \pm 25$ deg. and we would still have a phase congruency of approximately 0.9. Thus, using the cosine of the phase deviation is a rather insensitive measure of phase congruency.

We can construct a more sensitive measure of phase congruency by noting that, at a point of phase congruency, the cosine of the phase deviation should be large and the absolute value of the sine of the phase deviation should be small. The gradient of the sine function is a maximum at the origin. Therefore, making use of the sine of the phase deviation will increase our sensitivity. Accordingly, a more sensitive phase deviation function on which to base the calculation of phase congruency is


$$\Delta\Phi(x) = \cos(\phi_n(x) - \bar{\phi}(x)) - |\sin(\phi_n(x) - \bar{\phi}(x))|. \quad (10)$$

Figure 8 plots this function along with the cosine function for comparison. The function falls very nearly linearly as phase deviation moves from 0 to $\pm\pi/2$. Thus, a near-direct measure of phase deviation is obtained without having to resort to inverse trigonometric functions.

Using this new measure of phase deviation, $\Delta\Phi(x)$, a new measure of phase congruency can be defined as

$$PC_2(x) = \frac{\sum_n W(x)[A_n(x)\Delta\Phi_n(x) - T]}{\sum_n A_n(x) + \varepsilon},$$

Figure 8. Comparison between $\cos(x)$ (dotted line) and $\cos(x) - |\sin(x)|$ (solid line).



where, as before, ε is a small constant to avoid division by zero and T is the estimated noise influence⁴. Note that this expression for phase congruency is $PC_2(x)$ to distinguish it from the previous definition of phase congruency that will now be referred to as $PC_1(x)$.

The relationship between these two phase congruency measures can be seen in statistical terms. The measure $PC_2(x)$ is related to the weighted mean absolute deviation of phase from the weighted mean⁵ in that the phase deviation measure it uses varies almost linearly with angular deviation of phase. On the other hand, $PC_1(x)$ —in using the cosine of the phase deviation—is related to the approximate weighted variance with respect to the weighted mean phase.

The calculation of this new measure of phase congruency, $PC_2(x)$, can be done using dot and cross products between the filter output response vectors to calculate the cosine and sine of $(\phi_n(x) - \bar{\phi}(x))$. The unit vector representing the direction of the weighted mean phase angle, $\bar{\phi}(x)$ is given by

$$(\bar{\phi}_e(x), \bar{\phi}_o(x)) = \frac{1}{\sqrt{(F(x)^2 + H(x)^2)}}(F(x), H(x)). \quad (11)$$

Now, using dot and cross products one can form the quantities:

$$A_n(x) \cos(\phi_n(x) - \bar{\phi}(x)) = e_n(x) \cdot \bar{\phi}_e(x) + o_n(x) \cdot \bar{\phi}_o(x), \quad (12)$$

$$A_n(x) |\sin(\phi_n(x) - \bar{\phi}(x))| = |e_n(x) \cdot \bar{\phi}_o(x) - o_n(x) \cdot \bar{\phi}_e(x)|. \quad (13)$$

Thus,

$$A_n(x)(\cos(\phi_n(x) - \bar{\phi}(x)) - |\sin(\phi_n(x) - \bar{\phi}(x))|) = (e_n(x) \cdot \bar{\phi}_e(x) + o_n(x) \cdot \bar{\phi}_o(x)) - |e_n(x) \cdot \bar{\phi}_o(x) - o_n(x) \cdot \bar{\phi}_e(x)|, \quad (14)$$

which gives us the quantity needed to calculate this new version of phase congruency.

7 Extension to Two Dimensions

So far our discussion has been limited to signals in one dimension. Calculation of phase congruency requires the formation of a 90 deg. phase shift of the signal which we have done using odd-symmetric filters. As one cannot construct rotationally symmetric odd-symmetric filters, one is forced to analyze a two-dimensional image by applying our one-dimensional analysis over several orientations and combining the results in some way. Three issues must be resolved: the shape of the filters in two dimensions, the numbers of orientations to analyze, and the way in which the results from each orientation are combined.

7.1 2-D Filter Design

The one-dimensional filters described previously can be extended into two dimensions by simply applying some spreading function across the

4. This equation for phase congruency does not lend itself readily to the noise analysis described in Section 4. In practice, it is found that the analysis used in Section 4 approximately applies to the PC_2 measure, but the noise effect is typically overestimated. This can be compensated for by rescaling the value for T calculated in Equation (9) by a factor of 0.5 to 0.7.

