Phase singularities in scale-space

Allan D Jepson and David J Fleet*

This paper concerns the use of phase information from band-pass signals for the measurement of binocular disparity, optic flow and image orientation. Towards this end, one of the important properties of band-pass phase information is its stability with respect to small geometric deformations and contrast changes. However, in particular regions phase can also be very unstable due to the occurrence of phase singularities. We discuss the existence of phase singularities, and their relation to the neighbourhoods where phase is unreliable. Moreover, we present a simple method for detecting these regions of instability.

Keywords: phase information, stability, scale-space

INTRODUCTION

In order to compute image velocity or binocular disparity, it is necessary (in some sense) to localize structure in an image sequence and track it across frames, or match it between left and right stereo views. For example, differential-based velocity techniques measure the translation of level contours of either constant intensity, or constant amplitude of filter response[1–5], while zero-crossing approaches focus on the motion of zero-crossings in the output of band-pass filters[6,7]. Recently, the use of contours of constant phase has been suggested for the measurement of binocular disparity[8–11] and of image velocity[12]. In choosing what type of structure to track, it is important to consider its stability under common image deformations such as contrast variations, dilations, shears, and rotations, in addition to simple translation (cf. Verri and Poggio[13]). One main advantage of phase information is that, except near certain points referred to here as *singularities*, phase is generally stable with respect to small affine image deformations.

In this paper we illustrate the stability of phase information as compared to the amplitude of filter output for 1D signals through the use of a scale-space representation. Note that we are not concerned here with the representation of images, as in the traditional scale-space work[14], but rather with tracking structure across frames. Therefore, we are primarily concerned with small changes in scale (e.g. up to 20% dilations). In addition, we discuss the existence of phase singularities, the neighbourhoods about them where phase is unreliable, and we present a simple method for their detection. Given this detection scheme, highly accurate and robust approaches to the measurement of optic flow and binocular disparity are possible. For example, based on the spatiotemporal gradient of phase information, Fleet and Jepson[12] reported a technique for the measurement of component (normal) image velocity for which approximagely 90% of the accepted estimates had less than 5% relative error in cases of significant dilation and shear.

The results presented here are of general interest for several reasons. First, they apply to zero-crossings of the filter output in that zero-crossings in the response of band-passed filters can be viewed as lines of constant phase. Second, similar results apply to 2D signals. Third, the problems caused by deviations from image translation do not exist solely for phase-based techniques (cf. Schunck[15] and Verri and Poggio[13]). The fact that the geometric deformations can be handled robustly within a phase-based framework should be viewed as a major advantage.

BAND-PASS SCALE-SPACE EXPANSIONS

To demonstrate the robustness of phase and its singularities we first consider a *scale-space* expansion of a 1D signal that expresses the output of a band-pass filter $K(x, \lambda)$, as a function of spatial position and the principal wavelength λ, to which the filter is tuned. In what follows, the filters are assumed to be band-pass and constant phase. The kernels $K(x, \lambda)$ are complex where the real and imaginary parts form quadrature pairs (they are Hilbert transforms of one another). Furthermore, we assume translational invariance and, for convenience, self-similarity across scales so that the octave bandwidth is constant (cf. Mallat[16]). Bandwidths of approximately one octave are common in vision applications. Finally, it is assumed that the

Department of Computer Science, University of Toronto, Toronto, Canada M5S 1A4
*Department of Computing and Information Science, Queen's University, Kingston, Canada K7L 3N6
Paper received: 30 October 1990.

0262–8856/91/005338–06

filters' output should be subsampled in both space and scale to maintain an efficient, yet complete, representation of the input.

