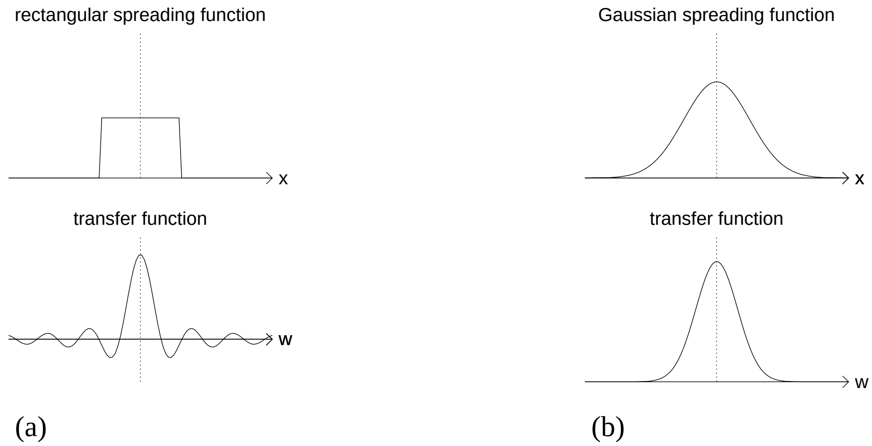

Figure 21:
(a) A rectangular spreading function has a transfer function that reverses some phase components.
(b) A Gaussian spreading function only modulates amplitudes.

2D spectrum.

A logical way to construct 2D filters in the frequency domain is to use polar-separable 2D Gaussians. In the radial direction, along the frequency axis the filters are designed in the same way as we have been designing 1D filters, that is, Gaussians with geometrically increasing centre 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

$$G(\theta) = e^{-\frac{(\theta-\theta_0)^2}{2\sigma_\theta^2}}, \tag{19}$$

where θ_0 is the orientation angle of the filter, and $\sigma_\theta = s\Delta\theta$ with s being a scaling factor and $\Delta\theta$ being the orientation spacing between filters. The scaling factor s, is set so that the overlap of the Gaussians in the angular direction is sufficient to ensure even coverage of the 2D frequency plane. In the results presented here a scaling factor value of 1.2 has been used. This value is not critical, though it has been found that when values greater than about 1.4 are used, the uneven spectral coverage that results starts to cause errors in the normalizing of local energy to produce phase congruency. A filter orientation spacing of 30^o 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 (see Appendix A). The final arrangement of filters results in a 'rosette' of overlapping polar-separable 2D Gaussians in the frequency plane (Figure 22). Simoncelli et al. [84] describe a systematic filter design technique for achieving uniform coverage of the frequency plane that could be applied here.

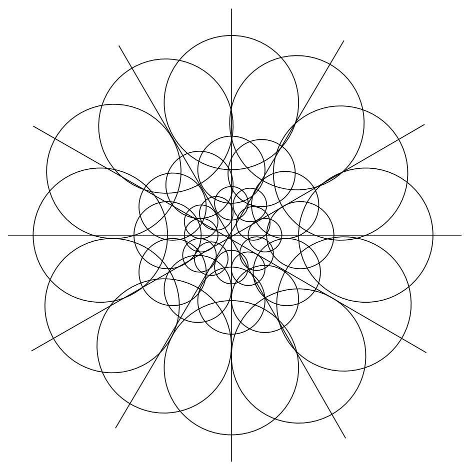

Figure 22: Tiling of the frequency plane with oriented filters at different scales.

3.5.3 Noise compensation in two dimensions

When compensating for noise in two dimensions one has to consider the spatial width of the filters in addition to their bandwidths. The spatial width of a filter affects its noise response. In the design of a filter rosette, the spatial width of a filter is varied with its centre frequency in order to maintain constant directional sensitivity. Viewed in the frequency plane the angular 'width' of the filters is proportional to their bandwidth resulting in the smallest scale filters gathering an even greater proportion of energy from the spectrum than they do in the 1D case (see Figure 22). Thus, in 2D, the square of the amplitude response in the spatial domain will be proportional to the *square* of the bandwidth multiplied by the power level of the noise. That is, in 2D, the relative amplitudes of the filter responses to noise will now be directly proportional to their bandwidth. Accordingly the noise compensation term, T now becomes

$$\begin{aligned} T &= k\,\overline{A}_0' \sum_{n=0}^{N-1} \frac{1}{m^n} \\ &= k\,\overline{A}_0' \frac{1-(\frac{1}{m})^N}{1-\frac{1}{m}} \,, \end{aligned} \tag{20}$$

where $\overline{A}_0' = e^{\overline{\log A_0(x,y)}}$ forms our estimate of the mean noise response of the smallest scale 2D filter pair over the image.

3.5.4 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, it is important to avoid making assumptions about the 2D form of features that one may encounter.