5. Normally the mean absolute deviation is calculated with respect to the median of a distribution; the median minimizes this quantity. However, the mean of the phase distribution is more accessible to us than the median, especially if we want to weight phase values by amplitude values.

filter perpendicular to its orientation. Such a 2-D filter is separable; image convolution can be accomplished by a 1-D convolution with the spreading function, followed by a 1-D convolution in the orthogonal direction with the wavelet function. Since we are interested in the phase information, the important thing to ensure is that convolution with the spreading function does not corrupt the phase data in the image.

The obvious spreading function to use is the Gaussian, and there are good reasons for choosing it. Any function smoothed with a Gaussian undergoes amplitude modulation of its components, but phase is unaffected. Thus, the phase congruency at any features will be preserved. If, on the other hand, we were to, say, use a rectangular spreading function, some phase angles in the signal would be reversed because the transfer function (a sine function) has negative values. Phase congruency at features would then be corrupted.

7.2 Filter Orientations

To detect features at all orientations, our bank of filters must be designed so that they tile the frequency plane uniformly. In the frequency plane, the filters appear as 2-D Gaussians symmetrically or antisymmetrically placed around the origin, depending on the spatial symmetry of the filters. The length-to-width ratio of the 2-D wavelets controls their directional selectivity. This ratio can be varied in conjunction with the number of filter orientations used in order to achieve an even coverage of the 2-D spectrum.

A logical way to construct 2-D filters in the frequency domain is to use polar-separable 2-D Gaussians. In the radial direction, along the frequency axis, the filters are designed in the same way as we have been designing 1-D filters (that is, log Gaussians with geometrically increasing center frequencies and bandwidths). In the angular direction, the filters have Gaussian cross-sections, where the ratio between the standard deviation and the angular spacing of the filters is some constant. This ensures a fixed length-to-width ratio of the filters in the spatial domain. Thus, the cross-section of the transfer function in the angular direction is

$$G(\theta) = e^{-\frac{(\theta-\theta_0)^2}{2\sigma_\theta^2}},$$

where θ_0 is the orientation angle of the filter, and σ_θ is the standard deviation of the Gaussian spreading function in the angular direction. This is set to be some fixed ratio of the orientation spacing between the filters to ensure even coverage of the spectrum in all orientations. A filter orientation spacing of 30 deg. has been found to provide a good compromise between the need to achieve an even spectral coverage while minimizing the number of orientations. The use of more filter orientations does not change the quality of the results significantly. The final arrangement of filters results in a rosette of overlapping polar-separable 2-D Gaussians in the frequency plane. Simoncelli et al. [35] describe a systematic filter design technique for achieving uniform coverage of the frequency plane that could also be applied here.

7.3 Combining Data over Several Orientations

The important issue here is to ensure that features at all possible orientations are treated equally, and all possible conjunctions of features (such as corners and “T” junctions) are treated uniformly. Indeed, we want

to avoid making any assumptions about the 2-D form of features that one may encounter. It is also important that the normalization of energy to form phase congruency is done after summing energies over all orientations. We want the final result to be a weighted normalized value, with the result from each orientation contributing to the final result in proportion to its energy.

The approach that has been adopted is as follows. At each location in the image, we calculate energy, $E(x)$, in each orientation, compensate for the influence of noise, apply the weighting for frequency spread, and then form the sum over all orientations. This sum of energies is then normalized by dividing by the sum over all orientations and scales of the amplitudes of the individual wavelet responses at that location in the image. This produces the following equation for 2-D phase congruency:

$$PC_2(x) = \frac{\sum_o \sum_n W_o(x) [A_{no}(x) \Delta \Phi_{no}(x) - T_o]}{\sum_o \sum_n A_{no}(x) + \varepsilon}, \quad (15)$$

where o denotes the index over orientations. Notice in the equation above that noise compensation is performed in each orientation independently. The perceived noise content as deduced from the average power response of the smallest scale wavelet pair can vary significantly with orientation due to the correlation in noise along scan lines that can occur in the digitization of an image.

8 Scale via High-Pass Filtering

The traditional approach to analyzing an image at different scales is to consider various low-pass or band-passed versions of the image. Versions of the image having only low frequencies left are considered to contain the “broad scale” features. This approach is inspired from the presence of receptive fields in the visual cortex that act as band-pass filters [19]. While this approach is intuitive, the justification for assuming the brain uses these band-passed versions of the image directly for multiscale analysis is perhaps somewhat circular. On being presented with a low-pass version of an image, one is asked “What features do you see?” Of course, you see the broad scale features: they are the only things left in the image to be seen.

