Pergamon

PII: S0042-6989(97)00128-4

Vision Res., Vol. 38, No. 2, pp. 259-266, 1998

Printed in Great Britain
0042-6989/98 $19.00 + 0.00

Is Human Motion Detection Subserved by a Single or Multiple Channel Mechanism?

ROBERT F. HESS,*† PETER J. BEX,* ERIC R. FREDERICKSEN,* NUALA BRADY*

Received 24 September 1996; in revised form 20 December 1996; in final form 15 April 1997

Two recent versions of a single channel model of motion perception have had impressive success in explaining direction discrimination by human observers for spatially filtered noise images in two-flash apparent motion. It has been argued that the dramatic breakdown in motion perception which occurs when one image in the two-flash sequence is low-pass filtered can be explained only by a single channel model. We show that neither version of the single channel model which has been proposed can explain performance for noise images chosen to provide comparable stimulation in the spatial channels known to subserve human vision. A multi-channel model of motion perception has little difficulty in explaining these results. © 1998 Elsevier Science Ltd

Single channel Motion Fractal Features D_{max} Apparent motion

INTRODUCTION

Until recently, the accepted models of human motion processing typically involved multiple mechanisms operating over a range of spatial and temporal scales (for review see Nakayama, 1985). The most popular examples involved the detection of motion energy (i.e. oriented energy in the spatio-temporal spectrum) using receptive fields that are narrowly tuned for orientation and spatial frequency (Adelson & Bergen, 1985; van Santen & Sperling, 1985; Watson & Ahumada, 1985 but also see Eagle, 1996 for a multi-channel feature model). In the case of a broad band image in two-flash apparent motion, the direction of displacement would be detected independently at a number of different spatial scales. The maximum displacement for which direction can be reliably reported (D_{max}), involves the combination of directional signals at a variety of spatial scales (Cleary & Braddick, 1990a). For band-pass filtered noise patterns, D_{max} is inversely related to the central frequency of the filter (Chang & Julesz, 1983, 1985; Bischof & Di Lollo, 1990; Cleary & Braddick, 1990a; Cleary, 1990). This has been assumed to reflect the activity of multiple band-pass mechanisms underlying motion processing.

Two recent studies have provided impressive support for an alternative model of motion processing comprising a single spatial channel (termed the single channel pre-filter model). In single channel models, image features in the output of a single pre-filter are matched between successive presentations. D_{max} is determined by the largest displacement consistent with a criterion percentage of false matches. The scaling of D_{max} with centre frequency in band-pass filtered noise images is consistent with the fact that the average separation between image features scales with centre spatial frequency.

When multi-resolution images are low-pass filtered, D_{max} is initially unaffected. However, when the filter cut-off is below 3.5 c/deg, D_{max} linearly scales with the cut-off frequency (Cleary & Braddick, 1990b; Bischof & Di Lollo, 1990; Morgan & Mather, 1994). The more convincing explanation for this initial lack of dependence comes from the single channel pre-filter model (a single bandpass filter preceding nearest neighbour matching of like-signed zero crossings in the filter output). The initial lack of effect on D_{max} of low-pass filtering is accounted for by the *lower* cut-off of the single pre-filter and so its output is unaffected. The linear scaling of D_{max} at lower cut-offs is consistent with the increasing distance between features in the output of the pre-filter (Morgan, 1992; Morgan & Mather, 1994). Multi-channel models resort to special pleading to explain this result by invoking inhibitory interactions among channels tuned to different spatial scales (Cleary & Braddick, 1990a).

The second and more critical piece of evidence in support of the single channel pre-filter model involves noise images in two-flash apparent motion, where only one of the two image sequences is low-pass filtered. Morgan & Mather (1994) showed that human direction discrimination collapsed for such stimuli when the filter cut-off of one image was below 1 c/deg. They showed that such a result is expected from their single channel model because there is little correlation between the two images in the output of the single pre-filter. However, the result was not predicted by multi-resolution motion

*McGill Vision Research, Department of Ophthalmology, McGill University, Montreal, Quebec, Canada.
†To whom all correspondence should be addressed [*Fax*: +1 514 843 1691; *Email*: Rhess@bradman.vsion.mcgill.ca].

models since there is always energy at common spatial scales (i.e. at low spatial scales) in the two images. Recently, Ledgeway (1996) also studied motion detection in two-flash random dot kinematograms in which the spatial structure of the noise in one or both frames was independently manipulated (band-pass filtered, in this case). Ledgeway found that there had to be some overlap in the spatial structure of both frames for motion to be detected. However, the result may be consistent with either single or multiple channel models: in one case (multiple channel) motion may be supported by mechanisms selective for the overlapping frequencies; in the other case (single channel pre-filter) motion may be supported by correlations in image structure arising only when the Fourier amplitude and phase spectra of the two patterns are similar.