It is 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. If energy was normalized in each orientation prior to being summed, then the contribution from each orientation would be independent of its energy, clearly an undesirable situation.

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, 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 2D phase congruency:

$$PC(x) = \frac{\sum_o \lfloor E_o(x) - T_o \rfloor}{\sum_o \sum_n A_{no}(x) + \varepsilon} \tag{21}$$

where o denotes the index over orientations.

On a straight edge segment, only filters with orientations roughly perpendicular to the edge will have significant output; responses from other orientations will be negligible and the equation above will reduce to the 1D expression for phase congruency developed earlier. At more complex feature types such as corners and 'T' junctions, several orientations will have energy at significant levels. However, the sum of the amplitudes of the individual wavelet responses will also increase, maintaining the appropriate normalization.

Notice in the equation above that noise compensation is performed in each orientation independently. In practice this has been found to give significantly better results as it allows effective noise compensation in the case when the image noise content is anisotropic. The perceived noise content as deduced from the average response of the smallest scale wavelet pair can vary significantly with orientation. This is due to the correlation in noise along scan lines that is often observed, and other processes that can occur in the digitization of an image. Appendix B discusses this issue in more detail.

3.6 The importance of frequency spread

Clearly a point of phase congruency is only significant if it occurs over a wide range of frequencies. As an extreme example of a signal with a narrow frequency spread consider a pure sine wave: its Hilbert transform will be a cosine, the sum of the squares of these two signals will be one everywhere and thus phase congruency will be one everywhere. Clearly, in this situation phase congruency is *not* acting as a feature detector. A more common situation is where a feature has undergone Gaussian smoothing, either intentionally or through image degradation. For these smoothed functions phase congruency can be high over extended regions of the signal resulting in poor localization. The Gaussian smoothing reduces the high frequency components in the signal and accordingly 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.

From Figure 23 we can see that to obtain good localization it is important to incorporate low frequency components in the calculation of phase congruency on smoothed profiles. These low frequency components are the least affected by the smoothing of the signal.

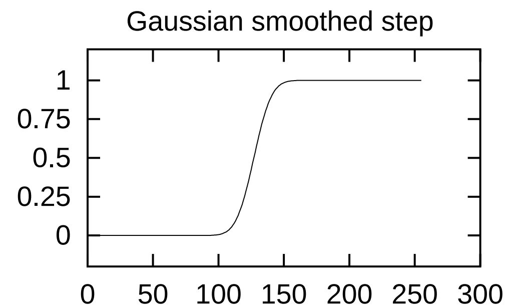

(a) Smoothed step profile.

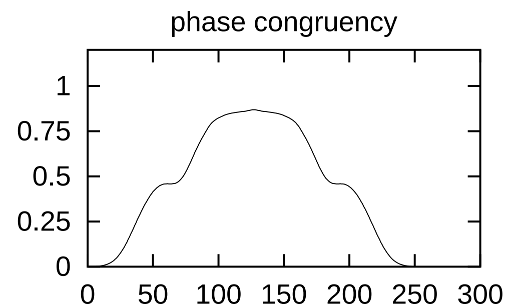

(b) Phase congruency considering wavelengths up to 65 units. The undulations in phase congruency are due to numerical effects.

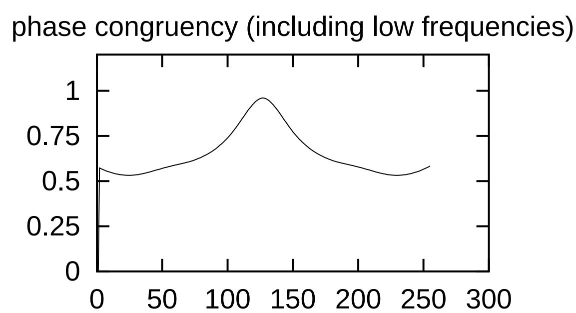

(c) Phase congruency considering wavelengths up to 150 units.

Figure 23: Phase congruency of a Gaussian smoothed step profile at two different scales of analysis.

Thus we are concerned with two conditions: one is where there is a very narrow range of frequencies present causing the values of energy and $\sum_n A_n$ to become equal, making phase congruency one. The other is where only a limited range of frequencies is considered, as in Figure 23(b), and one encounters no significant frequency content at all. In this situation both energy and $\sum_n A_n$ become very small, making the calculation of phase congruency poorly conditioned. This second situation is handled through the use of the parameter ε, in the denominator of the expression for phase congruency.