A major problem with the use of low- or band-pass filtering for multiscale analysis is that the number of features present in an image, and their locations, vary with the scale used. It seems very unsatisfactory for the location of a feature to depend on the scale at which it is analyzed.

The use of phase congruency to measure feature significance allows one to consider an alternative interpretation of feature scale. Phase congruency at some point in a signal depends on how the feature is built up from the local frequency components. Depending on the size of the analysis window, features some distance from the point of interest may contribute to the local frequency components considered to be present. Thus, features are not considered in isolation but in context with their surrounding features.

Therefore, as far as phase congruency is concerned, the natural scale parameter to vary is the size of the window in the image over which we perform the local frequency analysis. In the context of our use of wavelets to calculate phase congruency, the scale of analysis is specified by the spatial extent of the largest filter in the wavelet bank. With this approach, we are using high-pass filtering to specify the analysis scale. We cut out low-frequency components (those having wavelengths larger

than the window size) while leaving the high-frequency components intact.

If we use a small analysis window, each feature will be treated with a great degree of independence from other features in the image. We will be comparing each feature to only a small number of other features that are nearby; hence, each feature is likely to be perceived more important locally. At the largest scale (window size equal to image size), each feature is considered in relation to all other features, and we obtain a sense of global significance for each feature.

It should be noted that our original ideal that the significance of image features should be invariant to image magnification is not really attainable. In practice, we have to compute phase congruency using a finite number of spatial filters that cover a limited range of the spectrum. Changing the magnification of an image may alter the relative responses of individual filters and hence change the perceived phase congruency. However, the changes in measured phase congruency will, in general, be much smaller than any corresponding changes in intensity gradient.

In summary, it is proposed that multiscale analysis be done by considering phase congruency of differing high-passed versions of an image. The high-pass images are constructed from the sum of band-passed images, with the sum ranging from the highest frequency band down to some cutoff frequency. With this approach, no matter what scale we consider, all features are localized precisely and in a stable manner. There is no drift of features that occurs with low-pass filtering. All that changes with analysis at different scales is the relative significance of features. This provides an ideal environment for the use of coarse-to-fine strategies. Figure 9 illustrates a one-dimensional signal at two different scales of band-pass and high-pass filtering, along with phase congruency at the two high-pass scales.

