Measuring Temporal Phase Congruency

Peter J. Myerscough, Mark S. Nixon

Department of Electronics and Computer Science
University of Southampton, UK

January 15, 2004

Abstract

We describe a robust moving feature detector, that extracts feature points and feature velocities from a sequence of images. We develop a new approach based on phase congruency to include interpolation of feature orientation and improvements in robustness due to correlations in the image sequence. This new temporal phase congruency operator show improved capabilities on a series of different real image types, as well as a noise analysis on synthetic images.

1 Introduction

We present a new extension of phase congruency [2] taking it from single images to image sequences. An earlier approach by Mulet-Parada[3] also examines phase in images sequences, in application to echocardiography, but removes the multi-scale approach. We also extend the operator's ability to describe the orientation of features found from 1 of 6 orientations to a more granular scale. To this we also extend the orientation of a feature from simple 2D, to 2D + time, which is equivalent to describing a feature's velocity, or component velocity due to the aperture problem.

1.1 Temporal Phase Congruency

Phase congruency is a robust feature detector. It detects not only step and line responses, but also a broader set of features [1]. Its attributes include a high degree of invariance to lighting variation within images and image sequences.

To extend it to temporal phase congruency the first step is to convolve the image sequence with a set of log-Gabor filters at 'l' different orientations in space and time and 'm' different scales. Log-Gabor filters are chosen because they have zero DC response, and in cosine and sine based pairs they have a quadrature relationship. They can also cover large portions of the frequency space while keeping the DC component of the cosine part of the Gabor pair to zero. They are constructed in the frequency domain using a polar co-ordinate system, with the log-Gabor function along the wavelength axis and a Gaussian spreading function as the frequencies differ in angle in the spatio-temporal domain to that of the focus. The congruency between the different scales of filter about an orientation is then calculated, $C_l(x, y, t)$.

$$C_l(x,y,t) \quad = \quad \sum_{n}^{m} A_n(x,y,t)\Big(\cos\left(\Delta\phi_n\left(x,y,t\right)\right) - \left|\sin\left(\Delta\phi_n\left(x,y,t\right)\right)\right|\Big) \tag{1}$$

where $\Delta\phi_n\left(x,y,t\right)$ is the difference between the phase at each scale and the mean phase angle and $A_n(x,y,t)$ is the amplitude at scale 'n'. The $\cos(\ldots)$ and $\sin(\ldots)$ are calculated using the dot and cross products of the Gabor filter results with the mean of the filter results. This measure of congruency is then thresholded against an estimate of noise within the sequence, and weighted against the spread of the energy across each scale. It is then normalised by the total energy about all orientations. To calculate the orientation of the features, the orientation with the highest congruency for each point, is set as the rough orientation. From this the congruency values of surrounding orientations are formed into a patch and the congruency values are treated as 2D Gaussians whose sum forms a surface. The maximum of this surface is a more accurate value for the orientation of the features. The final phase congruency value is the sum of all the normalised congruences. (An alternative measure is the height of the surface used in interpolating the orientation.) Questions still remain as to the correct resolution of features that are more complex, such as T-junctions, where the concept of orientation is more abstract.

The new technique is likely to be more robust to noise due to the increased number of neighbours a feature has and the correlation between them. Typically features exist within a 2D line in images, but in image sequences they also persist over time. As such, feature detection can benefit from the correlation over the image sequence.

2 Results

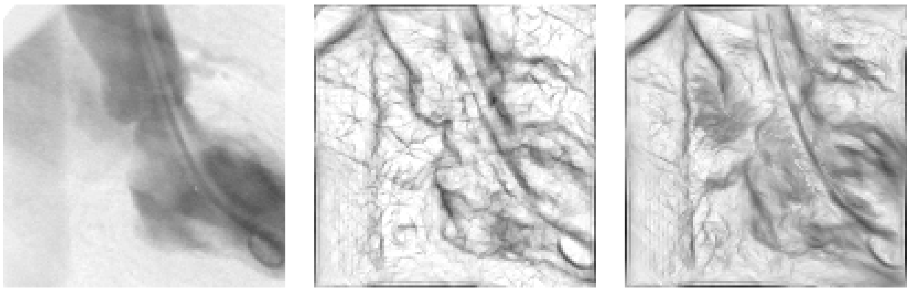

(a) Frame from a sequence of a pumping heart

(b) Phase congruency results

(c) Temporal phase congruency results

FIGURE 1: Output from the temporal phase congruency operator.

