

Phase Congruency Detects Corners and Edges

Peter Kovesi

School of Computer Science & Software Engineering
The University of Western Australia
Crawley, W.A. 6009
pk@csse.uwa.edu.au

Abstract. There are many applications such as stereo matching, motion tracking and image registration that require so called ‘corners’ to be detected across image sequences in a reliable manner. The Harris corner detector is widely used for this purpose. However, the response from the Harris operator, and other corner operators, varies considerably with image contrast. This makes the setting of thresholds that are appropriate for extended image sequences difficult, if not impossible. This paper describes a new corner and edge detector developed from the phase congruency model of feature detection. The new operator uses the principal moments of the phase congruency information to determine corner and edge information. The resulting corner and edge operator is highly localized and has responses that are invariant to image contrast. This results in reliable feature detection under varying illumination conditions with fixed thresholds. An additional feature of the operator is that the corner map is a strict subset of the edge map. This facilitates the cooperative use of corner and edge information.

1 Introduction

With the impressive reconstruction results that have been achieved by those working in projective geometry (see for example Hartley and Zisserman [1]) there has been a renewed interest in the detection of so called ‘corners’, or ‘interest points’. The success of these reconstructions depend very much on the reliable and accurate detection of these points across image sequences.

The definition of a corner is typically taken to be a location in the image where the local autocorrelation function has a distinct peak. A variety of operators have been devised to detect corners. These include those developed by Moravec [2], Harris and Stephens [3], Beaudet [4], Kitchen and Rosenfeld [5], and Cooper et al. [6]. Corner detectors based on the local energy model of feature perception have been developed by Rosenthaler et al. [7], and Robbins and Owens [8]. More recently the SUSAN operator has been proposed by Smith and Brady [9]. Of these the Harris operator probably remains the most widely used.

A common problem with all these operators, except the SUSAN operator, is that the corner response varies considerably with image contrast. This makes the setting of thresholds difficult. Typically we are interested in tracking features over increasingly extended image sequences. The longer the sequence the

greater the variations in illumination conditions one can expect, and the setting of appropriate thresholds becomes increasingly difficult, if not impossible. Inevitably thresholds have to be set at levels lower than ideal because detecting too many features is a lesser evil than detecting not enough. Stereo and motion reconstruction algorithms are then faced with the problem of dealing with very large clouds of noisy corner points, often greatly compromising their operation. Matching operations, which usually have rapidly increasing running time as a function of input size, suffer greatly in these conditions. Typically, considerable effort has to be devoted to the cleaning up of the output of the corner detector and to the elimination of outliers. The success of this cleaning up and outlier elimination process is usually crucial to the success of the reconstruction algorithm. Indeed, a very significant proportion of Hartley and Zisserman's book [1] is devoted to robust estimation techniques.

Another difficulty many of these operators have is that the Gaussian smoothing that is employed to reduce the influence of noise can corrupt the location of corners, sometimes considerably. The SUSAN operator deserves some special comment here because it does not suffer from these problems outlined above. It identifies features by determining what fraction of a circular mask has values the same, or similar, to the value at the centre point. Thresholds are therefore defined in terms of the size of the mask and no image smoothing is required. However, the SUSAN operator assumes that edges and corners are formed by the junctions of regions having constant, or near constant, intensity, and this limits the junction types that can be modeled.

To address the many problems outlined above this paper describes a new corner and edge detector developed from the phase congruency model of feature detection. The new operator uses the principal moments of the phase congruency information to determine corner and edge information. Phase congruency is a dimensionless quantity and provides information that is invariant to image contrast. This allows the magnitudes of the principal moments of phase congruency to be used directly to determine the edge and corner strength. The minimum and maximum moments provide feature information in their own right; one does not have to look at their ratios. If the maximum moment of phase congruency at a point is large then that point should be marked as an edge. If the minimum moment of phase congruency is also large then that point should also be marked as a 'corner'. The hypothesis being that a large minimum moment of phase congruency indicates there is significant phase congruency in more than one orientation, making it a corner.

The resulting corner and edge operator is highly localized and the invariance of the response to image contrast results in reliable feature detection under varying illumination conditions with fixed thresholds. An additional feature of the operator is that the corner map is a strict subset of the edge map. This facilitates the cooperative use of corner and edge information.

This paper is organized as follows: first the phase congruency model of feature perception is reviewed. We then examine how the phase congruency responses over several orientations can be analyzed in terms of moments to provide both

edge and corner information. Finally the performance is assessed relative to the commonly used Harris operator.

