

Figure 9: Phase congruency of a noisy signal calculated using a rectangular windowing function 32 points wide.

with high information content; a point of phase congruency indicates a point of high information content only if we have a wide range of frequencies present. We do not gain much information from knowing the phase congruency of a signal which has only one frequency component.

2.5 Summary

This chapter has briefly re-examined the aims of feature detection. The objective should be to find points of high information content in images. This objective is not necessarily satisfied by finding optimal ways of detecting step edges in the presence of noise. Ideally, the ability to detect features and assess their significance should

be independent of image contrast and spatial magnification. This implies that we need to measure feature significance via a dimensionless quantity.

The shortcomings of derivative based feature detectors have been briefly reviewed. The main problems included an inability to specify in advance what level of response corresponds to a significant feature, and the fact that they are generally only designed to detect step edges. The local energy model of feature perception has been introduced. This model has been inspired from psychophysical data and it detects a wide range of feature types. Local energy can be normalized to produce a measure of phase congruency; an approximation of the standard deviation in phase angles of the Fourier components at a point in the signal. Phase congruency is a dimensionless quantity and is thus an attractive way of detecting features and identifying their significance. It provides an absolute measure of the significance of feature points in an image and this offers the promise of allowing constant threshold values to be applied across wide classes of images. Thresholds could then be specified *in advance* of processing any image, and not have to be determined by trial and error *after* processing.

However, there are a number of issues to be addressed. How should phase congruency be calculated in 2D images? How should we calculate local frequency information and control the scale of analysis? How do we deal with the influence of noise, and how do we identify the range of frequencies present at a point of phase congruency? These issues, and others, are addressed in the following chapter which will describe how phase congruency can be calculated using wavelets.

Chapter 3

Phase congruency from wavelets

3.1 Introduction

This chapter describes a new way of calculating phase congruency using wavelets. In calculating phase congruency it is important to obtain spatially localized frequency information in images, wavelets offer perhaps the best way of doing this. The use of wavelets is also biologically inspired; the interest in calculating phase congruency is motivated from psychophysical results, hence, it would seem natural that one should try to calculate it using biologically plausible computational machinery. In this respect geometrically scaled spatial filters in quadrature pairs will be used. In addition to this it will be seen how the use of wavelets allows one to address the issues raised at the end of the previous chapter regarding the calculation of phase congruency.

The material that will be covered in this chapter is organized as follows: First, it will be shown how local frequency information can be obtained using quadrature pairs of wavelets, concentrating in particular on the use of Gabor wavelets. From this it is relatively straightforward to develop the ideas behind the calculation of phase congruency in one dimensional signals using wavelets. Material is then presented to address the difficulties regarding the calculation of phase congruency that were introduced in the previous chapter. First, the influence of noise in the calculation of phase congruency is considered and an effective method for identifying and compensating for noise is developed. This is followed by a section covering the

issues involved in extending the calculation of phase congruency to 2D images. It is then shown how the use of wavelets allow us to obtain a measure of the spread of frequencies present at a point of phase congruency. This helps one determine the degree of significance of a point of phase congruency and allows to improve feature localization. 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 conclusions are presented.

3.2 Using Wavelets for Local Frequency Analysis

Recently the Wavelet Transform has become one of the methods of choice for obtaining local frequency information. Most of the current literature on wavelets can be traced back to the work of Morlet et al. [59] Morlet and his co-workers were interested in obtaining temporally localized frequency data in their analysis of geophysical signals. The basic idea behind wavelet analysis is that one uses a bank of filters to analyze the signal. The filters are all created from rescalings of the one wave shape, each scaling designed to pick out particular frequencies of the signal being analyzed. An important feature is that the scales of the filters vary geometrically, giving rise to a logarithmic frequency scale.

However, many of these ideas were developed earlier by Granlund [30]. In this remarkable paper he developed many of the ideas behind what we would now call multi-scale wavelet analysis. He also proposed an image feature detector that is closely related to the local energy model. For some reason Granlund's paper has remained relatively unnoticed despite its innovative nature, though his work has been developed by Wilson, Calway and Knutsson (see, for example, Wilson, Calway and Granlund [96], Knutsson, Wilson and Granlund [43], Calway and Wilson [10] and Calway, Knutsson and Wilson [9]). From the initial work of Morlet and his colleagues wavelet theory has been subsequently developed by Grossmann and Morlet [31], Meyer [58], Daubechies [14], Mallat [53] and many others.

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 wavelets that are symmetric/anti-symmetric. This constraint means that the work on orthogonal wavelets (which dominates much of the literature) is not applicable to us. Chui [13] provides a proof that, with the exception of the Haar wavelet, one cannot have a wavelet of compact support that is both symmetric and orthogonal. The Haar wavelet is rectangular in shape and is clearly not appropriate for our needs.

For this work the approach of Morlet will be followed, that is, using wavelets based on complex valued Gabor functions - sine and cosine waves, each modulated by a Gaussian. Using two filters in quadrature enables one to calculate the amplitude and phase of the signal for a particular frequency at a given spatial location. It should be noted that these wavelets are not orthogonal; some conditions must apply in order to achieve reasonable signal reconstruction after decomposition. However, we only require approximate reconstruction up to a scale factor over a band of frequencies or wavelet scales.




Figure 10: Gabor wavelet: a sine and cosine wave modulated by a Gaussian.