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. On a spectrum of this shape each wavelet will gather an amount of power that is inversely proportional to its bandwidth. (The power level is inversely proportional to the centre frequency squared, but the bandwidth over which the power is gathered is proportional to centre frequency). In a manner analogous to the analysis of filter amplitude responses to noise we can now deduce the expected distribution of filter amplitude responses to a square waveform. However, unlike the situation with noise, a filter's response to a step feature is localized in space. The spatial width of the response is dictated by the spatial width of the filter, which is inversely proportional to its bandwidth. So for each filter, while the power it captures falls with bandwidth, the spatial width over which that power is expressed in terms of the square of the amplitude response, also falls with bandwidth. Therefore, to preserve the expected power gathered by each filter the magnitudes of the squared amplitude response have to remain constant, independent of filter centre frequency. Hence, the expected distribution of amplitude responses to a step feature will be a uniform one. This is illustrated in Figure 24. Field [20] points out that in many cases images of natural scenes have overall power distributions that are inversely proportional to the frequency squared, 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. This is likely to maximize the precision of any computation (numerical or neural) that we make with the filter outputs.

It should be noted that the Gaussian smoothing of a feature that arises whenever

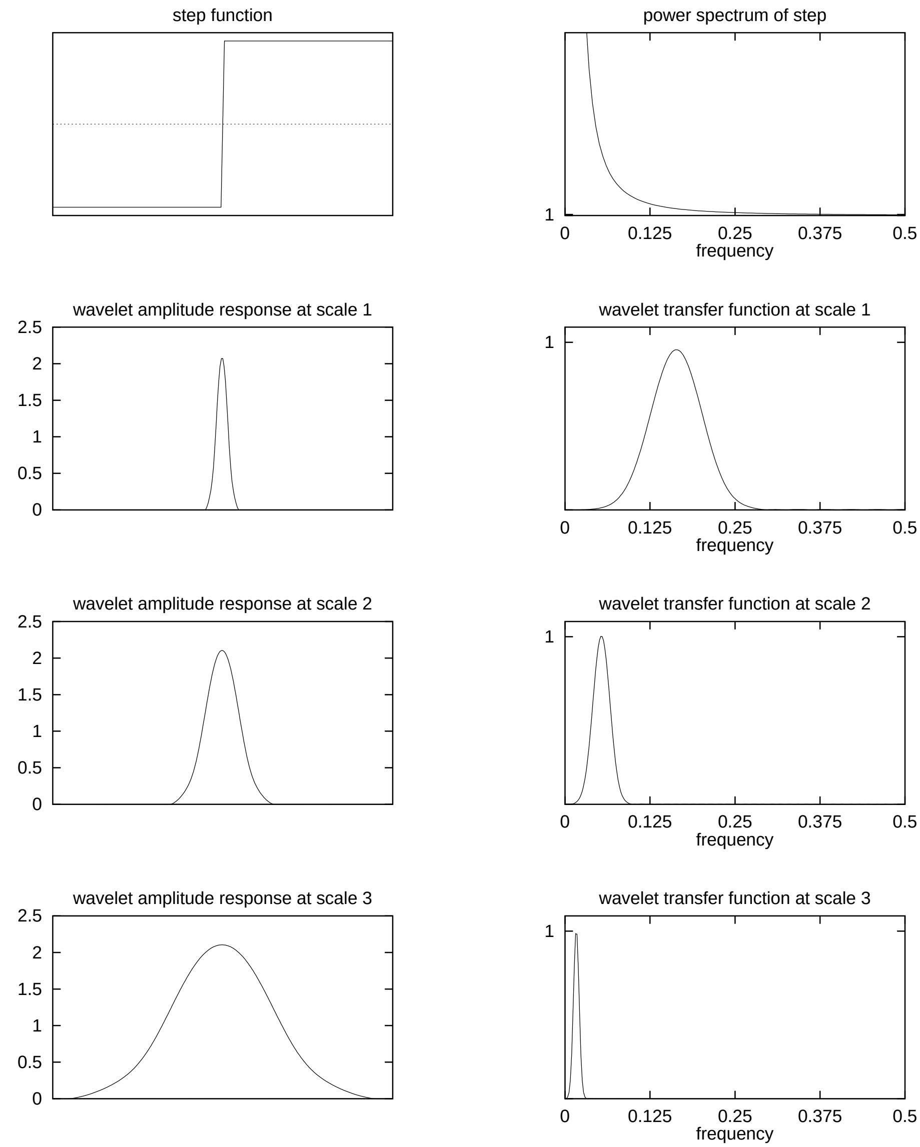

Figure 24: On a step feature, wavelet pairs at all scales will respond with the same amplitude.

a 2D filter having a Gaussian spreading function encounters a feature at some non-orthogonal angle (see Figure 20) *does not* significantly change the frequency distribution relative to what would be expected from an unsmoothed feature. The reason for this is that the degree of Gaussian smoothing, as seen by each filter, varies directly in proportion with the scale of each filter due to its fixed length/width ratio. Consider the convolution process in the frequency domain: A filter at some given centre frequency will see a feature spectrum that has been modulated by a Gaussian of some spread. A second filter tuned to a frequency of twice the first will see a feature spectrum that has been modulated by a Gaussian of twice the spread (spatially, its spreading function is a Gaussian of half the width of the first). The relative modulation of the feature spectrum as seen by each individual filter over its bandwidth will therefore remain roughly constant. This is illustrated in Figure 25. Thus, the overall distribution of filter responses, while modulated, will remain largely unchanged.