The scale-space expansion of an input signal $I(x)$ is defined by:

$$S(x, \lambda) = K(x, \lambda) * I(x), \tag{1}$$

where $*$ denotes convolution. Because $K(x; \lambda) \in$, the response $S(x, \lambda) = \mathrm{Re}[S(x, \lambda)] + i\,\mathrm{Im}[S(x, \lambda)]$ can be written in terms of amplitude and phase components as $R(x, \lambda) = \rho(x, \lambda)\, e^{i\phi(x,\lambda)}$, where:

$$\rho(x, \lambda) = |S(x, \lambda)| = \sqrt{\mathrm{Re}[S(x, \lambda)]^2 + \mathrm{Im}[S(x, \lambda)]^2}, \tag{2a}$$

$$\phi(x, \lambda) = \arg[S(x, \lambda)] = \mathrm{Im}[\log_e S(x, \lambda)] \in (-\pi, \pi]. \tag{2b}$$

The scale-space defined by (1) is similar to band-pass expansions defined by Gabor functions or the Laplacian of a Gaussian, although it is expressed in terms of analytic complex kernels instead of real kernels (cf. Mallat[16], Witkin[14], Koenderink[17] and Papoulis[18]).

Following Whitham[19] or Papoulis[18], the local (instantaneous) frequency of the response can be defined as the spatial derivative of the phase signal, $k(x) \equiv \phi_x(x, \lambda)$. If the phase of $S(x, \lambda)$ is linear, as in $\phi(x, \lambda) = k_0 x$, then $S(x, \lambda)$ is simply an amplitude-modulated sinusoid with constant frequency k_0. Otherwise, the phase derivative $\phi_x(x, \lambda)$ provides a local, constant-frequency approximation to $S(x, \lambda)$. In order to compute $\phi_x(x, \lambda)$ we can make use of the following identity:

$$\phi_x(x, \lambda) = \frac{\mathrm{Im}[S^*(x, \lambda)\, S_x(x, \lambda)]}{|S(x, \lambda)|^2}, \tag{3}$$

where $S^*(x, \lambda)$ denotes complex conjugate of $S(x, \lambda)$. Note that this equation shows how the phase derivative ϕ_x can be computed without first unwrapping the phase.

As mentioned above, it is important that the signal property to be tracked is *stable* with respect to scale perturbations. This suggests that its level contours should be vertical in scale-space. To see this, consider two 1D signals (e.g. left and right views), where one is a dilation and translation of the other, i.e. let:

$$I_r(a(x)) = I_l(x), \quad \text{where } a(x) = a_0 + a_1 x, \tag{4}$$

with a_1 near 1. Because the filters have constant bandwidth the two output signals will satisfy:

$$S_r(a(x), \lambda_1) = S_l(x, \lambda_2), \quad \text{where } \lambda_1 = a_1 \lambda_2, \tag{5}$$

i.e. the two outputs would have similar structure if filters tuned to λ_1 and λ_2 had been applied to $I_r(x)$ and $I_l(x)$, respectively. However, in measuring disparity (or velocity) it is common to apply the same filters to $I_r(x)$ and $I_l(x)$ because the scale factor a_1 is unknown. In other words, we attempt to recover $a(x)$ by matching structure (features) of $S_r(a(x), \lambda_1)$ and $S_l(x, \lambda_1)$. To be successful, the structure of $S_l(x, \lambda_2)$ that is used for matching must be well represented by the structure of $S_l(x, \lambda_1)$. Equivalently, its level contours should be nearly vertical in scale-space.

For illustrative purposes, let the band-pass kernel be $K(x, \lambda) = Gabor(x; \sigma(\lambda), k(\lambda))$, where:

$$Gabor(x; \sigma, k) \equiv e^{ixk} G(x; \sigma) \tag{6}$$

denotes a Gabor kernel[20] where $G(x; \sigma)$ is a Gaussian†. The peak tuning frequency of the filter is given by:

$$k(\lambda) = \frac{2\pi}{\lambda}. \tag{7a}$$

With a bandwidth of β octaves (assumed to be near 1), and with one standard deviation as a measure of the extent of the Gaussian envelope, the standard deviation of the amplitude spectra $\sigma_k = \sigma^{-1}$ should satisfy $\beta = \log_2[(k(\lambda) + \sigma_k)/)k(\lambda) - \sigma_k)]$. From this, it is easy to show that the approximate radius of spatial support is given by:

$$\sigma(\lambda) = \frac{1}{k(\lambda)} \left(\frac{2^\beta + 1}{2^\beta - 1}\right). \tag{7b}$$