If the bank of wavelet filters is designed so that the transfer function of each filter overlaps sufficiently with its neighbours in such a way that the sum of all the transfer functions forms a relatively uniform coverage of the spectrum one can reconstruct the decomposed signal over a band of frequencies up to a scale factor. (If the transfer functions are scaled so that when their sum is taken we obtain a

uniform transfer function of magnitude one, the reconstructed signal will have the original scale.) Therefore, a problem we have is determining the appropriate scaling factor between successive centre frequencies so that the overlap between transfer functions results in an even spectral coverage. Granlund [30] suggests that the upper cutoff frequency of one transfer function (where it falls to half its maximum value) should coincide with the lower cutoff frequency of the next function. However, in practice this does not produce particularly even coverage, and a closer spacing is generally desirable. In the results presented in this chapter the filters used have had bandwidths of approximately one octave with a scaling between successive centre frequencies of 1.5. This arrangement was arrived at by experimentation, the values are not critical and a wide range of parameters produce satisfactory results. Referring to Figure 11 one can see that, in this example, the sum of the spectra of the five wavelets produces a relatively ideal band-pass filter, especially when viewed on the log frequency scale. Design of the wavelet bank ends up being a compromise between wishing to form a smooth sum of spectra while at the same time minimizing the number of filters used so as to minimize the computation requirements.

Analysis of a signal is done by convolving the signal with each of the wavelets. If we let I denote the signal and M_n^e and M_n^o denote the even and odd wavelets at a scale n , the amplitude of the transform at a given wavelet scale is given by

$$A_n(x) = \sqrt{(I(x) * M_n^e)^2 + (I(x) * M_n^o)^2} \quad (7)$$

and the phase is given by

$$\Phi_n(x) = \text{atan2}(I(x) * M_n^e, I(x) * M_n^o). \quad (8)$$

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

The results of convolving a signal with a bank of wavelets can be displayed graphically via a scalogram (Figure 12). Each row of the scalogram is the result of convolving the signal with a quadrature pair of wavelets at a certain scale. Phase is plotted by mapping 0–360 degrees to 0–255 grey levels (note therefore, that the black/white discontinuities in the scalogram correspond to the wrap-around in

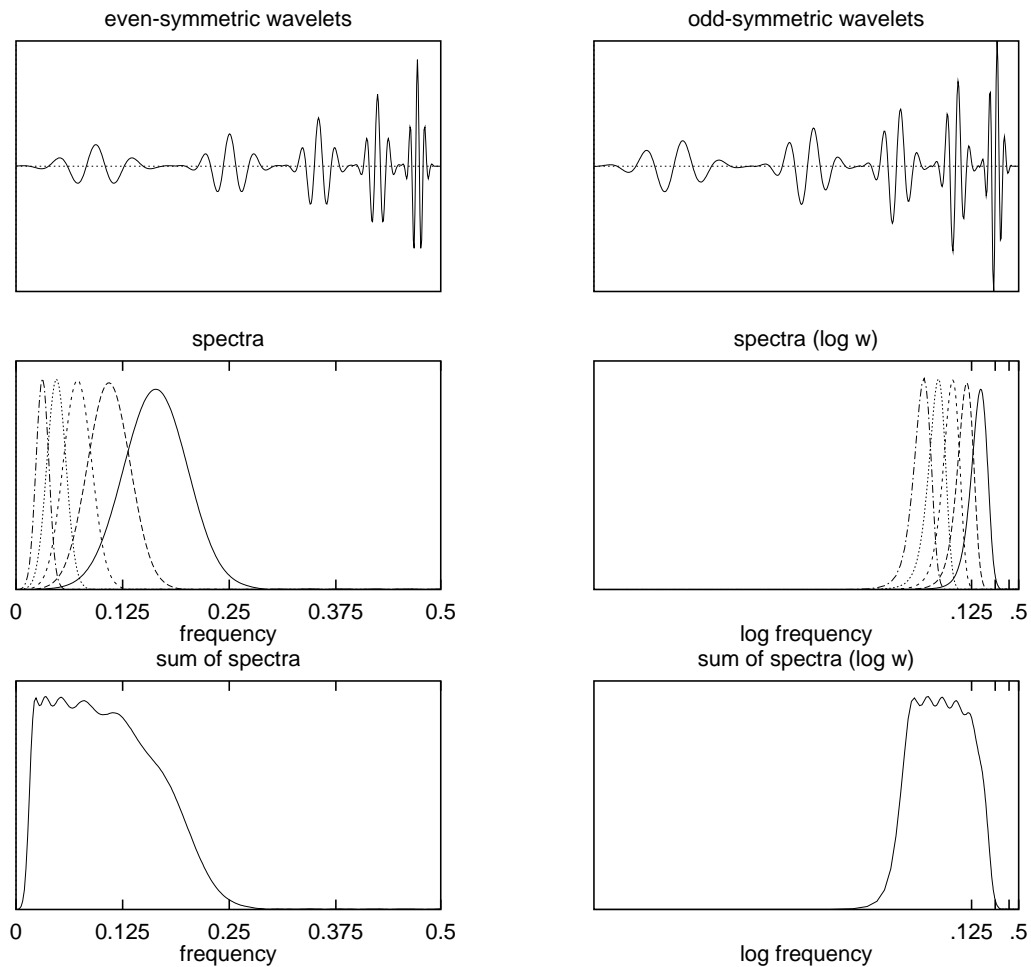


Figure 11: Five wavelets and their respective Fourier Transforms indicating which sections of the spectrum each wavelet responds to. Collectively the wavelets provide a wide coverage of the spectrum, though with some overlap. Note that on a logarithmic frequency scale the spectra are identical.

phase). The vertical axis of the scalogram is a logarithmic frequency scale, with the lowest frequency at the bottom. Each column of the scalogram can be considered to be a local Fourier spectrum for each point. Note that to achieve a dense scalogram such as shown here the scaling factor between successive filter center frequencies will be only slightly greater than 1.