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. Following an analysis similar to that done for the square waveform one can show that for a delta feature the amplitude of the wavelet filter responses will be proportional to their bandwidths, and hence their centre 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 like $1/\omega^4$. When considered in terms of the amplitude of filter responses on a logarithmic frequency scale this will produce a distribution strongly skewed to the low frequency end.

Thus, the difficulty faced here is that there is no one ideal distribution of filter responses. Given the very different types of distributions expected for delta and roof edges, we cannot even say the ideal distribution should be a function of the phase angle at which the congruency occurs. 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 a common feature type 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

$$s(x) = \frac{1}{N}\left(\frac{\sum_n A_n(x)}{A_{max}(x) + \varepsilon}\right) , \qquad (22)$$

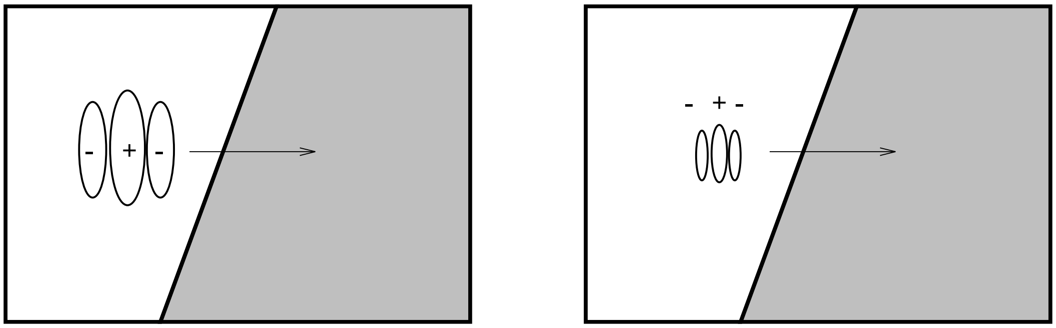

(a) Filters at scales differing by a factor of two encountering a step edge at an angle.

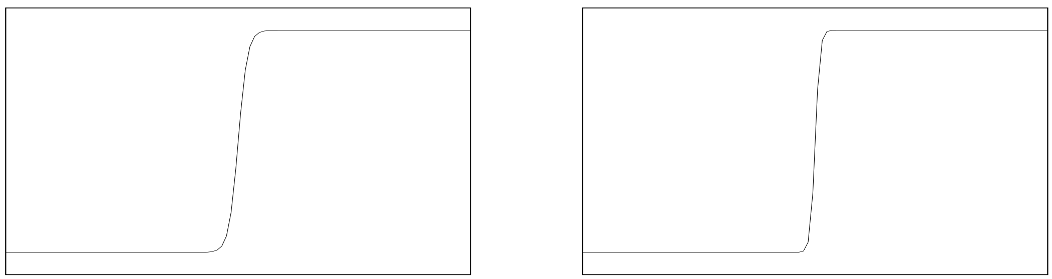

(b) Equivalent 1D signals seen by each filter after 1D convolution with each filter's vertical spreading function. The broad scale filter sees a more heavily smoothed feature.

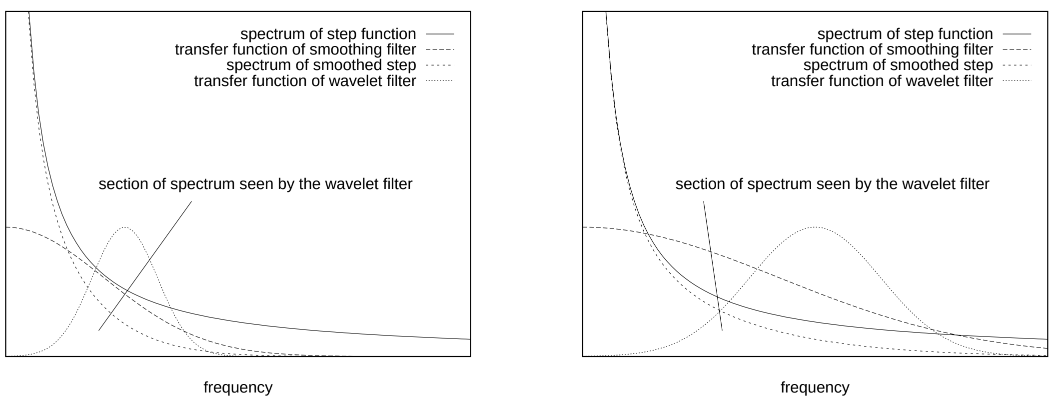

(c) Processing of the respective signals in the frequency domain.