Figure 1 (top) shows a signal composed of a sample of white Gaussian noise concatenated with a scanline from a real image. The two middle images show $\rho(x, \lambda)$ and $\phi(x, \lambda)$, generated with a Gabor kernel with $\beta = 1.0$, over an interval of two octaves for λ. Level contours of constant amplitude and phase are shown below. It is clear from Figure 1 that $\rho(x, \lambda)$ depends significantly on scale because its level contours are not generally vertical. As a consequence, the amplitude signal and the raw filter output (which has a significant amplitude component) are not stable candidate properties for matching. In contrast note that, except for several isolated regions, phase is generally stable with respect to scale perturbations. As explained below, the major source of this instability is the occurrence of singularities in the phase signal $\phi(x, \lambda)$. Moreover, a simple method is suggested for detecting these regions of instability.

SINGULARITY NEIGHBOURHOODS

For a general image $I(x)$, the scale-space defined by (1) is analytic, and contains a number of isolated zeros, where $S(x, \lambda) = 0$. In $\rho(x, \lambda)$ shown in Figure 1, zeros appear as black spots. The phase signal in (2b) is also analytic, except at the zeros of $S(x, \lambda)$. It can be shown that, for white noise input, the expected density of *phase singularities* about a given scale λ_0, is proportional to λ_0^{-2} [21].

In order to describe the effects of phase singularities, we begin with a simple model that generates singularities. This model consists of two complex exponentials, the amplitudes of which vary as a function of scale. For example, let the input $I(x)$ be the superposition of two waveforms at frequencies k_0 and k_1 where $k_1 > k_0$.

†Strictly speaking, the real and imaginary Gabor parts do not form quadrature pairs, i.e. $\langle \mathrm{Re}[Gabor], \mathrm{Im}[Gabor] \rangle \neq 0$. But for sufficiently small octave bandwidths (e.g. an octave or less), they are close enough to illustrate the main points clearly.

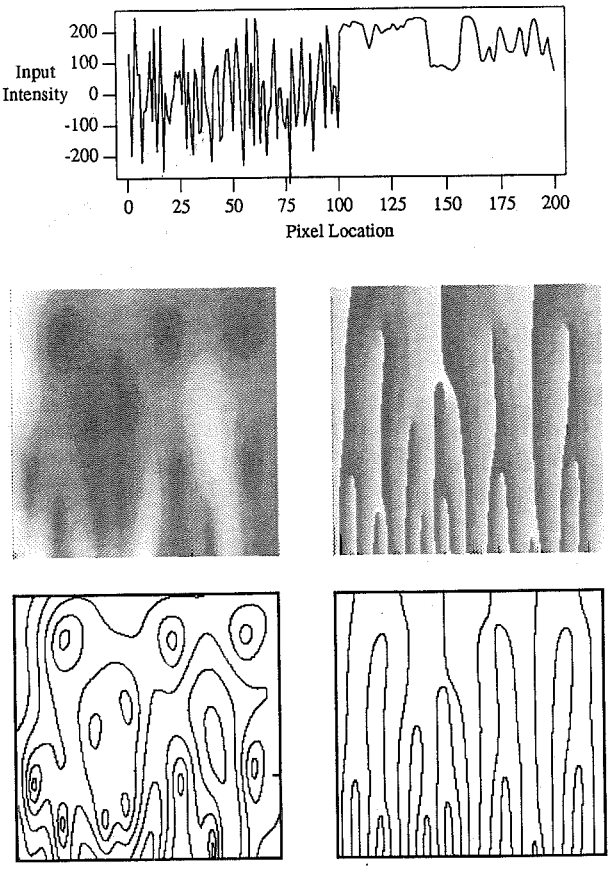

Figure 1. Gabor scale-space expansion: the input signal (top) consists of white Gaussian noise (left), and a scanline from a real image (right). The middle images show the amplitude and phase components of $S(x, \lambda)$ for $12 \leq \lambda \leq 48$ (in pixels). The vertical and horizontal axes represent scale and spatial position. The bottom two figures show level contours of constant $\rho(x, \lambda)$ and $\phi(x, \lambda)$. The small tick on the right side of the amplitude contours (lower left) marks scale $\lambda = 20$ for reference in other figures.