The phase plot of the scalogram is of particular interest because it enables one to actually see the points of high phase congruency. At locations in the signal where there are large step changes one can see a vertical line of constant grey value in the phase diagram indicating a constant phase angle over all frequencies at that point in the signal.

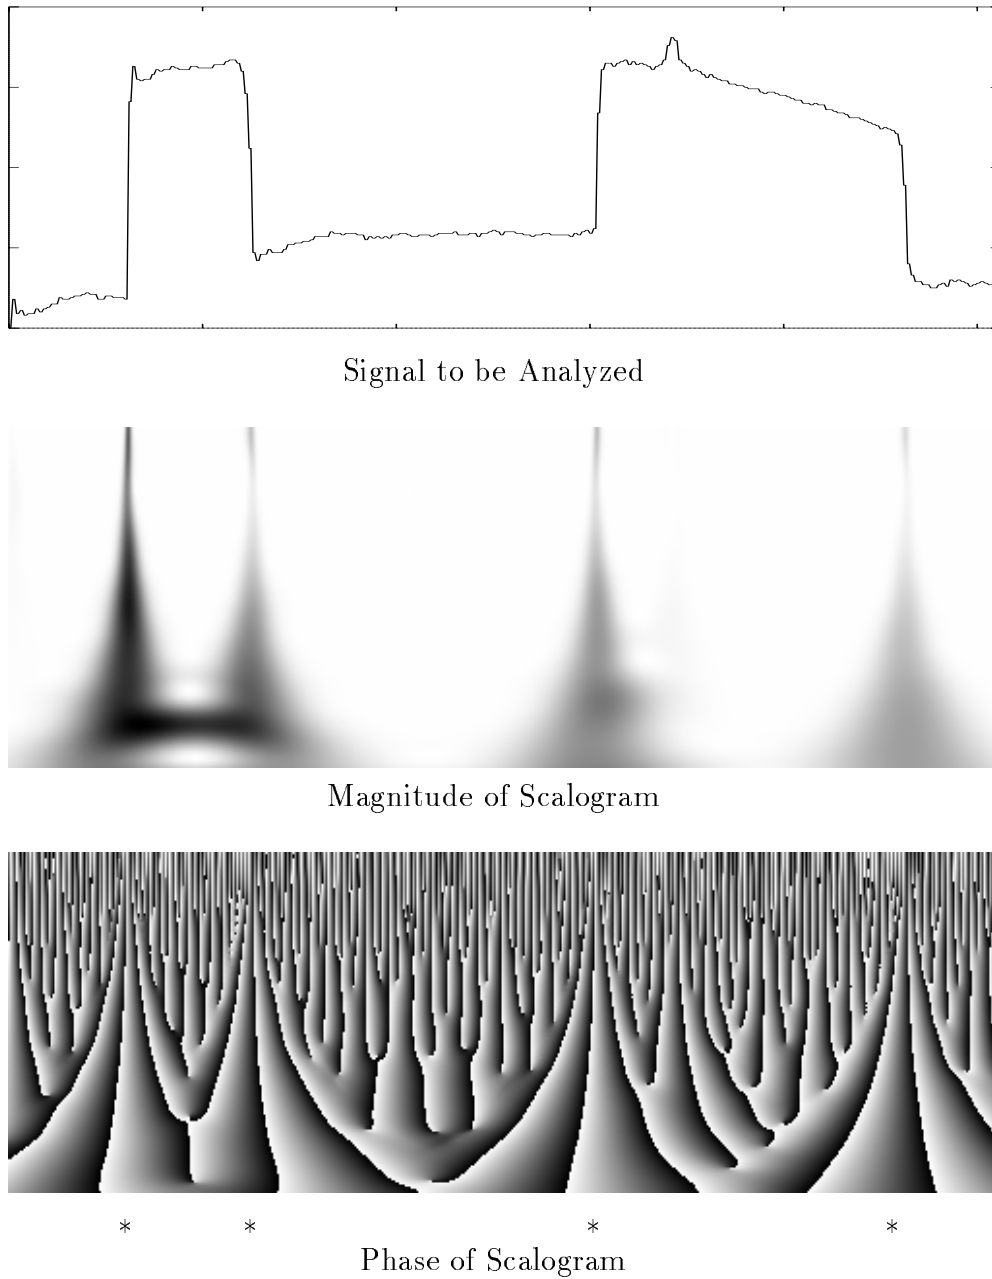


Figure 12: A one dimensional signal and its amplitude and phase scalograms. The horizontal axes of the scalograms correspond directly with the signal's horizontal axis. The vertical axes of the scalograms correspond to a logarithmic frequency scale with low frequencies at the bottom. The asterisks mark vertical lines of constant phase that occur at the step transitions in the signal. These are points of phase congruency. (Note: the phase scalogram is presented by mapping 0–360 degrees to 0–255 grey levels).