To date, broad band images with effectively flat power spectral density in linear coordinates have been used in two-flash motion experiments. It has been assumed that because of its flat spectrum, such a stimulus provides equal activation of the spatial mechanisms underlying human motion perception. However, the bandwidth of human spatial mechanisms tuned to different peak spatial frequencies is known to be approximately equal when plotted in *log* not *linear* frequency (Field, 1987). This assumes that different spatial mechanisms have equal contrast gain (see Brady & Field, 1995 for psychophysical evidence; and Croner & Kaplan, 1995 for neurophysiological evidence). This observation of the bandwidth of different spatial mechanisms has an important bearing on the above debate between multi-channel and single channel models of motion perception. The traditional use of broad band images with flat (linear) spectra (i.e., white noise) in two-flash apparent motion is far from ideal because it results in a greater activation of high spatial frequency mechanisms known to operate within the human visual system. A stimulus which provides more equal activation of human spatial mechanisms would have an amplitude spectrum which varies as 1/f (Field, 1987). This consideration is paramount when considering D_{max} since, according to multi-channel models, D_{max} depends on the combination of directional signals at a variety of spatial scales.

A stimulus with a 1/f amplitude spectrum (termed fractal) can provide clear predictions with which to differentiate between multi- and single channel models of motion detection (see Fig. 1). The overall directional response of multiple channel models depends on the rules for the combination of motion signals across spatial scales. Cleary & Braddick (1990a) suggest that direction discrimination is supported by the lowest spatial frequency at which there is a detectable, coherent motion signal. This is also supported by the results of Boulton & Hess (1990), Boulton & Baker (1991) and Eagle (1996). For 1/f images, this component is equally detectable in unfiltered and low-pass filtered 1/f images, so D_{max} should be relatively stable. However, the lowest detectable coherent motion signal increases with the cut-off of a high-pass filter and so D_{max} should linearly decrease with

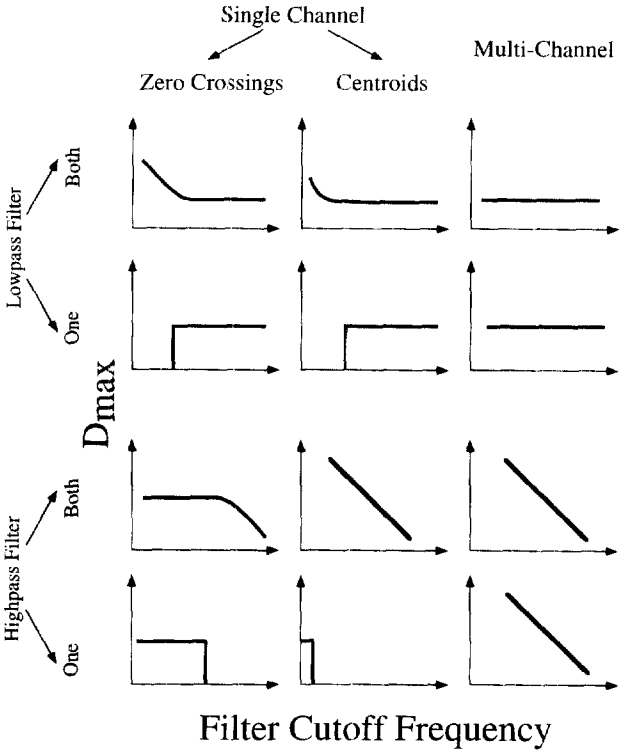

FIGURE 1. Schematic predictions for the relationship between D_{max} for direction discrimination and different types of spatial filtering of 1-D motion sequences seen in two-flash apparent motion for single (zero crossings and centroids) and multi-channels of motion detection. In the multi-channel model it is assumed that information at the lowest available spatial scale predominates.

the removal of low spatial frequency components (in accordance with the half-cycle limit).

For the single channel case, the combination of the profile of the pre-filter (band-pass or low-pass) and the matching token (zero crossings or centroids of zero bounded regions) determine the form of these relationships. For example, when zero crossings are matched after band-pass front end filtering, at some point D_{max} should scale with the low-pass filter cut-off because the average distance between features will be varying. In the high-pass case, D_{max} should be relatively independent of the filter cut-off because the average distance between features is not varying. But when the centroid of zero bounded regions are matched after low-pass front end filtering, there are opposing sets of predictions for D_{max}. In the low-pass case, D_{max} should be relatively invariant with filter cut-off and in the high-pass case, D_{max} should scale with the filter cut-off. This follows from the relative importance of low spatial frequencies to determining centroid locations in the low-pass pre-filter output.

When only one of the two 1/f images in two-flash apparent motion is filtered, the multi-channel predictions remain unchanged. This is because there is no coherent motion signal at the spatial scales present in only one image, but there is still energy or features at spatial scales common to both images. However, the predictions change for the single channel pre-filter models. Motion

perception should collapse when there is little or no correlation between the zero crossing or centroid locations in the output of the pre-filter stage for each image. In general, this will occur at the extremes of spatial frequency filtering, although, as we will see, the centroid model is particularly sensitive to even slight degrees of high-pass filtering of one of the two images. These predictions represent a strong test of the single channel model.