Then, the output of a band-pass filter with its peak tuning frequency $k(\lambda)$ somewhere between k_0 and k_1, is given by:

$$S(x, \lambda) = \rho_0(\lambda)\, e^{ik_0 x + \psi} + \rho_1(\lambda)\, e^{ik_1 x}, \tag{8}$$

where ψ is an arbitrary phase shift. Because of the roll-off of the amplitude spectrum away from the peak tuning frequency, an increase in scale, and therefore a decrease in $k(\lambda)$, causes variation in the two amplitude coefficients. In particular:

$$\frac{\mathrm{d}}{\mathrm{d}\lambda}[\rho_0/\rho_1] > 0, \tag{9}$$

so the relative amplitude of the lower frequency component becomes larger as the scale parameter increases. Singularities occur when $\rho_0(\lambda) = \rho_1(\lambda)$, where the signal passes through the origin. Figure 2 shows a simple example in which $k_1 = 2k_0$. The four panels show the signal in the complex plane as a function of x, for successively larger ratios of ρ_0/ρ_1.

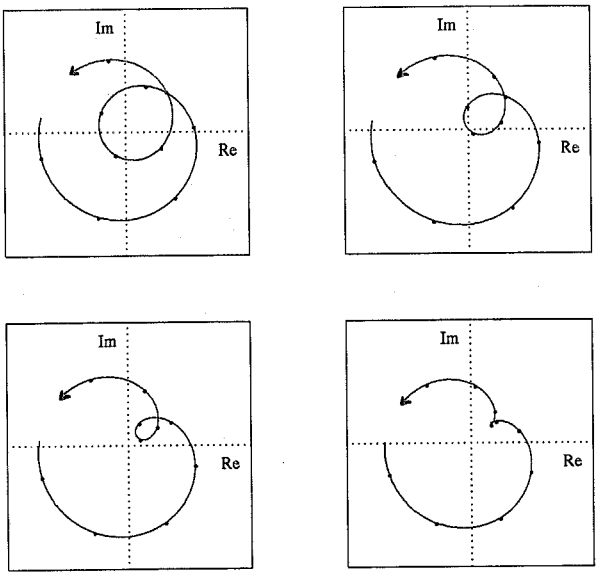

Figure 2. Simple model with a phase singularity: these four plots show four instances of the simple model (8) at successively larger scales. The signal shown is a function of spatial position in the complex plane. The actual singularity occurs in the top-right figure where the signal passes through the origin. The bullets marks fixed spatial positions for illustrative purposes, and the arrows indicate the direction for increasing x values

The arrows indicate the direction for increasing values of x. The upper-left panel illustrates the signal behaviour for a scale just below a singularity; that is for $\rho_0 < \rho_1$ and $\lambda < \lambda_0$, where λ_0 denotes the scale at which the singularity occurs. The upper-right panel illustrates the case $\lambda = \lambda_0$, for which the singularity is a result of $S(x, \lambda_0)$ passing through the origin. The bottom two panels show the behaviour at two scales above λ_0, where the relative amplitude of the higher frequency signal component decreases further. The sequence as a whole illustrates the general mechanism through which higher frequency information is lost as one moves up in scale-space.

This model also helps illustrate the behaviour of phase in the neighbourhood of singularities. As we show below, these neighbourhoods can be characterized in terms of properties of the complex logarithm of the filter response $S(x, \lambda)$, i.e. $\log_e S(x, \lambda) \equiv \log_e \rho(x, \lambda) + i\phi(x, \lambda)$. Here ρ and ϕ are the amplitude and phase of S as described in equation (2). In particular, we use the x-derivative of $\log_e S$, which satisfies:

$$\begin{aligned} \frac{\partial}{\partial x} \log_e S(x, \lambda) &= \frac{[S^*(x, \lambda)\, S_x(x, \lambda)]}{|S(x, \lambda)|^2} \\ &= \frac{\rho_x(x, \lambda)}{\rho(x, \lambda)} + i\phi_x(x, \lambda). \end{aligned} \tag{10}$$

Here the imaginary part is $\phi_x(x, \lambda)$, which is simply the local frequency as in (3), and the real part is the relative amplitude derivative $\rho_x(x, \lambda)/\rho(x, \lambda)$. The general behaviour of these two terms in the neighbourhood of a phase singularity is illustrated using the

simple model (8). In what follows, let (x_0, λ_0) denote the location of a typical singularity.