3.3 Calculating Phase Congruency Via Wavelets

To calculate phase congruency we need to construct the following quantities: Firstly we need to remove the DC component from our signal, $I(x)$ while at the same time retaining as many of the other frequency components as possible. We will denote this DC removed signal $F(x)$. Note, we wish to retain a broad range of frequencies in our signal because phase congruency is only of interest if it occurs over a wide range of frequencies. Secondly we have to construct $H(x)$, the Hilbert transform of $F(x)$, and finally we need to calculate the normalizing component $\sum A_n(x)$, the sum of the amplitudes of the frequency components in $F(x)$. Note that the sum of amplitudes becomes a function of x as our frequency analysis of the signal is now spatially localized.

The outputs from a quadrature pair of filters can be thought of as representing a vector in the complex plane; the real component coming from the even-symmetric filter output, and the odd-symmetric output representing the imaginary component (see Figure 13). If we sum the results of convolving our signal with the bank of even wavelets we will reconstruct a band passed version of our signal, amplified according to the scaling and overlap of the transfer functions of our filters. We can use this result for $F(x)$, an approximation of the original signal with the DC component removed. An approximation of the Hilbert transform, corresponding to $H(x)$, can be constructed from the sum of the convolutions of the signal with the odd wavelets; this is a signal covering the same bandwidth of the original signal and amplified in the same way as $F(x)$, but shifted in phase by 90° . Thus

$$F(x) = \sum_n I(x) * M_n^e, \text{ and} \quad (9)$$

$$H(x) = \sum_n I(x) * M_n^o. \quad (10)$$

The sum of the amplitudes of the frequency components in $F(x)$ is given by

$$\sum_n A_n(x) = \sum_n \sqrt{(I(x) * M_n^e)^2 + (I(x) * M_n^o)^2}. \quad (11)$$

With these three components we are able to calculate phase congruency. At this point the expression for phase congruency is slightly modified by adding a small positive constant to the denominator. This small constant, ε , prevents the

expression from becoming unstable when $\sum_n A_n(x)$ (and hence $E(x)$) becomes very small. Thus

$$PC(x) = \frac{E(x)}{\sum_n A_n(x) + \varepsilon}, \quad (12)$$

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. Large values of ε can be used to suppress the influence of noise, but as we shall see in the next section there is a much better way to compensate for noise. An ε value of 0.01 has been used for all the results presented here.

Note that the amplification of the signal due to the scaling and overlap in the filter transfer functions is cancelled out through the normalization process in the calculation of phase congruency. The use of wavelets allows one to obtain high spatial localization in the calculation of phase congruency. The frequency range over which phase congruency is calculated is easily controlled through the number of scales used in the wavelet bank. Figure 14 illustrates all the intermediate steps in the calculation of phase congruency on a one-dimensional signal. Note that this figure incorporates a measure of frequency spread in the calculation of phase congruency. Frequency spread is discussed in section 3.6.

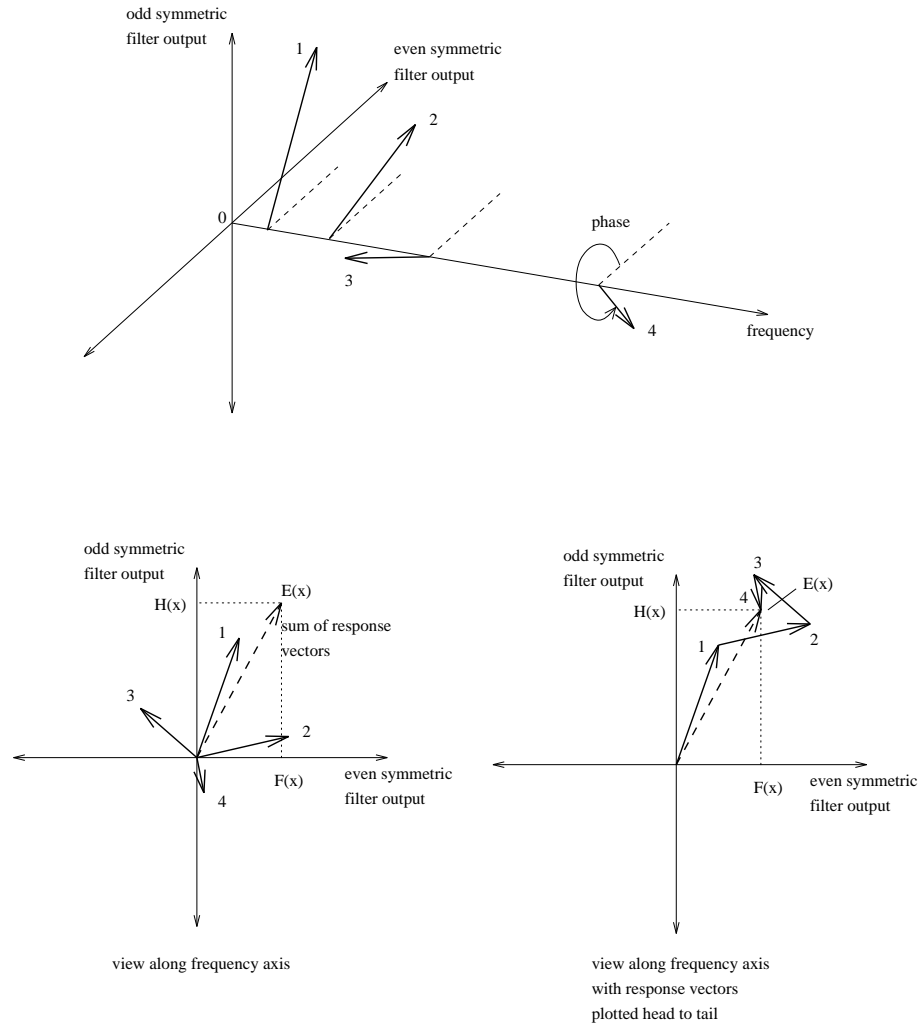


Figure 13: Calculation of phase congruency from convolution of the signal with quadrature pairs of filters. The convolution output from each quadrature pair of filters at a location in the signal can be considered to represent a response vector having length A_n and phase angle ϕ_n . When the response 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.