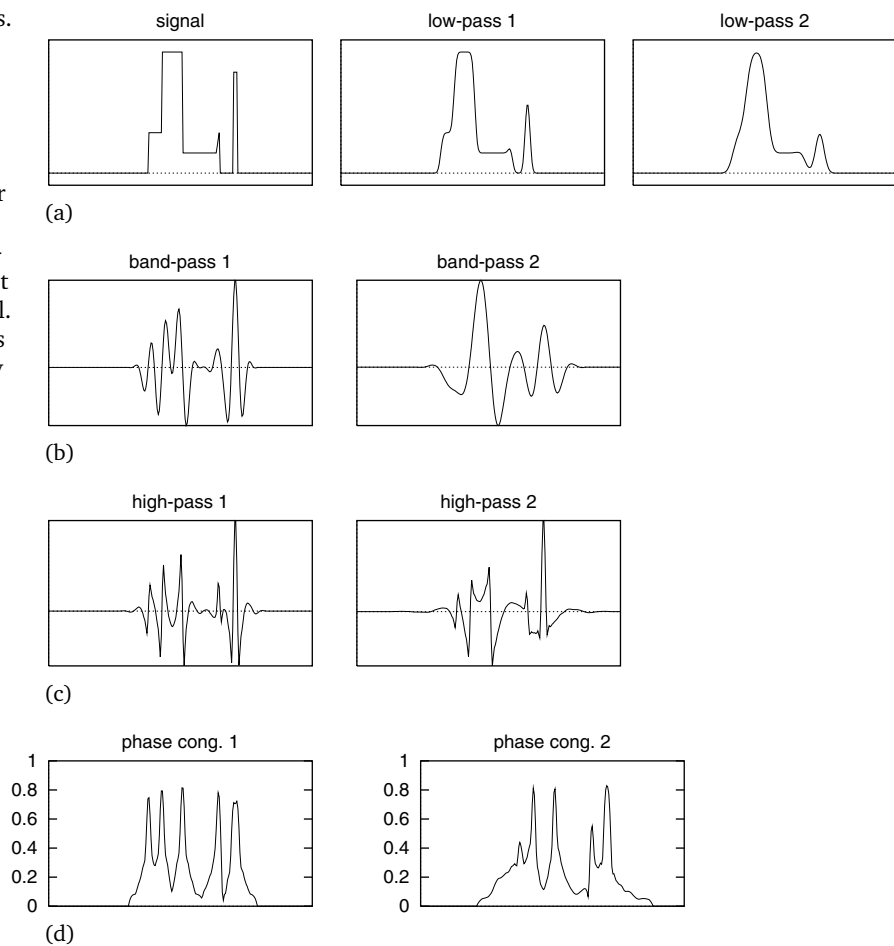
9 Experimental Results

A problem in discussing the performance of a feature detector is devising a sensible form of evaluation. Performance criteria have been used by a number of researchers to design edge operators, notably Canny [2, 3], Spacek [36], and Deriche [4]. These criteria generally measure the ability of a detector to produce a distinct local maximum at the point of a step discontinuity in the presence of noise. However, these criteria are limited in their usefulness, as they are concerned only with specific feature types, usually step discontinuities, and they are not concerned with the absolute value of the resulting maxima in the detector's output. They provide no guide as to one's ability to set general thresholds. A feature detector is of limited use if one does not know in advance what level of response corresponds to a significant feature.

One of the primary motivations for using phase congruency to detect image features is that it provides an absolute measure of the significance of features. This allows one to set thresholds that are applicable across wide classes of images. The other motivation for detecting features on the basis of phase congruency is that we are not required to make any assumptions about the luminance profile of the feature; we are simply looking for points where there is order in the frequency domain. Step discontinuities, lines, and roof edges are all detected.

The performance of the phase congruency detector is illustrated on two test images and on five real images on the following pages. For

Figure 9. Analysis at different scales. Note that the number and location of features as measured by phase congruency remains constant and that only their relative significance varies with scale. Under low-pass and band-pass filtering, the number and locations of features vary. (a) A 1-D signal and two different low-pass versions of it. (b) Two different band-pass versions of the 1-D signal. (c) Two different high-pass versions of the signal. (d) Phase congruency at both scales of high-pass filtering shown in (c).



comparison, the output of the Canny detector is also presented. The implementation of the Canny detector used here follows the modifications suggested by Fleck [8]. The raw, gradient magnitude image is displayed so that comparison can be made without having to consider any artifacts that may be introduced by nonmaximal suppression and thresholding processes. The purpose of providing this comparison is to illustrate some of the qualitative differences in performance between the two detectors. Quantitative comparisons are difficult because the design objectives of the two detectors are completely different. One is seeking to localize step edges, and the other is seeking to identify points of phase congruency.

9.1 Parameters

It should be emphasized that all the results presented on the following pages were obtained by applying the same parameter and threshold values to every image.

The raw phase congruency images were obtained by applying Equation (15) to the images with the following parameters. Local frequency information was obtained using two-octave bandwidth filters over four scales and six orientations. The wavelength of the smallest scale filters was 3 pixels, the scaling between successive filters was 2. Thus, over the four scales, the filter wavelengths were 3, 6, 12, and 24 pixels. The filters were constructed directly in the frequency domain as polar-separable

functions: a logarithmic Gaussian function in the radial direction and a Gaussian in the angular direction. In the angular direction, the ratio between the angular spacing of the filters and angular standard deviation of the Gaussians was 1.2. This results in a coverage of the spectrum that varies by less than 1%. A noise compensation k value of 2.0 was used. The frequency spread weighting function cutoff fraction, c , was set at 0.4, and the gain parameter, γ , was set at 10. The value of ε , the small constant used to prevent division by zero in the case where local energy in the image becomes very small, was set at 0.01. None of these parameter values are particularly critical. However, it is worth noting that the spatial extent of log Gabor filters appears to be minimized when they are constructed with a bandwidth of approximately two octaves [16].

The phase congruency feature maps were obtained by performing nonmaximal suppression on the raw phase congruency images followed by hysteresis thresholding with upper and lower hysteresis threshold values fixed at phase congruency values of 0.3 and 0.15.

MATLAB code for performing the calculation of phase congruency, nonmaximal suppression, and hysteresis thresholding is provided by Kovési [17] for those wishing to replicate the results presented here.