METHODS

Stimuli

1/f noise images (1-D) were generated on a PC (Gateway 20004DX2-66E) using custom software. The images were generated by Fourier transforming Gaussian distributed 8-bit noise images using an FFT (Press *et al.*, 1992), filtering their amplitude spectra to 1/f, and inverse transforming. Low-pass and high-pass filtered images were created by setting to zero the amplitude of all frequencies beyond the cut-off frequency of the filter before taking the inverse transform. With an 8 deg square display of 256×256 elements, the images spanned seven frequency octaves from 0.125 c/deg to a Nyquist frequency of 16 c/deg. The upper and lower cut-off of the high and low-pass filters were separated by one octave steps from 0.125 to 8 c/deg. The 1/f noise was generated at maximum contrast (r.m.s. contrast = 31%), the contrast of the filtered images was not normalized after filtering. This ensured that the contrast at all spatial scales was equal in filtered and broad band image pairs. The image contrast in a filtered image was proportional to the octave bandwidth of that image compared with the original unfiltered version. For example, an image filtered to a one octave band would have an r.m.s. contrast of 5.6%.

For each experimental run, the images were pre-computed in strips which were 2048 pixels wide by 256 pixels high and loaded to the graphics card, a Visual Stimulus Generator 2/3 (Cambridge Research Systems). Six novel image strips were used when both images in the motion sequence had the same spectral characteristics (e.g., high-pass – high-pass). Three novel image pairs were used when the spectral characteristics of the images differed (e.g., 1/f – low pass). On each trial, from a randomly selected image strip, a stimulus image measuring 256 by 256 pixels was presented to the display (using a randomly determined horizontal starting co-ordinate and full wrap-around). The displaced image was then chosen from either the same image strip (same spectrum condition) or from the matching image strip (different spectrum condition), at a starting horizontal coordinate offset from that of the first image by the displacement size. This method ensured that novel images were used throughout the experiment.

Display

The stimuli were displayed on a Nanao Flexscan 6500 monitor with a frame rate of 100 Hz. The mean luminance of the display was 26 cd/m^2. The luminance of the display was calibrated using a UDT Photometer and linearized using a lookup table method. Subjects viewed the screen binocularly in a dimly lit room. The stimulus images were 256 by 256 pixels and subtended 8 deg of visual angle at a viewing distance of 80 cm. They were presented in the centre of the screen and the remainder of the display was blank and at the mean luminance.

Procedure

One of the authors (PB) and a naïve subject (IM) served as observers, both had normal or corrected to normal visual acuity and practised the task before formal data collection. Each trial was preceded by a blank screen of mean luminance with a small, bright fixation square in the centre. This remained on until the trial was initiated by the observer. An image was then presented for 100 msec., and the displaced image was presented for a further 100 msec without an inter-stimulus interval. The direction of the displacement (right or left) was randomly determined on each trial, and the observer indicated the perceived direction of motion by pressing one of two keyboard buttons. The displacement size was determined using the method of constant stimuli. The magnitude and direction of displacement were varied from small displacements at which there were few or no errors in direction discrimination to large displacements at which performance fell to chance. The errors in direction discrimination were recorded at each displacement, and were later fitted with a Weibull function (Weibull, 1951) from which D_{max} was estimated as the displacement at which there were 20% errors. Each displacement size was presented at least 40 times.

Modelling

Responses of the two different versions of the single channel pre-filter model were simulated using custom functions written in MatLab (Mathworks, Inc.) software. Simulations were spatially one-dimensional (1-D) (as were the stimuli used in the psychophysics) with stimuli presented as space–time arrays with spatial positions along a fractal noise pattern. The 1-D fractal noise stimuli used in the model were constructed in an identical manner to those used in the experiments but along a single spatial dimension. Each model operates on a space–time array which specifies the two-flash sequence, and produced a scalar signal of motion strength. This signal was linked to a "left" versus "right" decision according to its sign and logged as a correct or erroneous trial according to whether it agreed with the actual direction of simulated motion. The simulation looped through 200 sets of 15 displacements and calculated the accumulated percent errors, as in a typical psychophysical experiment. The upper displacement threshold (D_{max}) was estimated at the displacement at which there were 20% errors. This was repeated for all the different conditions of spatial filtering. Either both fractal images were spatially filtering with the same cut-off (low-pass or

high-pass) or only one of the two images was spatially filtered (low-pass or high-pass). For all the model predictions there was only one free parameter, the vertical scaling.