2 The Phase Congruency Model of Feature Detection

Rather than assume a fixed phase, the *Local Energy Model* postulate: the Fourier components

of the intensity gradient, the *Local Energy Model* postulate: points in an image where the phase congruency is high in Figure 1 [10]. Notice




Fig. 1. Fourier series of a square wave and the sum of the first four terms.

how the Fourier components are all *in phase* at the point of the step in the square wave. Congruency of phase at *any angle* produces a clearly perceived feature [11]. The angle at which the congruency occurs dictates the feature type, for example, step or delta.

The Local Energy Model model was developed by Morrone et al. [10] and Morrone and Owens [12]. Other work on this model of feature perception can be found in Morrone and Burr [13], Owens et al. [14], Venkatesh and Owens [15], and Kovesi [16–20, 11]. The work of Morrone and Burr [13] has shown that this model successfully explains a number of psychophysical effects in human feature perception.

The measurement of phase congruency at a point in a signal can be seen geometrically in Figure 2. The local, complex valued, Fourier components at a location x in the signal will each have an amplitude $A_n(x)$ and a phase angle $\phi_n(x)$. Figure 2 plots these local Fourier components as complex vectors adding head to tail. The magnitude of the vector from the origin to the end point is the *Local Energy*, $|E(x)|$.

The measure of phase congruency developed by Morrone et al. [10] is

$$PC_1(x) = \frac{|E(x)|}{\sum_n A_n(x)}. \tag{1}$$

Under this definition phase congruency is the ratio of $|E(x)|$ to the overall path length taken by the local Fourier components in reaching the end point. If all the Fourier components are in phase all the complex vectors would be aligned and the ratio of $|E(x)|/\sum_n A_n(x)$ would be 1. If there is no coherence of phase




Fig. 2. Polar diagram showing the Fourier components at a location in the signal plotted head to tail. The weighted mean phase angle is given by $\bar{\phi}(x)$. The noise circle represents the level of $E(x)$ one can expect just from the noise in the signal.

the ratio falls to a minimum of 0. Phase congruency provides a measure that is independent of the overall magnitude of the signal making it invariant to variations in image illumination and/or contrast. Fixed threshold values of feature significance can then be used over wide classes of images.

It can be shown that this measure of phase congruency is a function of the cosine of the deviation of each phase component from the mean

$$PC_1(x) = \frac{\sum_n A_n(\cos(\phi(x) - \bar{\phi}(x)))}{\sum_n A_n(x)}. \quad (2)$$

This measure of phase congruency does not provide good localization and it is also sensitive to noise. Kovési [18, 19] developed a modified measure consisting of the cosine minus the magnitude of the sine of the phase deviation; this produces a more localized response. This new measure also incorporates noise compensation:

$$PC_2(x) = \frac{\sum_n W(x) [A_n(x)(\cos(\phi_n(x) - \bar{\phi}(x)) - |\sin(\phi_n(x) - \bar{\phi}(x))|) - T]}{\sum_n A_n(x) + \varepsilon}. \quad (3)$$

The term $W(x)$ is a factor that weights for frequency spread (congruency over many frequencies is more significant than congruency over a few frequencies). A small constant, ε is incorporated to avoid division by zero. Only energy values that exceed T , the estimated noise influence, are counted in the result. The symbols $[\]$ denote that the enclosed quantity is equal to itself when its value is positive, and zero otherwise. In practice local frequency information is obtained via banks of Gabor wavelets tuned to different spatial frequencies, rather than via the Fourier transform. The appropriate noise threshold, T is readily determined from the statistics of the filter responses to the image. For details of this phase congruency measure and its implementation see Kovési [19–21].

3 Combining Phase Congruency Information over many Orientations

A weakness of previous implementations of phase congruency has been the way in which information over many orientations is used and combined. The definition of phase congruency outlined above only applies to 1D signals. To obtain an overall measure of phase congruency in 2D local energy is first calculated in several orientations, typically six, using data from oriented 2D Gabor wavelets. Equation 3 is modified so that the numerator is the weighted and noise compensated local energy summed over all orientations, and the denominator is the total sum of filter response amplitudes over all orientations and scales. While this approach produces a phase congruency measure that results in a very good edge map it ignores information about the way phase congruency varies with orientation at each point in the image.