9.2 Discussion

The main qualitative difference between the two detectors is the wide range of response values from the Canny detector. For example, with the Canny detector, the low-contrast square in the circle at the top right of the first test image almost disappears, whereas under phase congruency it is marked prominently. This wide range of responses from the Canny detector makes threshold selection difficult. The other obvious difference is that the Canny detector produces responses on each side of line features, whereas the phase congruency detector produces a response centered on the line. (This problem was recognized by Canny, and he designed a separate operator to detect line features.) One problem with the phase congruency detector (or at least this implementation of it) is its behavior at junctions of features having greatly different magnitudes. Notice how the horizontal edges in the gray scale on the left-hand side of the first test image fade as they meet the strong vertical edge of the gray scale. At the junction between the low-magnitude horizontal edges and the high-magnitude vertical edge, the normalizing component of phase congruency, $\sum_n A_n$, is dominated by the magnitude of the vertical edge. Thus, at the junction, the significance of the horizontal edges relative to the vertical one is downgraded. This problem could possibly be overcome through a different approach to combining phase congruency information over several orientations. Freeman [9], in his use of a normalized form of energy for feature detection, encountered this same problem at junctions of varying contrast. His solution was to normalize energy in each orientation independently using only energy responses from filters in the same orientation. The energy values used for the normalization were blurred spatially in a direction perpendicular to the orientation being considered. This approach has not been adopted here. Robbins and Owens [31] also provide a detailed study of the detection of 2-D features, such as junctions and corners, using local energy.

10 Conclusion

Phase congruency is a dimensionless measure of feature significance. It provides an invariant measure of the significance of feature points in an image, and this allows constant threshold values to be applied across wide classes of images. Thresholds can be specified in advance; they do not have to be determined empirically for individual images.

The theory behind the calculation of phase congruency has been extended in a number of ways. This paper develops a new measure of phase congruency that provides greater localization of features. It is shown how phase congruency can be calculated via log Gabor wavelets, and the problems involved in extending the calculation of phase congruency from 1-D signals to 2-D are addressed. It is shown that, for a normalized measure of feature significance (such as phase congruency), it is crucial to be able to recognize the level of noise in an image and to compensate for it. An effective method of compensation is presented that requires only that the noise power spectrum be approximately constant.

Also presented is the importance of weighting phase congruency by some measure of the spread of the frequencies that are present at each point in an image. This prevents false positives being marked where the frequency spread is very narrow. It also improves the localization of features. While it is not possible to specify one ideal distribution of filter response amplitudes with frequency, it is shown that, when geometrically scaled filters are used, a uniform distribution of responses is a particularly significant one. This distribution matches typical spectral statistics of images and corresponds to the distribution that arises at step discontinuities.

Another contribution of this work is to offer a new approach to the concept of scale in image analysis. The natural scale parameter to vary in the calculation of phase congruency is the size of the analysis window over which to calculate local frequency information. Thus, under these conditions, scale is varied using high-pass filtering rather than low-pass or band-pass filtering. The significant advantage of this approach is that feature locations remain constant over scale, and only their significance relative to each other varies.