The neighbourhoods just above and below singular points can be characterized in terms of the behaviour of $\phi(x, \lambda)$ and $\phi_x(x, \lambda)$. Above singular points (for $\lambda > \lambda_0$), they are characterized by local frequencies that are significantly below the corresponding peak tuning frequencies $k(\lambda)$. Within these neighbourhoods there exist *retrograde regions* where local frequencies are negative, i.e. $\phi_x(x, \lambda) < 0$. In Figure 2 (bottom-left), where the signal normally winds counter-clockwise about the origin, the retrograde behaviour corresponds to the middle of the small loop where the signal winds clockwise with respect to the origin. Along the boundaries of retrograde regions the local frequency is zero. The significance of $\phi_x(x, \lambda)$ vanishing is that the level phase contours are horizontal, and not vertical as desired. Nearby this boundary, both inside and outside the retrograde regions, the level contours are generally far from vertical which, as discussed above, implies considerable phase instability. Below singular points (for $\lambda < \lambda_0$) the neighbourhoods are characterized by local frequencies of response that are significantly higher than the peak tuning frequencies. In addition, the local frequency changes rapidly as a function of spatial location. The extreme local frequencies above and below the singular point also suggest considerable difficulty for accurate signal interpolation in these regions.

To illustrate this behaviour with singularities that arise in practice, Figure 3 (left) shows a 1D slice of $\phi(x, \lambda)$ and $\phi_x(x, \lambda)$ from the scale-space of Figure 1 at a single scale. This slice of scale-space, at $\lambda = 20$, passes through three singularity neighbourhoods: two just above singularities (near locations 17 and 124), and one just below a singularity (near location 180). Notice the low (sometimes negative) and high local frequencies near the singularities. Figure 3 (right) shows the typical behaviour of level phase contours near a singularity, taken from the scale-space expansion of a scanline from a real image. The phase singularity is the point in the middle through which several of the phase contours pass. The small elliptical contour marks the retrograde boundary where $\phi_x(x, \lambda) = 0$. The instability above the singular point is clear from the nearly horizontal level phase contours. Directly below the singular point the high local frequencies are evident from the high density of phase contours.

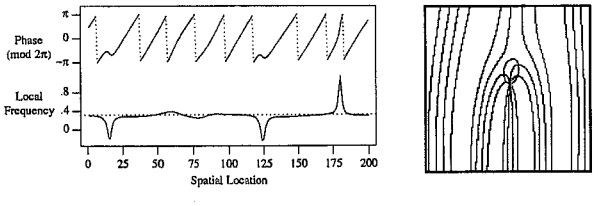

Figure 3. Phase and local frequency near singularities: (left) $\phi(x, \lambda)$ and $\phi_x(x, \lambda)$ are shown for a slice of the scale-space in Figure 1 (with $\lambda = 20$). Vertical dotted lines denote phase wrapping (not discontinuities), and the horizontal dotted line marks the filter's peak tuning frequency $k(\lambda) = 0.314$; (right) typical behaviour of level phase contours near a singularity. The singularity is the point in the centre through which several phase contours pass. The small ellipsoidal contour marks the retrograde boundary

Finally, the neighbourhoods to the left and right of (i.e. spatially adjacent to) singular points can be characterized in terms of amplitude variation. As we approach a singular point, $\rho(x, \lambda_0)$ goes to zero. Based on a simple linear model of amplitude, at a specific location x_1 near the singularity, the distance to the singular point is approximately $|\Delta x| = \rho(x_1, \lambda_0)/|\rho_x(x_1, \lambda_0)|$. Equivalently, as we approach the singularity, $|\rho_x(x_1|/\rho(x_1, \lambda_0)$ increases. As a consequence, the neighbourhoods to the left and right of phase singularities can be characterized by large values of $|\rho_x(x, \lambda)|/\rho(x, \lambda)$, which is provided by the real part of (10).