Figure 25: When filters having fixed length-width ratios are applied to an angled step edge the degree of smoothing of the feature resulting from the filter's spreading function will vary with the scale of the filter. The effect of this is that the shape of the distribution of filter amplitude responses will not be greatly different from that obtained when the filter encounters a perpendicularly oriented edge.

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[2]. If all the filter responses are equal we have the maximum possible frequency spread, and the spread measure becomes one. If only one filter response is non-zero (minimal spread) the spread measure falls to $1/N$.

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

$$W(x) = \frac{1}{1 + e^{g(c - s(x))}} , \tag{23}$$

where c is the 'cut-off' value of filter response spread below which phase congruency values become strongly penalized, and g is a gain factor that controls the sharpness of the cut-off. Typical values for c and g are 0.4 and 10 respectively; a plot of the weighting function with these parameters is shown in Figure 26. Note the sigmoid function has been merely chosen for its simplicity and ease of manipulation.

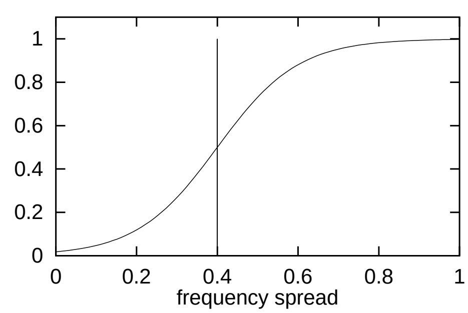

Figure 26: Frequency spread weighting function with a cut-off value of 0.4 and g value of 10.

In 2D, the weighting of phase congruency by frequency spread at each point in an image has to be done in each orientation independently. Thus the weighting factor is applied to the energy in each orientation prior to summing over all orientations and normalizing by the total sum of the filter response amplitudes. We modify equation 21 to produce the following:

$$PC(x) = \frac{\sum_o (W_o(x) \lfloor E_o(x) - T_o \rfloor)}{\varepsilon + \sum_o \sum_n A_{no}(x)} . \tag{24}$$

Figure 14 provides an example of how the frequency spread weighting function varies across a 1D signal. Weighting by frequency spread, as well as reducing

[2]The presence of the small constant ε to avoid division by zero will, of course, mean that $s(x)$ can never attain its ideal maximum value of 1.

spurious responses where the frequency spread is low, has the additional benefit of sharpening the localization of features, especially those that have been smoothed. Referring to Figure 14, it is of interest to note that peaks in $E(x)$, $\sum_n A_n(x)$, and to a lesser extent, $W(x)$ are all strong indicators of the presence of a feature, though it is phase congruency that provides a dimensionless quantity that is invariant to contrast.

3.7 Scale via high-pass filtering

3.7.1 Difficulties with low-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 [54]. While this approach is intuitive, the justification for assuming the brain uses these band-passed versions of the image *directly* for multi-scale 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 only 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 multi-scale 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. This is a problem for differential based feature detection schemes that require smoothing to suppress the influence of noise. It also creates difficulties for applications that make use of coarse-to-fine strategies for feature matching. Feature positions drift with scale, localization at broad scales is poor and the number of features proliferate as scale is reduced. This problem was recognized by Marr and a good illustration of this is found in his book[3] which shows zero crossings in an image of a Henry Moore sculpture at three different scales. In many parts of the image zero crossings at one scale do not correspond to any other

[3]Marr. *Vision*. 1982, Figure 2-20, page 69

zero crossings at other scales. In his caption for this sequence of images Marr says

> "... This set of figures neatly poses the next problem - How do we combine all this information into a single description?"

Marr proposes a coarse-to-fine strategy to resolve this problem, but he has difficulties in constructing a truly systematic approach and recognizes that many special cases would have to be accounted for.

Witkin [97] introduces the idea of a *scale-space* image in an attempt to resolve this problem, that is, different settings of the scale parameter result in different descriptions of the scene, and notes that we have no way of knowing what scale of analysis is 'correct'. He argues that signals should be analyzed over a *continuum* of scales to produce a scale-space image. He also specifies the requirement that while additional zero crossings in the scale-space image may appear as scale is decreased, existing ones should not disappear. Witkin shows that the Gaussian is the only smoothing kernel that guarantees this behaviour. Then, by considering the trajectories of zero crossings over scale he suggests that the scale-space image can be collapsed into a tree structure that provides a concise, qualitative description of the signal over all scales. Witkin also makes the observation that features that we consider most salient in a signal seem to correspond to zero crossings that remain stable over wide ranges of scale.

Koenderink [46] develops Witkin's analysis of 1D signals and extends them to 2D images. He recognizes the connection between smoothing in scale-space and the diffusion equation and this allows him to develop a detailed study of the behaviour of scale-space in 2D.