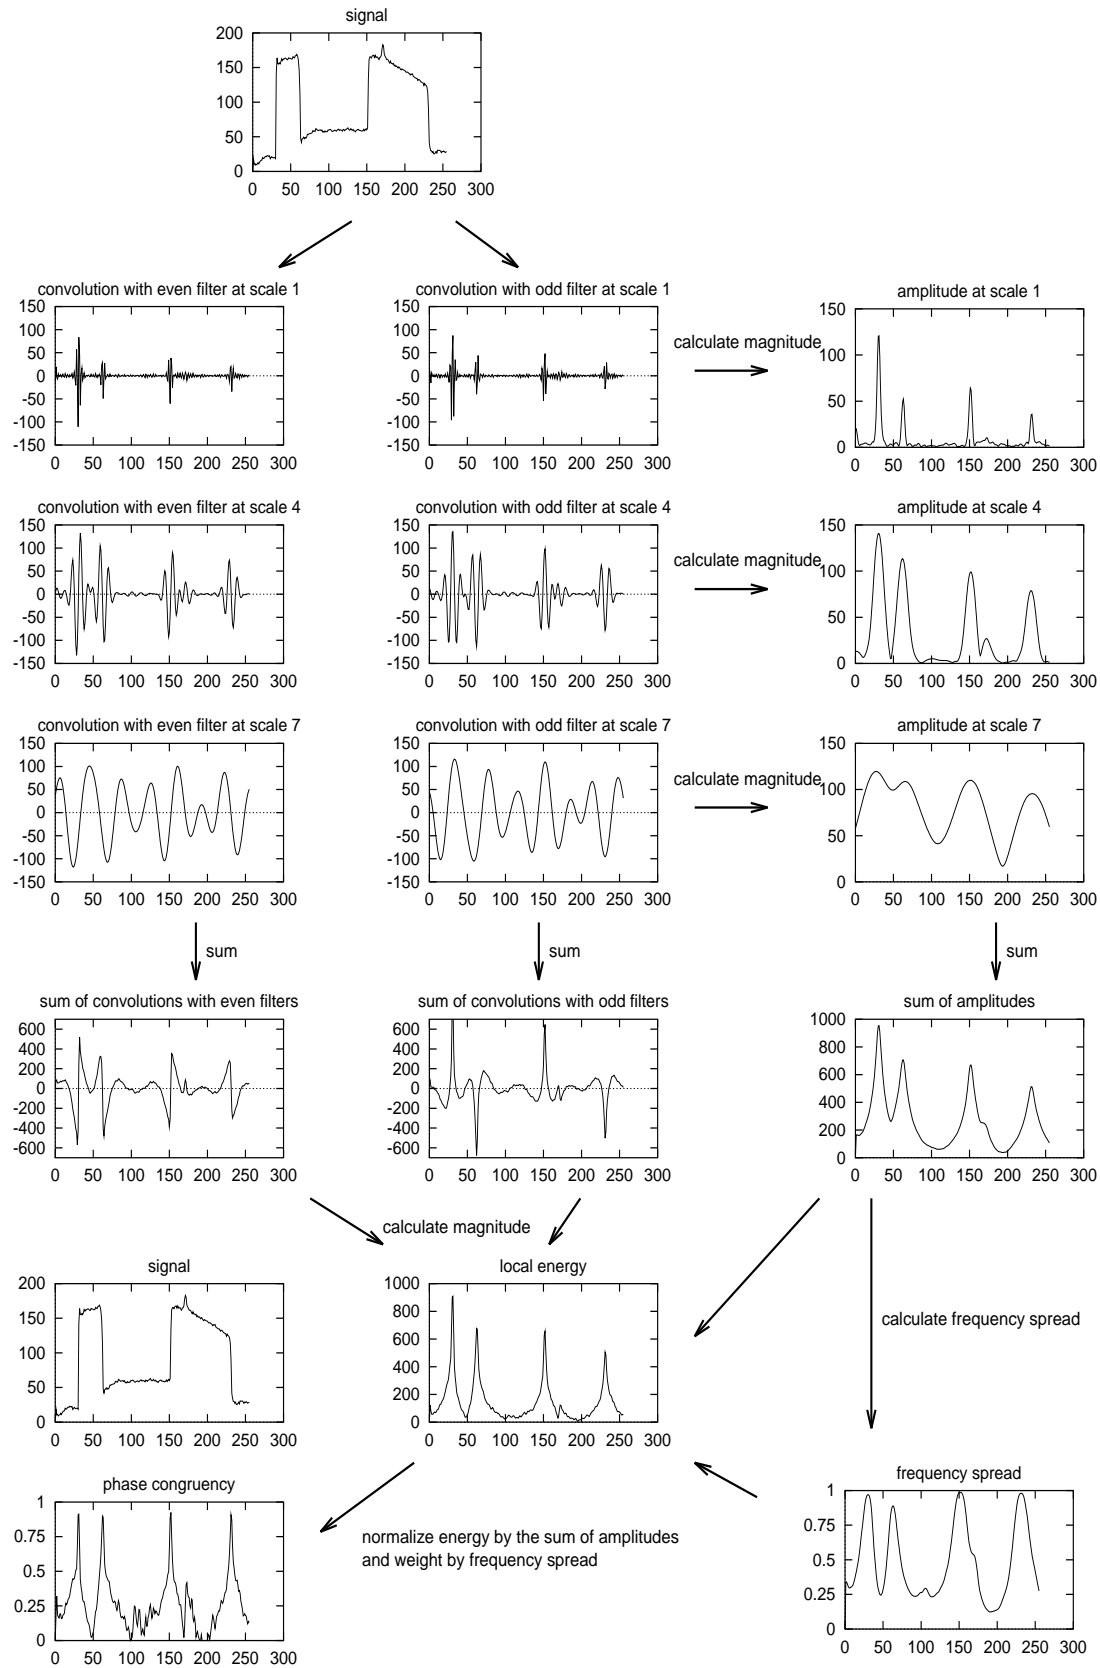


Figure 14: Steps in the calculation of phase congruency. Note: Not all intermediate convolutions are displayed.

3.4 Noise

A difficulty with phase congruency is its response to noise. Figure 15 illustrates the phase congruency of a step function with and without noise. In this example the signal-to-noise ratio is 80 (here SNR is defined as $step\ size/\sigma_{noise}$). In the vicinity of the step, phase congruency is only high 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.

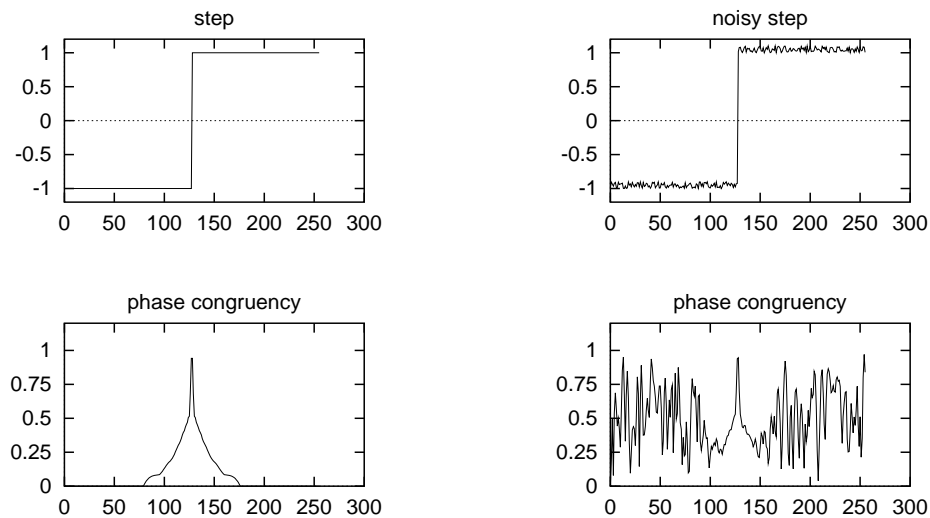


Figure 15: Phase congruency of a step function with and without noise.

What is it that makes noise noise?

Trying to recognize noise by its spectral characteristics is of no use to us. The amplitude spectrum of noise is typically flat, effectively indistinguishable from the spectral characteristics of a line (delta) feature in an image. However, we can make an attempt to identify the level of noise in an image if we make use of the following two observations:

1. Noise is *everywhere* in an image and its level is generally constant.
2. Features, such as edges, occur sparsely in an image (in fact on a set of measure zero).


If one assumes these two conditions (and assumes that the noise is additive) we can deduce that the response of the smallest scale (highest centre frequency) filter pair in our wavelet bank will be almost entirely due to noise. Only at feature points will the response differ from the background noise response, and the regions where it will be responding to features will be small due to the small spatial extent of the wavelet. The distribution of response amplitudes from the smallest scale filter pair across the whole image will be asymmetric. The peak of the distribution, representing response to noise, will be strongly skewed to the left and a long tail will extend at the upper end of