Two different versions of the single channel feature matching model were examined. In the first, zero-crossings were used as the matching token in the output of a Laplacian of gaussian pre-filter with a space constant of 10 min. (Morgan, 1992). Like signed zero-crossings were matched in the 1-D profiles by searching for instances in the profile having successive points one below $-\theta$ and an adjacent one above $+\theta$, the threshold offset. Pilot studies showed that as long as the threshold offset θ was small relative to the amplitude of the stimuli (i.e., 0.01) the simulation results were insensitive to the value of this threshold. In the second version of the pre-filter model (Morgan & Mather, 1994), centroids of like signed zero bounded regions in the output of a single low-pass pre-filter (gaussian with a sigma of 6.75 min) were used as matching tokens. For each token in the first profile the nearest neighbour to the left and right of it in the second profile was determined and the nearer of the two taken as the correspondence match. This match contributed +1 or −1 to a tally, according to whether it corresponded to a left or right displacement. The sign of the final tally indicated whether the net response was leftward or rightward. If the final tally was equal to zero, the direction was randomly chosen. This was repeated for each of 15 displacements and the whole process replicated 200 times to derive a psychometric function for the model. A multi-channel LoG prediction was generated for the case where the image sequences were differentially filtered. The peak position of the LoG filter was either set at 1 cycle/image (low-pass filtering case) or to coincide with the position of the cut-off of the ideal filter (high-pass filtering case). The matching of nearest neighbour centroids determined D_{max}, in exactly the same way as described above for the single channel model.

RESULTS

Figure 2 shows the results (filled symbols) for low-pass filtering of both of the fractal images in a two-flash apparent motion sequence in which subjects had to discriminate the direction of motion. It can be seen that D_{max} is relatively independent of the cut-off frequency of the filter down to around 0.5 c/deg, which represents 4 cycles/image, although for observer IM, there is an initial decrease. At lower filter cut-offs there was a modest increase in D_{max}. Model predictions are shown for both versions of the single channel pre-filter model (Morgan, 1992; Morgan & Mather, 1994). The zero crossing model (following Laplacian filtering) is shown by the dashed line and the centroid model (following gaussian filtering) is shown by the solid line. It can be seen that while both models capture the initial independence of D_{max} with filter cut-off, they both display strong dependence on filter cut-off when it goes below 1 c/deg. This departure is more obvious for the zero crossing version of the model. The stability of the centroid model is a result of the stable low spatial frequency contribution to the centroid locations derived from the output of the single low-pass filter.

In Fig. 3, results of a similar experiment are displayed but this time for high-pass filtered image sequences. When both images were high-pass filtered to the same cut-off (unfilled symbols), D_{max} decreased steadily with the cut-off of high-pass filter. Model predictions are shown for both versions of the single channel pre-filter model (Morgan, 1992; Morgan & Mather, 1994). The

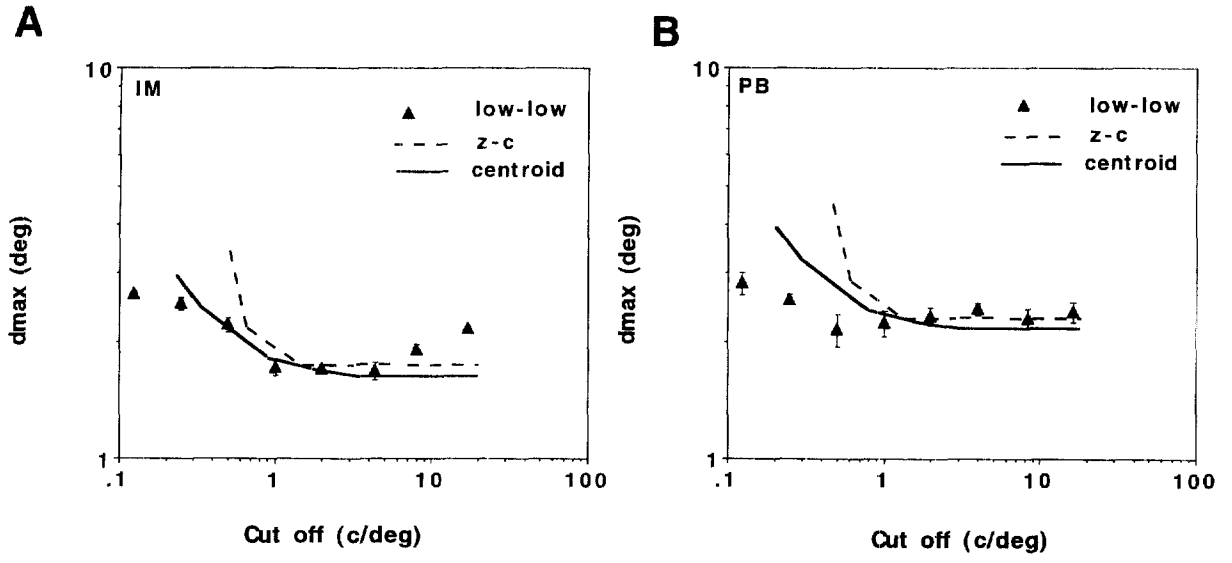

FIGURE 2. Direction discrimination data (symbols) for fractal images seen in two-flash apparent motion when both images are subjected to low-pass filtering. D_{max} in degrees is plotted against the cut-off of the ideal low-pass filter. The results are compared with two versions of the single channel pre-filter model; nearest neighbour matching of zero-crossings (dashed) and centroids of zero-bounded distributions (solid).