To include information about the way phase congruency varies with orientation we can proceed as follows: calculate phase congruency *independently* in each orientation using equation 3, compute moments of phase congruency and look at the variation of the moments with orientation. The principal axis, corresponding to the axis about which the moment is minimized, provides an indication of the orientation of the feature. The magnitude of the maximum moment, corresponding to the moment about an axis perpendicular to the principal axis, gives an indication of the significance of the feature. If the minimum moment is also large we have an indication that the feature point has a strong 2D component to it, and should therefore be additionally classified as a ‘corner’.

Following the classical moment analysis equations [22] we compute the following at each point in the image:

$$a = \sum (PC(\theta) \cos(\theta))^2 \quad (4)$$

$$b = 2 \sum (PC(\theta) \cos(\theta)) \cdot (PC(\theta) \sin(\theta)) \quad (5)$$

$$c = \sum (PC(\theta) \sin(\theta))^2, \quad (6)$$

where $PC(\theta)$ refers to the phase congruency value determined at orientation θ , and the sum is performed over the discrete set of orientations used (typically six). The angle of the principal axis Φ is given by

$$\Phi = \frac{1}{2} \text{atan2} \left(\frac{b}{\sqrt{b^2 + (a - c)^2}}, \frac{a - c}{\sqrt{b^2 + (a - c)^2}} \right). \quad (7)$$

The maximum and minimum moments, M and m respectively, are given by

$$M = \frac{1}{2} (c + a + \sqrt{b^2 + (a - c)^2}) \quad (8)$$

$$m = \frac{1}{2} (c + a - \sqrt{b^2 + (a - c)^2}). \quad (9)$$

This calculation of the maximum and minimum moments, along with the principal axis, corresponds to performing a singular value decomposition on a phase congruency covariance matrix. The moments correspond to the singular values.

3.1 Comparison with the Harris Operator

A comparison with the Harris operator is appropriate given its widespread use. The analysis described above is similar to that adopted by Harris and Stephens [3]. However they consider the minimum and maximum eigenvalues, α and β , of the image *gradient* covariance matrix in developing their corner detector. The gradient covariance matrix is given by

$$G = \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} \quad (10)$$

where I_x and I_y denote the image gradients in the x and y directions. A ‘corner’ is said to occur when the two eigenvalues are large and similar in magnitude. To avoid an explicit eigenvalue decomposition Harris and Stephens devise a measure using the determinant and trace of the gradient covariance matrix

$$R = \det(G) - k(\text{tr}(G))^2, \quad (11)$$

where $\det(G) = \alpha\beta$ and $\text{tr}(G) = \alpha + \beta$, the parameter k is traditionally set to 0.04. This produces a measure that is large when both α and β are large. However we have the problem of determining what is large. Noting that elements of the image gradient covariance matrix have units of intensity gradient squared we can see that the determinant, and hence the measure R will have units of intensity gradient to the fourth. This explains why the Harris operator is highly sensitive to image contrast variations which, in turn, makes the setting of thresholds exceedingly difficult. Some kind of sensitivity to image contrast is common to all corner operators that are based on the local autocorrelation of image intensity values and/or image gradient values.

Unlike image intensity gradient values phase congruency values are normalized quantities that have no units associated with them. If the moments are normalized for the number of orientations considered we end up with phase congruency moment values that range between 0 and 1. Being moments these values correspond to phase congruency squared. Accordingly we can use the maximum and minimum phase congruency moments *directly* to establish whether we have a significant edge and/or corner point. It should be emphasized that the minimum and maximum moments provide feature information in their own right; one does not have to look at their ratios. We can define a priori what a significant value of phase congruency moment is, and this value is independent of image contrast.

4 Results

The performance of the phase congruency operator was compared to the Harris operator on a synthetic test image, and on a real scene containing strong

shadows. The results are shown in figures 3 and 4. Raw Harris corner strength and raw phase congruency corner and edge strength images are displayed for comparison. It should be noted that the Harris corner strength values varied by many orders of magnitude across the image depending on contrast. To facilitate the display of the Harris corner strength image what is actually shown here is the *fourth root* of the image. Even after this transformation large variations are still evident. In contrast the phase congruency edge and corner strength images are minimally affected by image contrast and are readily thresholded (in this case with a value of 0.4 on both images) to produce a clear set of features.