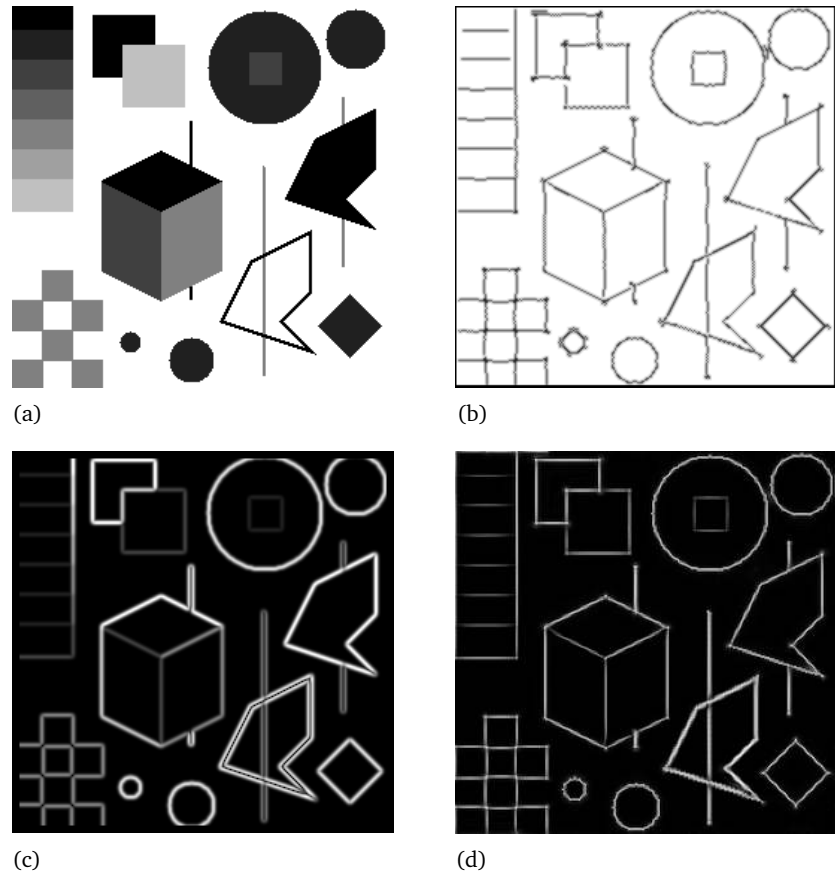
Acknowledgment

The author would like to acknowledge many useful discussions with John Ross, James Trevelyan, Robyn Owens, Ben Robbins, Chris Pudney, and Concetta Morrone. I am indebted to Miguel Mulet-Parada for pointing out some errors in the noise-compensation section, and to Adrian Baddeley for his help in fixing them.

References

- [1] E. H. Adelson and J. R. Bergen. Spatiotemporal energy models for the perception of motion. *Journal of the Optical Society of America A*, 2(2):284–299, 1985.
- [2] J. F. Canny. Finding edges and lines in images. Master's thesis, MIT. AI Lab. TR-720, 1983.
- [3] J. F. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 8(6):112–131, 1986.
- [4] R. Deriche. Using Canny's criteria to derive an optimal edge detector recursively implemented. *The International Journal of Computer Vision*, 1:167–187, April 1987.

Figure 10. Test1 image. Notice how phase congruency marks lines with a single response, not two as the Canny operator does, and that the magnitude of the phase congruency response is largely independent of local contrast. (a) original image; (b) phase congruency feature map; (c) Canny edge strength image; (d) raw phase congruency image.



- [5] D. L. Donoho. De-noising by soft thresholding. Technical Report 409, Department of Statistics. Stanford University, 1992.
- [6] D. J. Field. Relations between the statistics of natural images and the response properties of cortical cells. *Journal of The Optical Society of America A*, 4(12):2379–2394, December 1987.
- [7] Margaret M. Fleck. Multiple widths yield reliable finite differences. *IEEE T-PAMI*, 14(4):412–429, April 1992.
- [8] Margaret M. Fleck. Some defects in finite-difference edge finders. *IEEE T-PAMI*, 14(3):337–345, March 1992.
- [9] W. T. Freeman. *Steerable Filters and Local Analysis of Image Structure*. PhD thesis, MIT Media Lab. TR-190, June 1992.
- [10] A. Grossman. Wavelet transforms and edge detection. In S. Albeverio, P. Blanchard, M. Hazewinkel, and L. Streit (eds) *Stochastic Processes in Physics and Engineering*, pages 149–157. D. Reidel Publishing Company, 1988.
- [11] D. J. Heeger. Optical flow from spatiotemporal filters. In *Proceedings, 1st International Conference on Computer Vision*, pages 181–190, London, June 1987.
- [12] D. J. Heeger. Optical flow using spatiotemporal filters. *International Journal of Computer Vision*, 1:279–302, 1988.
- [13] D. J. Heeger. Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, 9:181–197, 1992.
- [14] P. D. Kovesi. A dimensionless measure of edge significance. In *The Australian Pattern Recognition Society, Conference on Digital*

Figure 11. Illustration of noise compensation. In this example, additive Gaussian noise with a standard deviation of 40 gray levels has been applied to a test image of 256 gray levels. The noise has been successfully ignored in the smooth regions of the image although the confidence of the measured phase congruency at features has been reduced. (a) Test2 image with additive Gaussian noise; (b) raw phase congruency on noise-free image; (c) Canny edge strength on noisy image; (d) raw phase congruency on noisy image.





Image Computing: Techniques and Applications, pages 281–288, Melbourne, 4–6 December 1991.

- [15] P. D. Kovesi. A dimensionless measure of edge significance from phase congruency calculated via wavelets. In *First New Zealand Conference on Image and Vision Computing*, pages 87–94, Auckland, August 1993.
- [16] P. D. Kovesi. *Invariant Measures of Image Features From Phase Information*. PhD thesis, The University of Western Australia, May 1996.
- [17] P. D. Kovesi. MATLAB code for calculating phase congruency and phase symmetry/asymmetry, April 1996. <http://www.cs.uwa.edu.au/~pk/Research/research.html>.
- [18] M. K. Kundu and S. K. Pal. Thresholding for edge detection using human psychovisual phenomena. *Pattern Recognition Letters*, 4:433–411, 1986.
- [19] D. Marr. *Vision*. Freeman: San Francisco, 1982.
- [20] D. Marr and E. C. Hildreth. Theory of edge detection. *Proceedings of the Royal Society, London B*, 207:187–217, 1980.
- [21] A. M. McIvor. A test of camera noise models. In *British Machine Vision Conference*, pages 355–359, Oxford, 1990.
- [22] J. Morlet, G. Arens, E. Fourgeau, and D. Giard. Wave propagation and sampling theory - Part II: Sampling theory and complex waves. *Geophysics*, 47(2):222–236, February 1982.