DETECTION OF SINGULARITY NEIGHBOURHOODS

In order to use phase information reliably toward the measurement of image velocity or binocular disparity, singularity neighbourhoods must be detected so that measurements in them may be discarded. Here we introduce constraints on local frequency and amplitude that can be used to identify locations within singularity neighbourhoods, while avoiding the explicit localization of the singular points. It is important in practice that this can be done using information available at only one scale, and hence a dense set of samples across many scales is not needed.

To detect the neighbourhoods above and below the singular points, we constrain the distance between the local frequency of response and the peak tuning frequency. This can be expressed as a function of the extent of the amplitude spectrum (measured at one standard deviation $\sigma_k(\lambda)$) *as follows:*

$$\frac{|\phi_x(x, \lambda) - k(\lambda)|}{\sigma_k(\lambda)} < \tau_k, \quad \sigma_k(\lambda) = k(\lambda)\frac{(2^\beta - 1)}{(2^\beta + 1)}. \quad (11)$$

The neighbourhoods adjacent to singular points can be detected with a local amplitude constraint:

$$\sigma(\lambda)\frac{|\rho_x(x, \lambda)|}{\rho(x, \lambda)} < \tau_\rho, \quad (12)$$

where $\sigma(\lambda)$ defines the radius of filter support. Level contours of (11) for different values of τ_k form 8-shaped regions with the singular points as their centres, while level contours of (12) form ∞-shaped regions. This is evident in Figure 4 (top panels). As τ_k and τ_ρ decrease, the constraints become tighter and larger neighbourhoods are detected. Figure 4 (middle-left) shows the combined behaviour of (11) and (12) as applied to the scale-space in Figure 1, with $\tau_k = 1.2$ (i.e. local frequencies are accepted up to 20% outside the nominal tuning range of the filters) and $\tau_\rho = 1.0$ (i.e. points within $\sigma(\lambda)$ of a singularity are discarded). These constraints typically remove about 15% the scale-space area. Finally, Figure 4 (middle-right) shows the original level phase contours as for Figure 1, while in the bottom row we present the contours that survive the constraints, and the contours in those regions removed. Notice the stability of the contours outside the singularity neighbourhoods.