Example results are provided for application to a sequence of image of a living heart, figure 1. The features in the heart have been highlighted successfully by the new technique. The temporal phase congruency technique shows more blurring around the features, but extracts more of the features within the heart compared to the phase congruency operator. It also provides extra velocity information for the features they detect, which is not given by the original version.

2.1 Noise testing

The new temporal phase congruency has also been tested against the original phase congruency technique on a synthetic sequence of a moving circle. The first test was a simple visual comparison. Both techniques extracted the edges of this simple image sequence very well. To gain a deeper insight, two types of noise were added in increasing amounts. The first was salt and pepper noise, where the 'noisy' pixels were set arbitrarily to black or white. The second was additive Gaussian noise, where a zero-mean Gaussian distributed random number was added to each pixel. The resulting feature maps are then compared with ground truth, where acceptable positions for a feature are compared with the actual feature map. These are plotted with the number of features that are falsely detected.

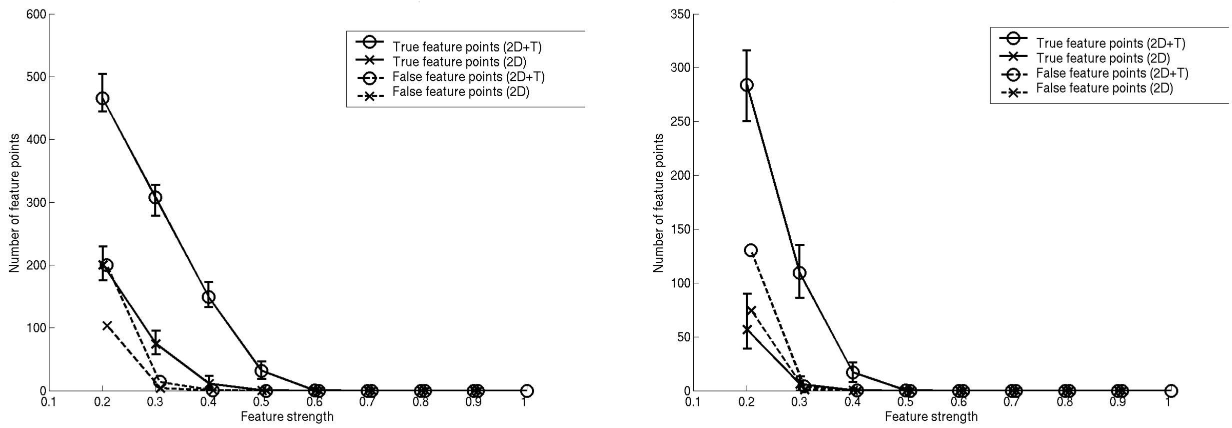

(a) Results for Gaussian noise with zero mean Gaussian with sigma set to 50% of the difference between black and white

(b) Results for 50 % salt and pepper noise

FIGURE 2: Results from testing a simple white disc moving on a black background with different types of noise added. Solid lines represent true features detected, dashed lines represent false features detected. Error bars show the range of results due to the test being run fifty times.

The results show that the temporal phase congruency measure is detecting more feature pixels of the circle as it moves through the sequence, and with a higher measure of feature strength or significance. These results are consistent for both types of additive noise that range from 10% to 90%, in each case the new technique shows better detection rates.

3 Conclusions

A new temporal phase congruency operator shows improvements over the original operator in noisy conditions and also provides velocity information for the moving features detected, which should allow for object outlines to be grouped more easily by higher level processes. Further work into different measures for phase congruency may be needed so that the temporal phase congruency detector has less blurring around some features.

Acknowledgements

This research has be funded by the EPSRC, with additional support from European Research Office of the US Army, Contract No.N68171-01-C-9002. Thanks also go to Ohio State university for the heart sequence shown in figure 1.

References

[1] Y. K. Aw, R. Owens, and J. Ross. A catalog of 1-d features in natural images. *CVGIP: Graphical Models and Image Processing*, 56(2):173–181, March 1994.

[2] P Kovesi. Image features from phase congruency. *Videre : Journal of Computer Vision Research*, 1(3):1–27, 1999.

[3] M. Mulet-Parada and J.A. Noble. 2d+t acoustic boundary detection in echocardiography. *Medical Image Analysis*, 4:21–30, 2000.