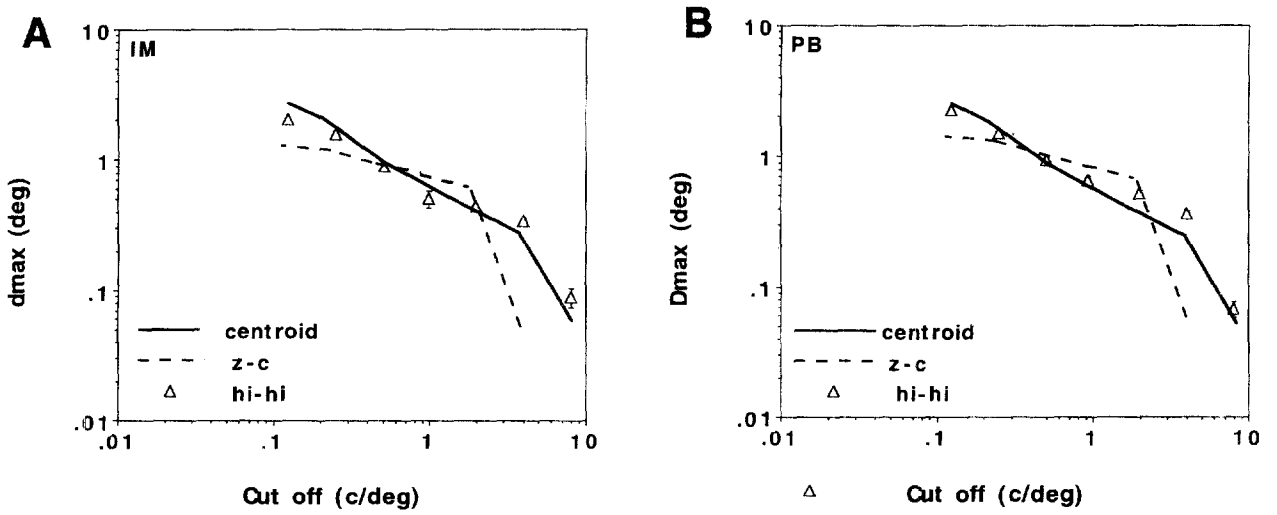

FIGURE 3. Direction discrimination data (symbols) for fractal images seen in two-flash apparent motion when both images are subjected to high-pass filtering. D_{max} in degrees is plotted against the cut-off of the ideal high-pass filter. The results are compared with two versions of the single channel pre-filter model; nearest neighbour matching of zero-crossings (dashed) and centroids of the zero-bounded distributions (solid).

zero crossing following Laplacian filtering model is shown by the dashed line and the centroid model following gaussian filtering is shown by the solid line. It can be seen that the centroid model was best able to explain these particular psychophysical results. The zero crossing model predicts a value of D_{max} that is initially relatively independent of filter cut-off and later strongly dependent on it. The dependence of D_{max} on high-pass filtering for the centroid model is a result of the importance of the low spatial frequency contribution to centroid locations.

In Figs 4 and 5 we show psychophysical data and model predictions for the conditions where only one image in the two-flash sequence was filtered. Figure 4 shows the low-pass data and Fig. 5 the high-pass data. In each figure, data are shown for both when the filtered image was presented first (circles) or second (squares) in the two-flash sequence. There was no effect of presentation order on performance for either type of filtering. Comparison of Figs 3 and 5 reveals that D_{max} was approximately equal when either both or a single image was high-pass filtered. However, D_{max} was slightly lower when only one image of the pair was low-pass filtered (Figs 2 and 4). Model predictions are shown as dashed lines (zero crossings following Laplacian filtering) and solid lines (centroids following gaussian filtering). Neither model was able to indicate the correct direction of displacement for the entire range over which human observers could perform this task (filter cut-off from 1 c/image or 0.125 c/deg to 12 c/image or 16 c/deg). The