the distribution as a result of responses to features. A simple mean of such a distribution will not provide us with a good estimate of the average noise response because the response to features will strongly bias the result. The approach that has been taken is to use the exponent of the mean (over x) of the log of the response amplitudes as an estimate of the average response to noise. That is, the estimate of the mean noise response of the smallest scale filter pair over the signal is given by

$$\overline{A}'_0 = e^{\overline{\log A_0(x)}}. \quad (13)$$

Taking the log of the response amplitudes tends to discount the influence of the tail at the upper end of the distribution¹. As illustrated in the example shown in Figure 16 this estimate is not perfect but is certainly better than using a simple mean. Another candidate for estimating the mean noise response is the *mode* of the distribution. However, in experiments it has been found that use of the mode can produce unreliable results. The response amplitude distribution is not always a smooth curve, and may not have a unique or clearly identifiable maximum. Spurious spikes can also cause the mode to be a very poor estimator.

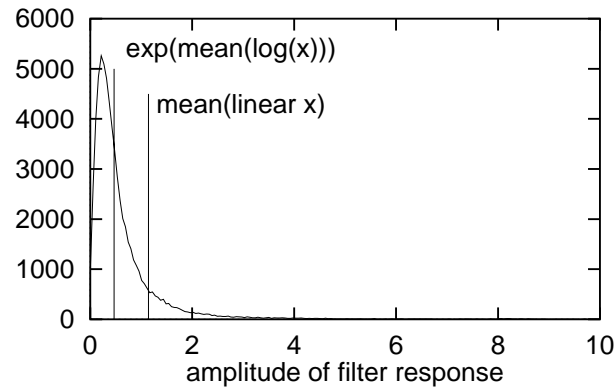
If the noise distribution is uniform the typical maximum amplitude response to noise will be given by twice this estimated average noise response as the minimum possible response level will always be 0. If the noise distribution is Gaussian

¹Canny [12] encounters a similar problem in estimating the noise response of his optimal detector. He convolves the output of his edge detector output with the second derivative of an impulse function. The distribution of the output consists of a Gaussian noise distribution plus an extended tail due to valid edge responses. Canny chooses to take the mean of the lower 80% of the distribution to estimate the noise.