Figure 12. Mandrill image. This image is largely composed of line features, and this highlights the difference between phase congruency and first-derivative edge operators. The Canny detector marks edges around all the hairs, while phase congruency marks the hairs directly as line features. (a) original image; (b) phase congruency feature map; (c) Canny edge strength image; (d) raw phase congruency image.



- [23] 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.
- [24] M. C. Morrone, A. Navangione, and D. Burr. An adaptive approach to scale selection for line and edge detection. *Pattern Recognition Letters*, 16:667–677, 1995.
- [25] M. C. Morrone and R. A. Owens. Feature detection from local energy. *Pattern Recognition Letters*, 6:303–313, 1987.
- [26] M. C. Morrone, J. R. Ross, D. C. Burr, and R. A. Owens. Mach bands are phase dependent. *Nature*, 324(6094):250–253, November 1986.
- [27] R. A. Owens. Feature-free images. *Pattern Recognition Letters*, 15:35–44, 1994.
- [28] R. A. Owens, S. Venkatesh, and J. Ross. Edge detection is a projection. *Pattern Recognition Letters*, 9:223–244, 1989.
- [29] P. Perona and J. Malik. Detecting and localizing edges composed of steps, peaks and roofs. In *Proceedings of 3rd International Conference on Computer Vision*, pages 52–57, Osaka, 1990.
- [30] K. K. Pringle. Visual perception by a computer. In A. Grasselli (ed.), *Automatic Interpretation and Classification of Images*, pages 277–284. Academic Press, New York, 1969.
- [31] B. Robbins and R. Owens. 2-D feature detection via local energy. *Image and Vision Computing*, 15(5):353–368, May 1997.

Figure 13. Goldhill image. This image illustrates the ability of the PC_2 measure to pick out fine features. The window panes and roof tiles in the nearer houses are clearly marked. (a) original image; (b) phase congruency feature map; (c) Canny edge strength image; (d) raw phase congruency image.



- [32] C. Ronse. On idempotence and related requirements in edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(5):484–491, May 1993.
- [33] C. Ronse. The phase congruency model for edge detection in multi-dimensional pictures. Technical report, Département d'Informatique, Université Louis Pasteur, Strasbourg, France. Report 97/16, October 1997.
- [34] L. Rosenthaler, F. Heitger, O. Kubler, and R. von der Heydt. Detection of general edges and keypoints. In *ECCV92, Springer-Verlag Lecture Notes in Computer Science*, volume 588, pages 78–86. Santa Margherita Ligure, Italy, Springer-Verlag, May 1992.
- [35] E. P. Simoncelli, W. T. Freeman, Ed. H. Adelson, and D. J. Heeger. Shiftable multiscale transforms. *IEEE Transactions on Information Theory*, 38(2):587–607, March 1992.
- [36] L. A. Spacek. *The Detection of Contours and their Visual Motion*. PhD thesis, University of Essex at Colchester, December 1985.
- [37] S. Venkatesh and R. Owens. On the classification of image features. *Pattern Recognition Letters*, 11:339–349, 1990.
- [38] S. Venkatesh and R. A. Owens. An energy feature detection scheme. In *The International Conference on Image Processing*, pages 553–557, Singapore, 1989.
- [39] Z. Wang and M. Jenkin. Using complex Gabor filters to detect and localize edges and bars. In Colin Archibald and Emil Petriu (eds.),

Figure 14. Venice image. Here the value of phase congruency's invariance to contrast can be seen. Notice how some of the buildings along the canal are partly in shadow. The output of the Canny detector almost disappears in the shadowed regions. However, phase congruency successfully picks out many of the features in these regions. (a) original image; (b) phase congruency feature map; (c) Canny edge strength image; (d) raw phase congruency image.



(a)



(b)



(c)



(d)

Advances in Machine Vision: Strategies and Applications, World Scientific Series in Computer Science, volume 32, pages 151–170. Singapore, World Scientific Press, 1992.

[40] E. W. Weisstein. *The CRC Concise Encyclopedia of Mathematics*. CRC Press, 1998.