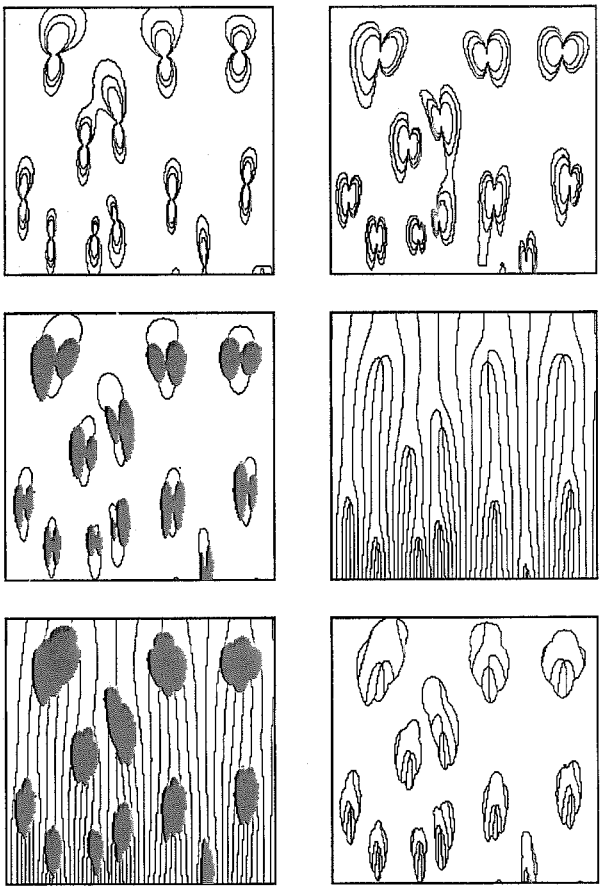

Figure 4. Detection of singularity neighbourhoods: (top-left) level contours of (11) for $\tau_k = 1$, 1.5 and 2 for the scale–space in Figure 1; (top-right) level contours of (12) for $\tau_\rho = 0.75$, 1, and 1.50; (middle-left) neighbourhoods removed by (12) with $\tau_\rho = 1$ are shown in black, while the contours show the remaining regions marked by (11) with $\tau_k = 1.2$; (middle-right) level phase contours for scale-space (cf. Figure 1); and (bottom row) contours that survive the constraints, and the phase contours in those regions removed with $\tau_\rho = 1$ and $\tau_k = 1.2$

MEASUREMENT OF BINOCULAR DISPARITY

To illustrate the problems caused by phase instability and the rapid variation of local frequency that occur in singularity neighbourhoods, we compare the results of a technique for disparity measurement with and without the removal of singularity neighbourhoods. Following Jenkin and Jepson[8] and Sanger[9], estimates of binocular disparity are computed as:

$$d(x) = \frac{[\phi_l(x) - \phi_r(x)]_{2\pi}}{k_0}, \qquad (13)$$

where $\phi_l(x)$ and $\phi_r(x)$ denote the phase responses of the left and right views, k_0 denotes the peak tuning frequency of the filter, and $[\theta]_{2\pi} \in (-\pi, \pi]$ denotes the principal part of θ. This computation presumes a model of local phase given by $\phi(x) = k_0 x + \phi_0$; when the left and right signals are shifted versions of one another, and the filter outputs have constant frequency k_0, then (13) yields the exact result. Toward a more general model, k_0 in (13) can be replaced by the average local frequency in the left and right outputs $(\phi'_l(x) + \phi'_r(x))/2$ (cf. Langley *et al.*[11]). This allows frequencies other than k_0 with accurate results[10]. Making the local model of the signal explicit is important because measurement accuracy and reliability depend on the appropriateness of the local model. For example, in neighbourhoods above and below singular points, which are characterized by a high variation in local frequency, the linear phase model is inappropriate and the numerical approximation of $\phi'(x)$ will be poor. The removal of these regions is therefore important.

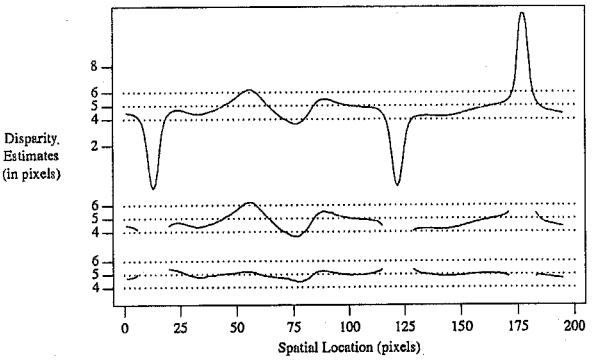

Figure 5. Disparity measurement: the top two plots show the disparity estimates based on (13) without, and then with, the removal of singularity neighbourhoods. Notice the substantial errors in the first case versus the second. The bottom plot shows the improved technique in which the local frequency is used instead of the peak frequency in (13). The same neighbourhoods have been removed

To illustrate this, assume a simple situation in which the left and right views are shifted versions of the 1D signal shown in Figure 1 (top). Let the disparity be 5 pixels, and let the Gabor filters be turned to a wavelength of 20 pixels. Thus the left and right phase signals are shifted versions of the scale-space slice shown in Figure 3 (left), which crosses three singularity neighbourhoods. Figure 5 (top) shows the results of (13) with the crude linear model[8,9] and without the removal of singularity neighbourhoods. Figure 5 (middle) shows the consequence of removing any disparity measurement for which the left or the right filter responses did not satisfy (11) or (12) with $\tau_\rho = 1.0$ and $\tau_k = 1.2$ (as in Figure 4). In Sanger[9] a heuristic constraint on amplitude differences between left and right signals and subsequent smoothing were used to lessen the effects of such errors. Unfortunately, this smoothing sacrifices the resolution and accuracy of nearby estimates. Finally, Figure 5 (bottom) shows the improvements obtained with the more general linear model.