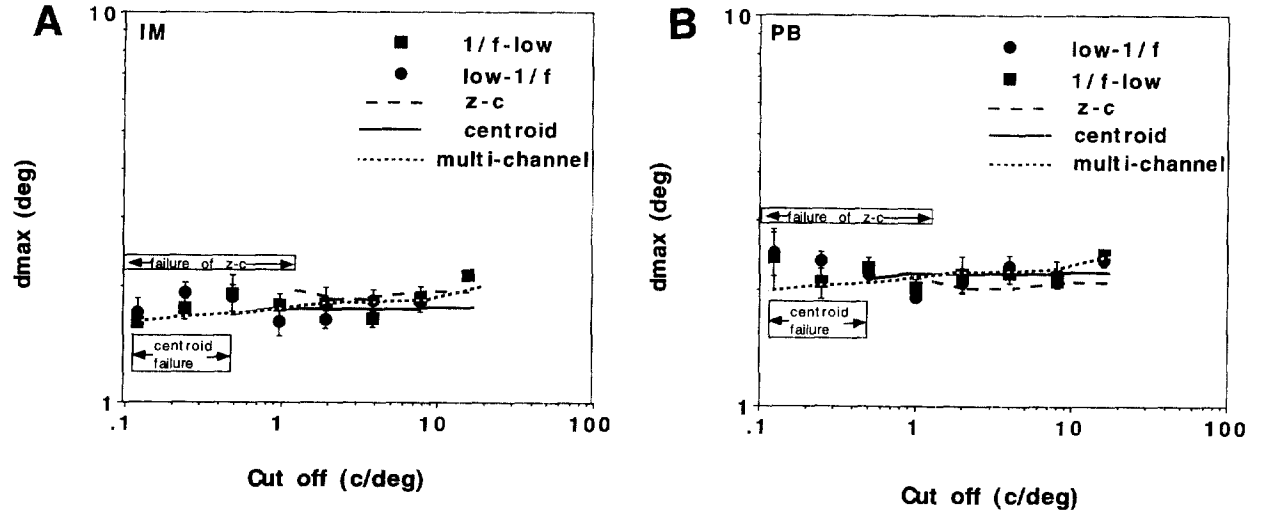

FIGURE 4. Direction discrimination data (symbols) for fractal images seen in two-flash apparent motion when only one image is subjected to low-pass filtering. D_{max} in degrees is plotted against the cut-off of the ideal low-pass filter. The results are compared with two versions of the single channel pre-filter model: nearest neighbour matching of zero-crossings (dashed) and centroids of zero-bounded distributions (solid) and one version of a multi-channel model (dotted lines).

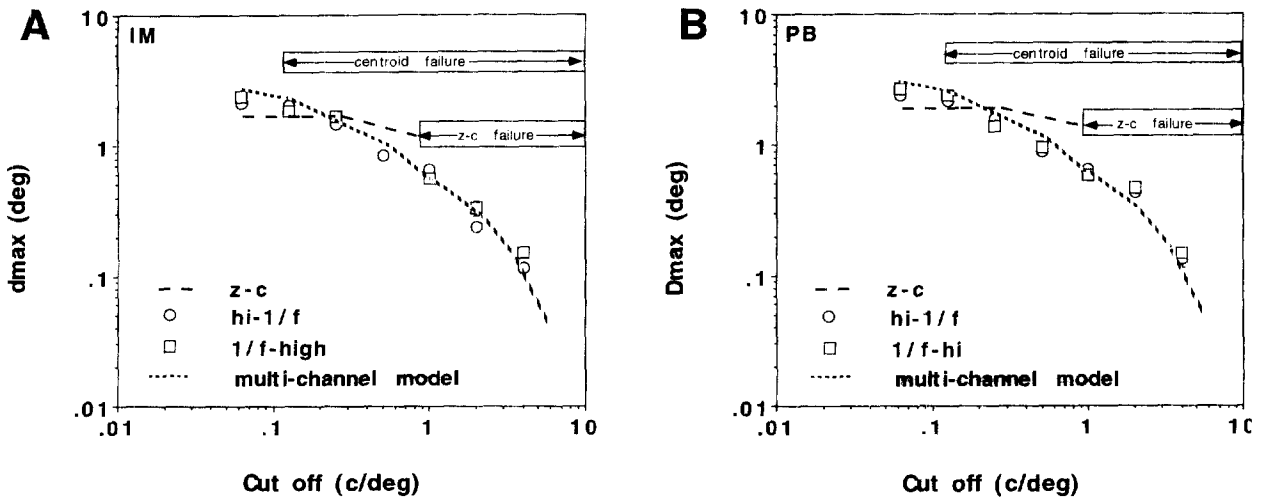

FIGURE 5. Direction discrimination data (symbols) for fractal images seen in two-flash apparent motion when only one image is subjected to high-pass filtering. D_{max} in degrees is plotted against the cut-off of the ideal high-pass filter. The results are compared with two versions of the single channel pre-filter model; nearest neighbour matching of zero-crossings (dashed) and centroids of zero-bounded distributions (threshold predictions could not be obtained over the region demarked by arrows) and one version of a multi-channel model (dotted lines).

zero crossing model failed when either type of filtering was severe (below 1 c/deg in the low-pass case and above 1 c/deg in the high-pass case). The centroid model which failed below 0.5 c/deg in the low-pass case was unable to cope with even a slight degree (i.e. above 1 c/image or 0.125 c/deg) of high-pass filtering.

In order to show the sufficiency of a multi-channel model we modelled the psychophysics for which the single channel model had most clearly failed to explain the measured human performance, namely when one image of the two-flash sequence is differentially filtered. We show results for the centroid feature matching (see Eagle, 1996 for a multi-channel zero crossing model). We assume that D_{max} is determined at the lowest spatial scale at which there is coherent motion (Cleary & Braddick, 1990a; Boulton & Hess, 1990; Boulton & Baker, 1991; Eagle, 1996). For the low-pass filtering case, we set the LoG filter to have a peak at 1 c/image. For the high-pass filtering case, we set the filter to coincide with the cut-off of the ideal filter used to high-pass one of the images. We show in Fig. 6 that filters positioned much lower than this will fail to see coherent motion because they will be matching features of different scale and filters much higher than this will not only have a lower D_{max} but also will signal motion in the opposite direction at moderate jump sizes (Eagle, 1996).