In applying the Harris operator to the synthetic image the standard deviation of the smoothing Gaussian was deliberately set to a large value of 3 pixels to illustrate the problem the operator has in localizing 'T' junctions. Notice how the detected Harris corners are displaced inwards on the left hand grey scale, note also the double response where the line intersects the circle. The phase congruency operator, on the other hand, locates 'T' junctions precisely. Changing the number of filter scales used to compute phase congruency does not affect the localization, it only affects the relative significance of features at different scales. For the real scene the standard deviation of the smoothing Gaussian for the Harris operator was 1 pixel. For both images phase congruency was computed using Gabor filters over 4 scales (wavelengths of 4, 8, 16 and 32 pixels) and over 6 orientations.

Another important point to note is that the phase congruency edge map *includes* the corner map, this is unusual for an edge detector! The fact that the phase congruency corner map is a strict subset of the phase congruency edge map greatly simplifies the integration of data computed from edge and corner information. This facilitates the process of building a model of the scene from point and edge data matched over two or more views. In contrast, if the edge and corner information is computed via separate means, say the Canny and Harris operators respectively, the edge data is unlikely to include the corner data. The Gaussian smoothing that is applied to reduce the influence of noise results in the Canny edges being weakened in strength, and rounded at corner locations.

MATLAB code is available for those wishing to replicate the results presented here [21].

5 Conclusion

Phase congruency provides a contrast invariant way of identifying features within images. By combining phase congruency information over multiple orientations into a covariance matrix, and calculating the minimum and maximum moments we produce a highly localized operator that can be used to identify both edges and corners in a contrast invariant way. The contrast invariance facilitates the tracking of features over extended image sequences under varying lighting conditions. An additional advantage of the operator is that the phase congruency corner map is a strict *subset* of the phase congruency edge map. This simplifies the integration of data computed from edge and corner information.