(a) image

(b) filter response amplitudes



(c) histogram of amplitudes

Figure 16: Example histogram of response amplitudes of a quadrature pair of small scale filters applied in one orientation to an image.

(a) Image being analyzed.

(b) Amplitude of smallest scale filter pair response (filters oriented horizontally).

(c) Histogram of response amplitudes. Marked on the histogram is the simple mean of amplitudes and the exponent of the mean of the log of the amplitude values. (Note the response amplitude image contains floating point values. Values in the image were quantized into 1000 levels in order to construct the histogram. Note also that the histogram plot has been truncated; the maximum amplitude was approximately 53.)

the typical maximum amplitude response will be approximately three times the estimated average noise response (the mean amplitude being an estimate of the standard deviation). It remains for us to estimate the response of the wavelets at other scales in our filter bank to the noise in the image, and then from this deduce the influence of noise on our measure of phase congruency.

If it is assumed that the frequency spectrum of the noise is flat we can estimate

the noise response of other wavelets relative to the response of the smallest scale wavelet pair from the relative bandwidths of the wavelets. The square of a filter pair's response amplitude in the spatial domain will be related to the power spectrum of the signal via Rayleigh's theorem: The integral of the squared modulus of a function is equal to the integral of the squared modulus of the spectrum. That is

$$\int_{-\infty}^{\infty} |f(x)|^2 dx = \int_{-\infty}^{\infty} |F(\omega)|^2 d\omega .$$

If we assume the noise power level is constant with frequency then the square of a filter's response amplitude will be proportional to its bandwidth. Thus, the amplitude of a filter's response to noise will be proportional to the square root of its bandwidth (and hence the square root of its centre frequency).

$$\begin{aligned} A_{n_{noise}}^2 &\propto \text{bandwidth} \propto f_n \\ A_{n_{noise}} &\propto \sqrt{f_n} \end{aligned} \quad (14)$$

where $A_{n_{noise}}$ is the amplitude of the noise response of the filter at scale n and f_n denotes the centre frequency of the filter at scale n .

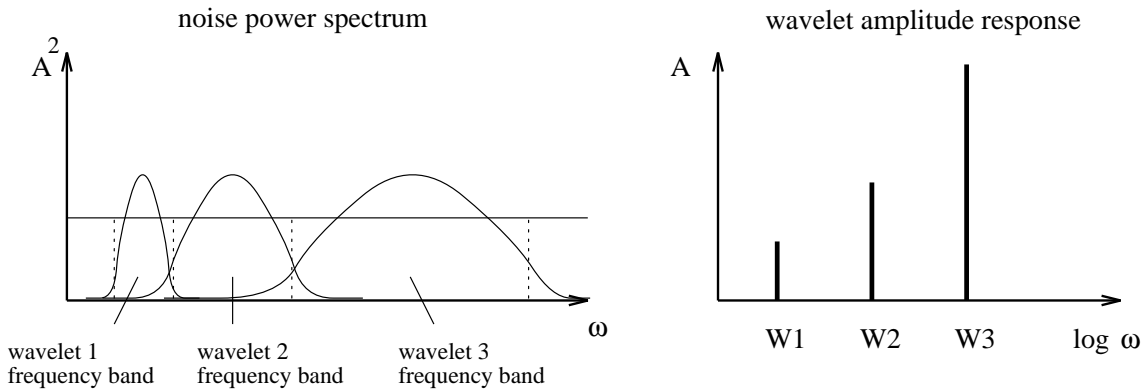


Figure 17: The amplitude response of a 1D wavelet to noise is proportional to the square root of its bandwidth (and hence the square root of its centre frequency) if the noise has a uniform power spectrum.

If the estimate of the mean noise response of the smallest scale wavelet pair over the whole image is given by \overline{A}'_0 , then the amplitude of the influence of noise on a filter at another scale can be estimated from the relative square roots of the filter bandwidths/centre frequencies.

$$A'_{n_{noise}} = \overline{A}'_0 \frac{\sqrt{f_n}}{\sqrt{f_0}} , \quad (15)$$

where $A'_{n_{noise}}$ is the estimated amplitude of the influence of noise on the filter at scale n , f_0 is the centre frequency of the smallest scale (highest centre frequency) filter and f_n is the centre frequency of the filter at scale n .

If m is the scaling factor between successive wavelets then the filter at the n^{th} scale will have a centre frequency given by

$$f_n = \frac{1}{m^n} f_0 . \quad (16)$$

If equations 15 and 16 are combined the magnitude of the total noise influence over all filter scales can be estimated with the following expression

$$T = k \overline{A}_0 \sum_{n=0}^{N-1} \frac{1}{\sqrt{m^n}} , \quad (17)$$

where N is the number of wavelet scales used, m is the scaling factor between successive filters and k is a scaling factor used to estimate the maximum amplitude of the noise response from the mean amplitude (typically $k \sim 2.5$).