The multi-channel model predictions are shown as dotted lines in Figs 4 and 5. In the low-pass case, this filter is fixed at the lowest scale and motion is supported (dotted prediction) across the same range as that found for human observers. In the high-pass case, the optimal filter depends on the extent of high-pass filtering and it is seen to scale (dotted prediction) with the extent of high-pass filtering in good agreement with the human psychophysics. Note that neither of the two versions of the single channel can account for human performance over most of this range (dashed and solid predictions in Figs 4 and 5).

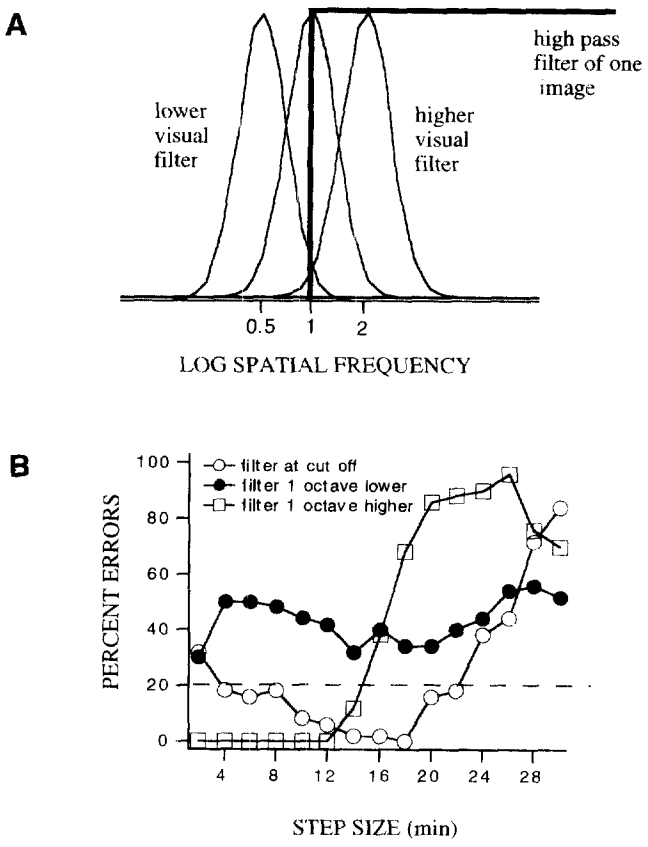

FIGURE 6. For the case where only one of the two image sequences is high-pass filtered. Model responses (B) are shown for a band-pass filter positioned either 1 octave below or 1 octave above the high-pass cut-off (A). It is argued that to optimize D_{max} it is best to use a filter whose peak is located near the high-pass cut-off. In implementing the multi-channel model, the response of the filter located at the high-pass cut-off of the image filter was used, the choice was, however, arbitrary.

DISCUSSION

The results (Figs 2 and 3) show that for both observers, D_{max} was relatively stable after low-pass filtering, but declined steadily with high-pass filtering when both images in a two-flash apparent motion sequence were equally filtered. Both versions of the single channel pre-filter model captured the general form of the psychophysical data.

When only one image of the two-flash apparent motion sequence was filtered, observers were able to detect motion between the two images despite their differing appearance. The psychophysical results were similar to the conditions in which both images were filtered. Neither the single channel zero-crossing nor the centroid model could provide a satisfactory prediction of the performance of human observers under these conditions. The similarity between the low-pass and high-pass results for the same and mixed spectrum sequences is in very good agreement with the predictions outlined for a multiple channel motion detection model (see Fig. 1) and shown to be supported by the multi-channel modelling using centroid features (dotted predictions in Figs 4 and 5). However, the results show that the cornerstone prediction of the single channel pre-filter model, namely the failure of motion detection when one of the two images in the two-flash sequence is low-pass filtered (Morgan & Mather, 1994), is not satisfied for human motion detection using fractal images.

Morgan & Mather (1994) found that when only one of the images was low-pass filtered, direction discrimination was sometimes worse when the unfiltered image was presented first. They suggested that the effect could arise from temporal processing differences for high and low spatial frequencies. Our results show that direction discrimination did not depend on the order of presentation for filtered and unfiltered 1/f noise images. Ledgeway (1996) also found no effect of temporal order for motion detection between filtered and unfiltered images. Ledgeway's images were flat spectrum white noise patterns (as were Morgan and Mather's) and so the difference may not be attributed to the fractal scaling of our patterns. The reason for these differences remains unclear.