While Witkin and Koenderink advocate the use of zero crossings in scale-space for representing signals Hummel's work [38] suggests that there are considerable difficulties in this approach. Hummel investigates the representation of signals in terms of their zero crossings in scale-space and one's ability to reconstruct the signal from this data. He concludes that for general signals this representation is unstable. Further, he shows that even when gradient information at the zero crossings is used the representation is still unstable.

Bergholm [5] adopts the scale-space model in developing his *edge focusing* approach to edge detection (described in Chapter 2). He makes a detailed study of the way feature positions are corrupted by Gaussian smoothing, this information being essential to the design of his coarse to fine strategy of finding edges. However, he reports difficulties in tracking blurred edges where the final edge points end up being scattered and broken up, presumably because scale-space becomes degenerate at the finest scales on a heavily blurred feature.

Perona and Malik [69] make a strong critique of the standard scale-space model. They are particularly concerned with the fact that the position of a feature is not known unless it is tracked through scale-space to the finest level and they argue that feature positions should be stable over scale. They develop an edge detection technique that is inspired by Koenderink's suggestion that the diffusion equation could possibly be run backwards in time to enhance an image; this however, is an ill-posed problem. Instead Perona and Malik devise a method to run the diffusion equation forwards in time to enhance edges. Their approach was described in Chapter 2.

3.7.2 High-pass filtering

It is clear that there are many problems with the concept of scale which have not been satisfactorily resolved. In particular, the idea that a feature's position should vary with the scale of analysis is very unsatisfactory. It is argued here that *feature location should not be a function of the scale of analysis, only the relative significance of features should change.* Subsequent image operations such as coarse-to-fine strategies would be greatly facilitated (or not even needed) if only the relative significance of features varied with scale, with the number of features and their locations remaining stable.

Much of the thinking about scale has been strongly influenced by the problems of applying differential operators to images. 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 the 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. (Filters in the wavelet bank range from some minimum spatial scale, say corresponding to the Nyquist frequency, up to a scale matching the window size). With this approach *high-pass* filtering is being used to specify the analysis scale. Low frequency components are being cut out (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 only be comparing each feature to a small number of other features that are nearby, and hence each feature is likely to be perceived as being 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. Something that may have high phase congruency/significance when analyzed over a small window may end up having a low measure of phase congruency if we look at it within a larger analysis window and consider the phase of the lower frequencies present. A feature having high phase congruency when analyzed over a wide range of spatial scales from small to large is clearly of greater significance than one having congruency over a limited number of small spatial scales.

It should be noted that the 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 the corresponding changes that would occur in the intensity gradients of the image, and hence the in output of a gradient based edge detector.

In summary, it is proposed that multi-scale 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 cut-off 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. Figure 27 illustrates a one-dimensional signal at three different scales of low-pass, band-pass and high-pass filtering, along with phase congruency at the three high-pass scales.

The low-pass signals were obtained by convolving the signal with Gaussian masks of different widths. The band-pass signals were obtained through convolution with even Gabor filters. The high-pass signals were obtained by summing the convolution results from even Gabor filters ranging from the finest scale (wavelength 2 pixels) down to varying levels of coarser scale. In addition to the 1D example shown here Appendix A contains experimental results where the stability of phase congruency on 2D images at different levels of high-pass filtering is demonstrated.