SUMMARY

Phase-based techniques for the measurement of binocular disparity and image velocity are encouraging, especially because of the stability of band-pass phase information with respect to deviations from image translation that are typical in projections of 3D scenes.

Despite this stability, phase is unreliable in the neighbourhoods of phase singularities. This instability was described, and it was shown that singularity neighbourhoods may be detected using simple constraints on the local frequency and the amplitude of the filter output. Finally, these results were discussed briefly in the context of binocular disparity measurement.

ACKNOWLEDGEMENTS

This research has been supported by grants from NSERC Canada and ITRC.

REFERENCES

1 **Fennema, C L and Thompson, W** 'Velocity determination in scenes containing several moving objects', *Comput. Vision, Graph. & Image Process.* Vol 9 (1979) pp 301–315

2 **Horn, B K P and Schunck, B G** 'Determining optic flow', *Artif. Intell.* Vol 17 (1981) pp 185–204

3 **Grazer, F** 'Hierarchical gradient-based motion detection', *Proc. DARPA IUW,* LA, USA (1987) pp 733–748

4 **Nagel, H H** 'On the estimation of optical flow: Relations between different approaches and some new results', *Artif. Intell.* Vol 33 (1987) pp 299–324

5 **Nagel, H H and Enkelmann, W** 'An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences', *IEEE Trans PAMI* Vol 8 (1988) pp 565–593

6 **Waxman, A M, Wu, J and Bergholm, F** 'Convected activation profiles: Receptive fields for real-time measurement of short-range visual motion', *Proc. IEEE CVPR*, Ann Arbor, MI, USA (1988) pp 717–723

7 **Marr, D and Ullman, S** 'Directional selectivity and its use in early visual processing', *Proc. R. Soc. Lond. B* Vol 211 (1981) pp 151–180

8 **Jenkin, M and Jepson, A D** 'The measurement of binocular disparity', in **Z Pylyshyn (ed.)**, *Computational Processes in Human Vision*, Ablex Press, NJ, USA (1988) pp 69–98

9 **Sanger, T** 'Stereo disparity computation using Gabor filters', *Biological Cybernetics* Vol 59 (1988) pp 405–418

10 **Fleet, D J, Jepson, A D and Jenkin, M** 'Phase-based disparity measurement' *Comput. Vision, Graph. & Image Process.*, Vol 53 No 2 (1991) pp 198–210

11 **Langley, K, Atherton, T J, Wilson, R G and Larcombe, M H E** 'Vertical and horizontal disparities from phase', *Proc. 1st ECCV*, Antibes, France (1990) pp 315–325

12 **Fleet, D J and Jepson, A D** 'Computation of component image velocity from local phase information', *Int. J. Comput. Vision*, Vol 5 (1990) pp 77–105

13 **Verri, A and Poggio, T** 'Against quantitative optic flow', *Proc. IEEE ICCV*, London, UK (1987) pp 171–179

14 **Witkin, A** 'Scale-space filtering', *Proc. IJCAI*, Karlsruhe, Germany (1987) pp 1019–1022

15 **Schunck, B G** 'Image flow: fundamentals and future research', *Proc. IEEE CVPR*, San Francisco, CA, USA (1985) pp 560–571

16 **Mallat, S** 'A theory for multiresolution signal decomposition: The wavelet representation', *IEEE Trans. PAMI* Vol 11 (1989) pp 674–693

17 **Koenderink, J J** 'The structure of images', *Biological Cybernetics* Vol 50 (1984) pp 363–370

18 **Papoulis, A** *Probability, Random Variables, and Stochastic Process*, McGraw-Hill, Singapore (1965)

19 **Whitham, G B** *Linear and Nonlinear Waves*, John Wiley, NY, USA (1974)

20 **Gabor, D** 'Theory of communication', *J. IEE,* Vol 93 (1946) pp 429–457

21 **Fleet, D J** *Measurement of Image Velocity*, PhD Dissertation, Department of Computer Science, University of Toronto (1990)