Our psychophysical results differ from those of Morgan & Mather (1994) because we use multi-resolution *fractal* images which produce comparable stimulation of the underlying spatial mechanisms. The use of multi-resolution images which provide comparable stimulation across the range of different spatial mechanisms known to underlie human visual processing greatly clarifies the relationship between D_{max} and spatial filtering (see also Eagle, 1996). If we assume that D_{max} depends on the half-cycle limit of the lowest detectable spatial frequency in the image (Boulton & Hess, 1990; Boulton & Baker, 1991; Eagle, 1996) then it is not surprising that for fractal images D_{max} is relatively invariant with low-pass filtering and it directly depends on the cut-off of high-pass filtering, whether one or both images are filtered. However, the psychophysical data show that D_{max} was much lower than one half-cycle of the lowest spatial frequency in the image. The lowest spatial frequency in our patterns was 0.125 c/deg, with a half-cycle limit of 4 deg, the estimates of D_{max} were around 2 deg. Furthermore, even using low-pass filtered 1/f images, there was a slight increase in D_{max} when the cut-off of the low-pass filter was below about 0.5–1 c/deg, especially for the naïve observer [IM, Fig. 2(a)]. The increase is much smaller and occurs at a lower cut-off than that observed by Cleary & Braddick (1990b) (at 3–4 c/deg) using white noise patterns.

These results suggest that either 1/f noise images do not provide *exactly* equal stimulation to all spatial channels or Cleary and Braddick's proposal that high spatial frequencies mask low may still be necessary even when the relative sensitivity of the human visual system is considered. Note that when the high spatial frequencies were present on a single frame (Fig. 4) and were therefore not moving, the increase in D_{max} was lost. This is also consistent with high spatial frequency masking if it is mediated by a temporal transient at onset and offset of the unfiltered image.

REFERENCES

Adelson, E. H. & Bergen, J. R. (1985). Spatio-temporal energy models for the perception of motion. *Journal of the Optical Society of America A, 2*, 284–299.

Bischof, W. F. & Di Lollo, V. (1990). Perception of directional sampled motion in relation to displacement and spatial frequency: evidence for a unitary motion system. *Vision Research, 9*, 1341–1362.

Boulton, J. C. & Baker, C. L. (1991). Motion detection is dependent on spatial frequency not size. *Vision Research, 31*, 77–87.

Boulton, J. C. & Hess, R. (1990). The optimal displacement for the detection of motion. *Vision Research, 32*, 61–72.

Brady, N. & Field, D. G. (1995). What's constant in contrast constancy? The effects of scaling on the perceived contrast of bandpass patterns. *Vision Research, 35*, 739–756.

Chang, J. J. & Julesz, B. (1983). Displacement limits, directional anisotropy and direction versus form discrimination in random-dot cinematograms. *Vision Research, 23*, 639–646.

Chang, J. J. & Julesz, B. (1985). Cooperative and non-cooperative processes of apparent motion of random-dot cinematograms. *Spatial vision, 1*, 39–45.

Cleary, R. (1990). Contrast dependence of short range apparent motion. *Vision Research, 30*, 463–478.

Cleary, R. & Braddick, O. J. (1990a). Discretion discrimination for bandpass random-dot kinematograms. *Vision Research, 23*, 303–316.

Cleary, R. & Braddick, O. J. (1990b). Masking of low frequency information in short-range apparent motion. *Vision Research, 30*, 317–327.

Croner, L. J. & Kaplan, E. (1995). Receptive fields of P and M ganglion cells across primate retina. *Vision Research, 35*, 7–24.

Eagle, R. A. (1996). What determines the maximum displacement limit for spatially broadband kinematograms. *Journal of the Optical Society of America A, 13*, 408–418.

Field, D. J. (1987). Relation between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A, 4*, 2379–2394.

Ledgeway, T. (1996). How similar must the Fourier spectra of the frames of a random-dot kinematogram be to support motion perception? *Vision Research, 36*, 2489–2495.

Morgan, M. J. (1992). Spatial filtering precedes motion detection. *Nature (Lond.), 365*, 344–346.

Morgan, M. J. & Mather, G. (1994). Motion discrimination in two-

frame sequences with differing spatial frequency content. *Vision Research, 34*, 197–208.

Nakayama, K. (1985). Biological image motion processing: a review. *Vision Research, 25*, 625–660.

Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. (1992) *Numerical recipes in C*. (2nd edn). Cambridge, U.K.: Cambridge University Press.

van Santen, J. P. H. & Sperling, G. (1985). Elaborated Reichardt detectors. *Journal of the Optical Society of America A, 2*, 300–321.

Watson, A. B. & Ahumada, A. J. (1985). Model of human visual-motion sensing. *Journal of the Optical Society of America A, 2*, 322–341.

Weibull, W. (1951) A statistical distribution function of wide applicability. *Journal of Applied Mechanics, 18*, 292–297.

Acknowledgements—This work was supported by an NSERC grant to RFH (#OGP0046528) and CLB (#OGP0001978). We are grateful to Curtis Baker for discussions and for use of his modelling software.