Figure 15. Boat image. Note the faithful recording of the rigging by phase congruency. (a) original image; (b) phase congruency feature map; (c) Canny edge strength image; (d) raw phase congruency image.

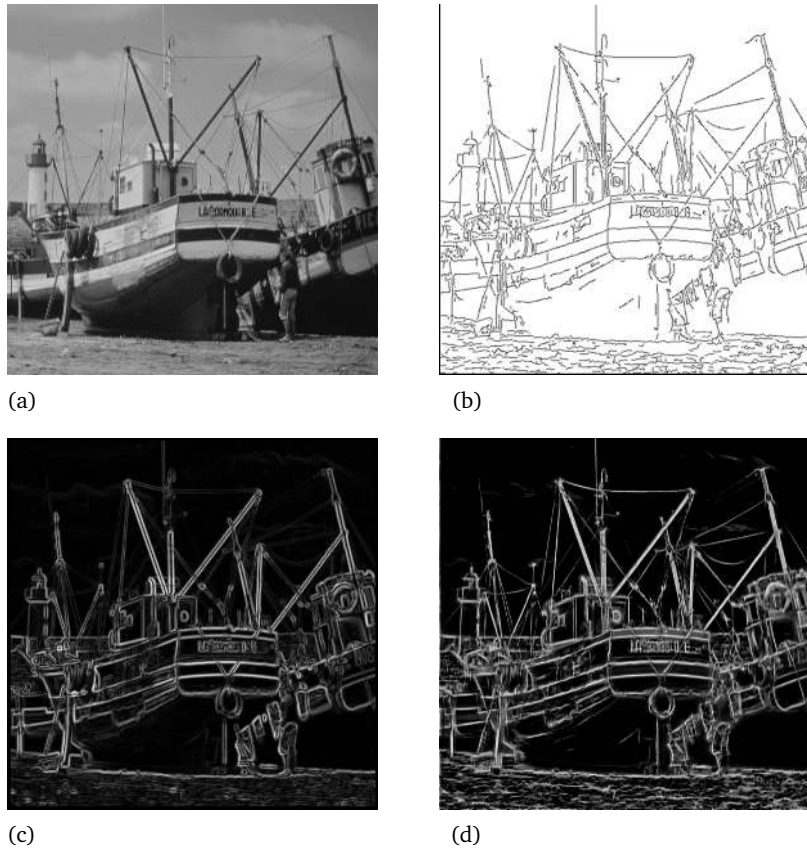



Figure 16. VDU image. The low contrast edges of the VDU against the table are marked quite clearly by phase congruency. Notice also, the marking of the right-hand vertical edge of the monitor. (a) original image; (b) phase congruency feature map; (c) Canny edge strength image; (d) raw phase congruency image.



Editors in Chief

Christopher Brown, *University of Rochester*

Giulio Sandini, *Università di Genova, Italy*

Editorial Board

Yiannis Aloimonos, *University of Maryland*

Nicholas Ayache, *INRIA, France*

Ruzena Bajcsy, *University of Pennsylvania*

Dana H. Ballard, *University of Rochester*

Andrew Blake, *University of Oxford, United Kingdom*

Jan-Olof Eklundh, *The Royal Institute of Technology (KTH), Sweden*

Olivier Faugeras, *INRIA Sophia-Antipolis, France*

Avi Kak, *Purdue University*

Takeo Kanade, *Carnegie Mellon University*

Joe Mundy, *General Electric Research Labs*

Tomaso Poggio, *Massachusetts Institute of Technology*

Steven A. Shafer, *Microsoft Corporation*

Demetri Terzopoulos, *University of Toronto, Canada*

Saburo Tsuji, *Osaka University, Japan*

Andrew Zisserman, *University of Oxford, United Kingdom*

Action Editors

Minoru Asada, *Osaka University, Japan*

Terry Caelli, *Ohio State University*

Adrian F. Clark, *University of Essex, United Kingdom*

Patrick Courtney, *Z.I.R.S.T., France*

James L. Crowley, *LIFIA—IMAG, INPG, France*

Daniel P. Huttenlocher, *Cornell University*

Yasuo Kuniyoshi, *Electrotechnical Laboratory, Japan*

Shree K. Nayar, *Columbia University*

Alex P. Pentland, *Massachusetts Institute of Technology*

Ehud Rivlin, *Technion—Israel Institute of Technology*

Lawrence B. Wolff, *Johns Hopkins University*

Zhengyou Zhang, *Microsoft Research, Microsoft Corporation*

Steven W. Zucker, *Yale University*