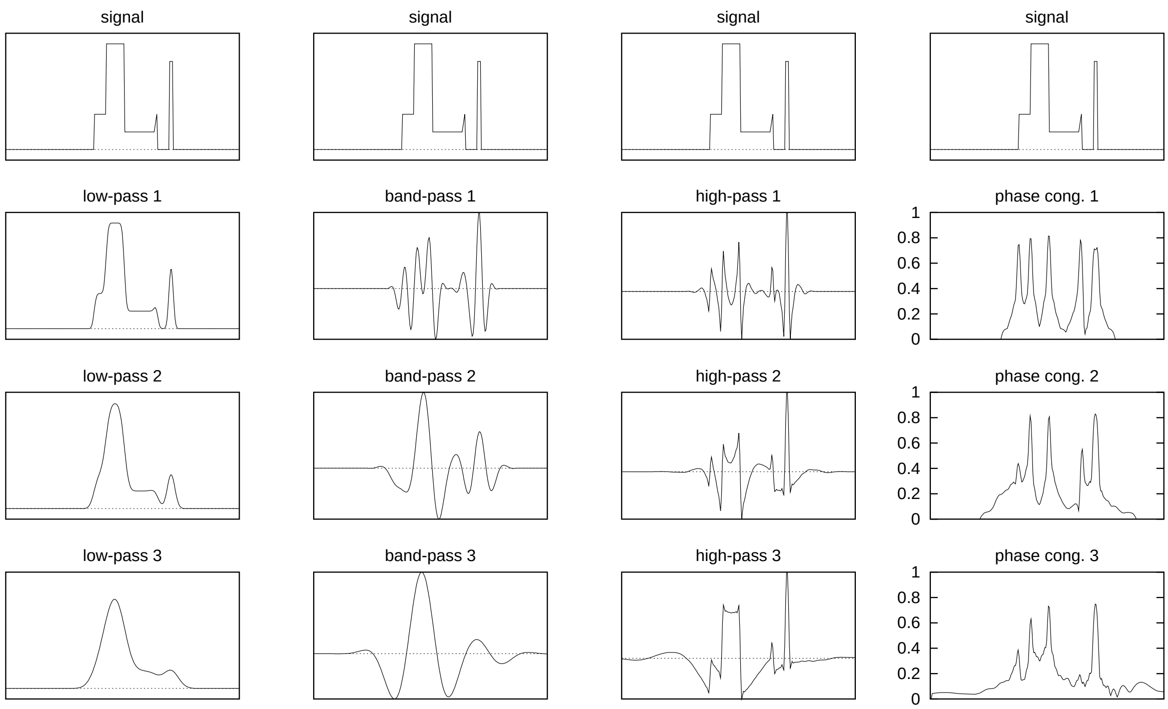

(a) low-pass (b) band-pass (c) high-pass (d) phase congruency

Figure 27: Analysis at different scales.
(a) Three different low-pass versions of a 1D signal.
(b) Three different band-pass versions of a 1D signal.
(c) Three different high-pass versions of the signal.
(d) Phase congruency at the three scales of high-pass filtering shown in (c).
Note that the number and location of features as measured by phase congruency remains constant and only their relative significance varies. Under low-pass and band-pass filtering the number and locations of features varies.

3.7.3 High-pass filtering and scale-space

There is a mirror symmetry in the relationship between the classical low-pass filtering model of scale-space and the high-pass model. In the low-pass model the low scale image is the one that has been heavily smoothed, the highest scale image corresponds to the original image. On the other hand, the high-pass approach to scale considers the original image as representing the lowest scale of image with higher scales obtained via high-pass filtering. This is illustrated by Figure 28.

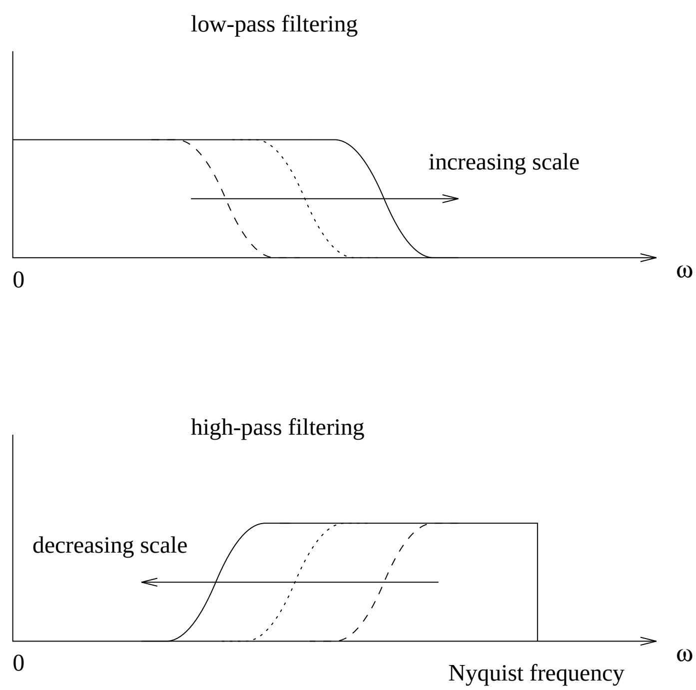

Figure 28: Filter transfer functions for low-pass and high-pass approaches to scale-space. With low-pass filtering the finest scale is achieved when the whole of the image spectrum is included. For high-pass filtering this situation is considered to represent the *lowest* scale of image.

If we look at the phase scalogram of a test profile (Figure 29) we can gain further insights in the differences between using high and low-pass filtering for multi-scale analysis. The lines of phase congruency corresponding to the profile's feature points are readily seen in the scalogram. They all start at the finest scale (at the top of the scalogram) and extend downwards. The extent to which these lines of phase congruency extend down specifies the overall significance of a feature.

High-pass filtering corresponds to extracting sections of the scalogram from the top downwards. Thus, at a fine scale of analysis with high-pass filtering one will