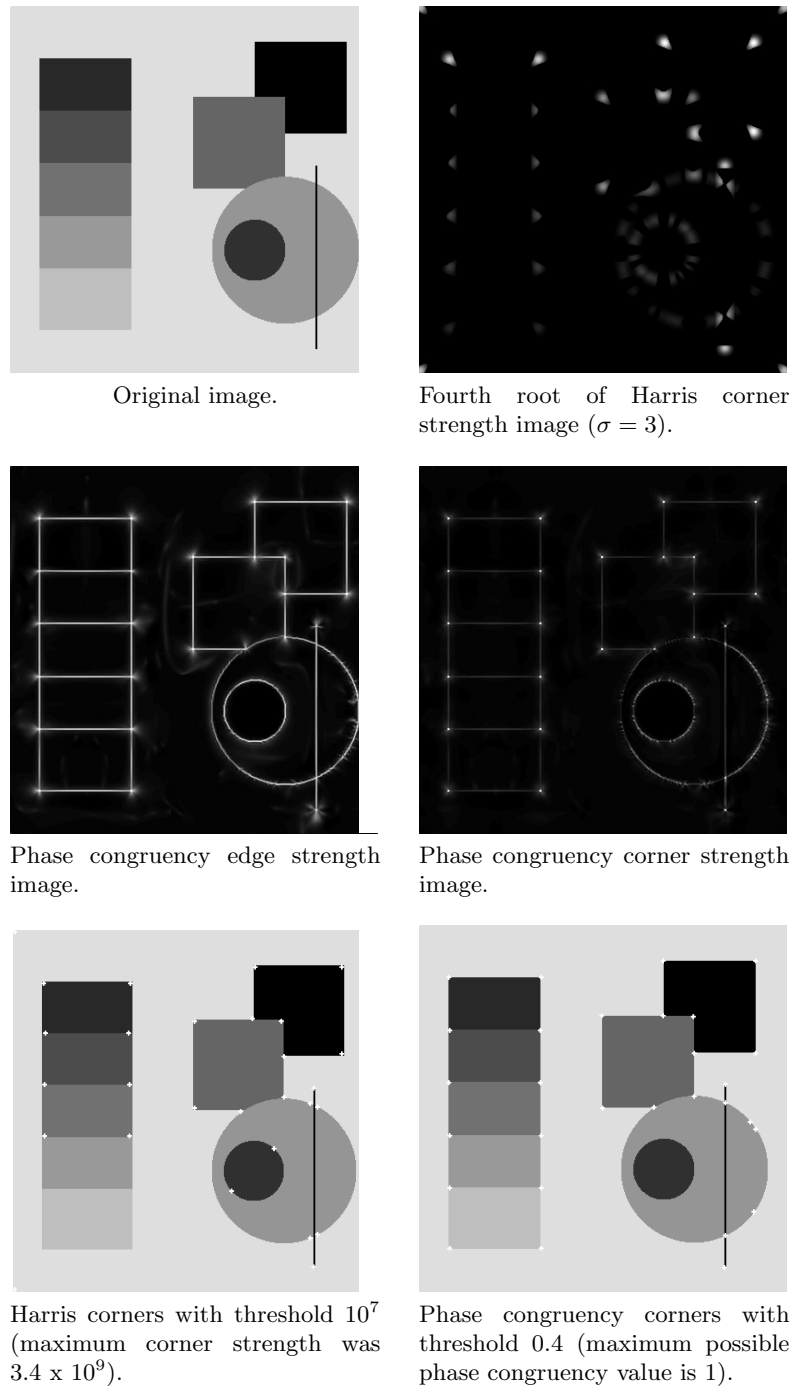





Fig. 3. Comparison of the Harris and phase congruency operators on a test image.




Original image.




Fourth root of Harris corner strength image ($\sigma = 1$).




Phase congruency edge strength image.



Phase congruency corner strength image.



Harris corners with threshold 10^8 (maximum corner strength was 1.25×10^{10}).



Phase congruency corners with threshold 0.4 (maximum possible phase congruency value is 1).

Fig. 4. Comparison of the Harris and phase congruency operators on an image with strong shadows.

References

1. Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge University Press (2000)
2. Moravec, H.P.: Robot rover visual navigation. UMI Research Press (1981)
3. Harris, C., Stephens, M.: A combined corner and edge detector. In: Proceedings, 4th Alvey Vision Conference. (1988) 147–151 Manchester.
4. Beaudet, P.R.: Rotationally invariant image operators. In: International Joint Conference on Artificial Intelligence. (1987) 579–583
5. Kitchen, L., Rosenfeld, A.: Grey-level corner detection. Pattern Recognition Letters (1982) 95–102
6. Cooper, J., Venkatesh, S., Kitchen, L.: Early jump-out corner detectors. PAMI **15** (1993) 823–828
7. Rosenthaler, L., Heitger, F., Kubler, O., Heydt, R.v.: Detection of general edges and keypoints. In: ECCV92, Springer-Verlag Lecture Notes in Computer Science. Volume 588., Springer-Verlag (1992) 78–86 Santa Margherita Ligure, Italy.
8. Robbins, B., Owens, R.: 2D feature detection via local energy. Image and Vision Computing **15** (1997) 353–368
9. Smith, S.M., Brady, J.M.: SUSAN - a new approach to low level image processing. International Journal of Computer Vision **23** (1997) 45–78
10. Morrone, M.C., Ross, J.R., Burr, D.C., Owens, R.A.: Mach bands are phase dependent. Nature **324** (1986) 250–253
11. Kovese, P.D.: Edges are not just steps. In: Proceedings of the Fifth Asian Conference on Computer Vision. (2002) 822–827 Melbourne.
12. Morrone, M.C., Owens, R.A.: Feature detection from local energy. Pattern Recognition Letters **6** (1987) 303–313
13. Morrone, M.C., Burr, D.C.: Feature detection in human vision: A phase-dependent energy model. Proc. R. Soc. Lond. B **235** (1988) 221–245
14. Owens, R.A., Venkatesh, S., Ross, J.: Edge detection is a projection. Pattern Recognition Letters **9** (1989) 223–244
15. Venkatesh, S., Owens, R.: On the classification of image features. Pattern Recognition Letters **11** (1990) 339–349
16. Kovese, P.D.: A dimensionless measure of edge significance. In: The Australian Pattern Recognition Society, Conference on Digital Image Computing: Techniques and Applications. (1991) 281–288 Melbourne.
17. Kovese, P.D.: A dimensionless measure of edge significance from phase congruency calculated via wavelets. In: First New Zealand Conference on Image and Vision Computing. (1993) 87–94 Auckland.
18. Kovese, P.D.: Invariant Measures of Image Features From Phase Information. PhD thesis, The University of Western Australia (1996)
19. Kovese, P.D.: Image features from phase congruency. Videre: Journal of Computer Vision Research **1** (1999) 1–26 <http://mitpress.mit.edu/e-journals/Videre/>.
20. Kovese, P.D.: Phase congruency: A low-level image invariant. Psychological Research **64** (2000) 136–148
21. Kovese, P.D.: MATLAB functions for computer vision and image analysis (1996-2003) <http://www.csse.uwa.edu.au/~pk/Research/MatlabFns/>.
22. Horn, B.K.P.: Robot Vision. McGraw-Hill (1986) New York.