The sum of the estimated noise responses over all the wavelet scales, T , will give us an upper bound on the effect of noise on the sum of the wavelet response amplitudes, $\sum_n A_n(x)$. It is an upper bound because in general the noise responses will not all be in phase; some will cancel. This bound on the effect of noise on $\sum_n A_n(x)$ is in turn an upper bound of its effect on the value of local energy, $E(x)$, as $\sum_n A_n(x)$ is always greater than, or equal to, $E(x)$ by a triangle inequality. 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 to noise. Thus, the expression for phase congruency is modified to the following:

$$PC(x) = \frac{[E(x) - T]}{\sum_n A_n(x) + \varepsilon} , \quad (18)$$

where $[\]$ denotes that the enclosed quantity is equal to itself when its value is positive, and zero otherwise.

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. This approach proves to be highly effective. Figure 19 shows the results of




Figure 18: Compensating for noise in the calculation of phase congruency. The background value of energy generated from the image noise results in a ‘noise circle’ of radius T . Phase congruency is now calculated by only using the amount by which the signal energy exceeds T before normalizing it by $\sum_n A_n(x)$.

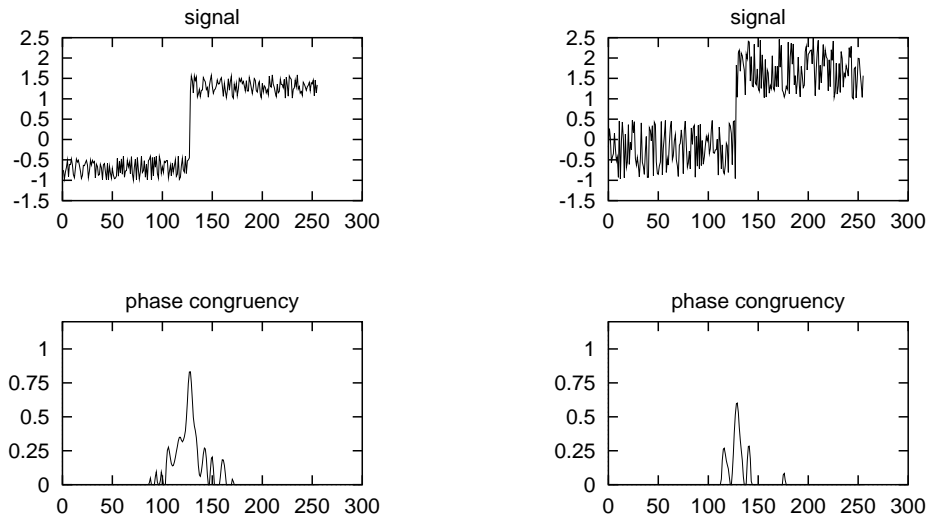


Figure 19: Noise compensated phase congruency of two step profiles; SNR 13.3 and 5.3 respectively.

processing two noisy step profiles. In both cases a k value of 2 was used to estimate the maximum influence of noise on local energy.

It should be re-emphasized here that this approach to noise compensation makes the assumption that the noise spectrum is uniform. In practice this will not be

exactly the case, particularly so with images from CCD cameras. ‘Smear’ in pixel values along the rows of an image from a CCD camera result in the noise content having greater amplitude at lower frequencies in the horizontal direction. However, looking in the vertical direction one does find that the noise spectrum is close to uniform as assumed. In all the results that are subsequently presented for 2D images in this thesis it has been assumed that the noise spectrum is uniform in all orientations. This has produced satisfactory results though no doubt improvements could be made if the spectral noise model was modified to match data measured from CCD cameras. Appendix B describes some experiments that were conducted to qualitatively assess the spectral properties of noise in images from CCD cameras.

3.5 Extension to two dimensions

So far the discussion has been limited to signals in one dimension. Calculation of phase congruency requires the formation of a 90° 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. There are four issues to be resolved: The shape of the filters in two dimensions, the numbers of orientations to analyze, the way in which the results from each orientation are combined, and the changes for noise compensation that are required in two dimensions.

3.5.1 2D filter design

The one dimensional filters described previously can be extended into two dimensions by simply applying some spreading function across the filter perpendicular to its orientation. The obvious spreading function to use is the Gaussian and there are good reasons for choosing it. Consider a 2D wavelet being convolved with a step edge feature that is not aligned with the orientation of the filter, as shown in Figure 20. If the filter is separable the convolution can be accomplished by a 1D convolution in the vertical direction with the spreading function, followed by a

1D convolution horizontally 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.




Figure 20: Convolution of a 2D wavelet with an angled edge.

In this example, the result of convolving the image vertically with a Gaussian spreading function will be to blur the edge so that on the subsequent convolution with the wavelet function we encounter a Gaussian smoothed step. Looking in the frequency domain, any function smoothed with a Gaussian suffers amplitude modulation of its components, but phase is unaffected (the transfer function of a Gaussian is a Gaussian). Thus, the phase congruency at the point of the feature is preserved. Ronse provides a formal proof of this result [78]. 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 sinc function) has negative values. The phase values would be corrupted and the step edge would be changed to a ramp; we would then perceive two Mach bands along the ramp edges instead of one step edge.

3.5.2 Filter orientations

To detect features at all orientations the bank of filters must be designed so that it tiles the frequency plane uniformly. In the frequency plane the filters appear as 2D Gaussians symmetrically or anti-symmetrically placed around the origin, depending on the spatial symmetry of the filters. The length to width ratio of the 2D 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