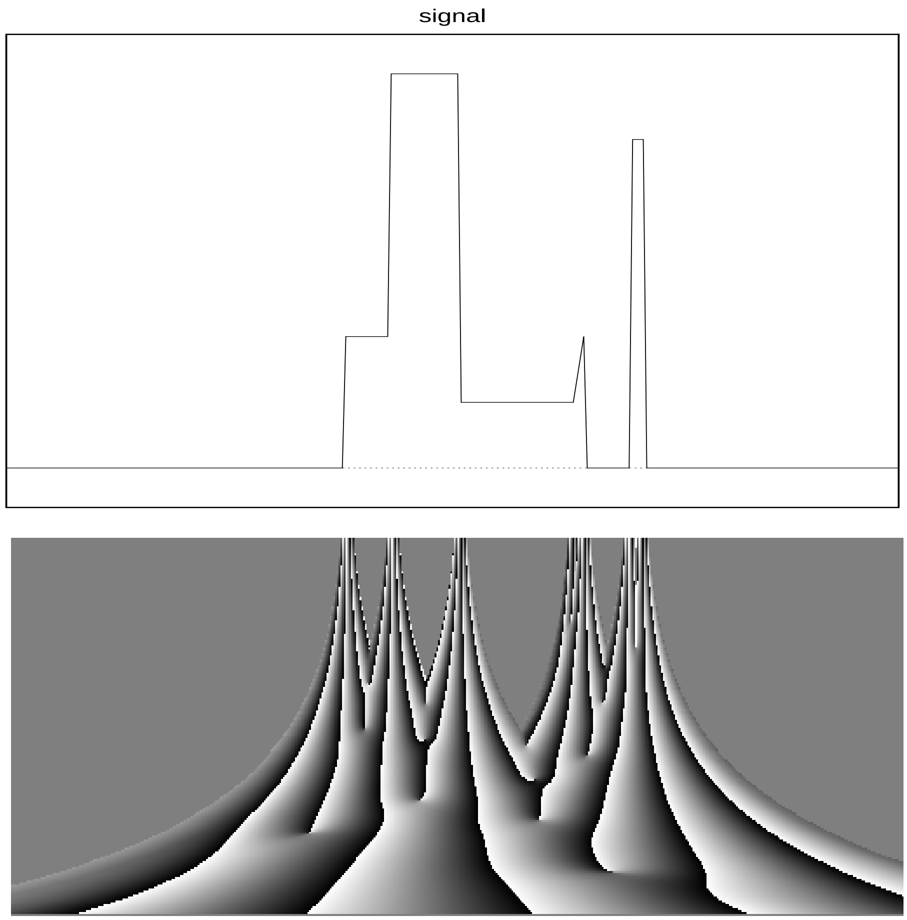

Figure 29: Test profile and its phase scalogram. Highest frequencies are at the top of the scalogram and lowest at the bottom.

be considering a narrow section of the scalogram at the top. All lines of phase congruency will span the full width of the section being considered and hence will all be considered to be equally significant. As the analysis section of the scalogram is progressively widened downwards some lines of phase congruency will terminate. For these features the measured phase congruency will be reduced and their relative significance will fall. The perceived location of these features, given by the point of maximal phase congruency, remains stable. Thus we have an interpretation of scale where only relative feature significance changes with the scale of analysis, not their locations.

In contrast to this approach low-pass filtering corresponds to extracting sections of the scalogram from the bottom upwards. Analysis of a heavily smoothed image corresponds to the analysis of a narrow section at the bottom of the scalogram.

Here one has only the vaguest idea of what features are present and what their locations are. As the analysis section is progressively extended upwards (in a coarse to fine strategy) new features appear and existing feature positions drift. Clearly, interpretation of the signal with this approach is far more difficult.

3.8 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 [11, 12], Spacek [87] and Deriche [16]. 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. They are only concerned 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 is unaffected by image contrast and scale. 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.

Figure 30 illustrates some results on a test image containing line features and step discontinuities of various magnitudes and orientations. For comparison, the output of the Canny detector is also presented [11]. The implementation of the Canny detector used here follows the modifications suggested by Fleck [24]. The raw, gradient magnitude image is displayed so that comparison can be made without having to consider any artifacts that may be introduced by non-maximal suppression and thresholding processes. The purpose of providing this comparison is to illustrate some of the qualitative differences in performance between the two detectors. It is hard to make meaningful quantitative comparisons as 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.

In all the results presented here a wavelet bank over 6 scales was used, the smallest filters having a period of 3 pixels, and the scaling factor between successive wavelets being 1.5. Filters were constructed directly in the frequency domain as 2D Gaussians in the frequency plane. For each filter the ratio between the centre frequency and the Gaussian standard deviation was set at 3 (see Appendix D for implementation details). The wavelets had a length to width ratio of one and their orientational spacing was 30^o. The final output was non-maximally suppressed and hysteresis thresholding applied between phase congruency values of 0.5 to 0.3. A noise compensation k value of 2.5 was used, and ε was set at 0.01. For the frequency spread weighting function the cut-off fraction, c was set at 0.4 and the gain parameter, g was set at 10. These parameters have been found to give reasonable results over a wide range of images. The non-maximal suppression technique that was used to produce the phase congruency edge maps is described in Appendix C. The Canny edge detector output was calculated after smoothing the images with a Gaussian mask having a standard deviation of 1 pixel.

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 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 centred on the line (this problem was recognized by Canny and he designed a separate operator to detect line features).