



Effects of Dot Density, Patch Size and Contrast on the Upper Spatial Limit for Direction Discrimination in Random-dot Kinematograms

RICHARD A. EAGLE,*† BRIAN J. ROGERS*

Received 3 October 1995; in revised form 26 April 1996

Two-frame random-dot kinematograms (RDKs) of different dot density, area and contrast were used to study the spatial properties of the human visual motion system. It was found that the maximum spatial displacement at which observers could reliably discriminate the direction of motion (d_{\max}) increased gradually by a factor of up to 6.4 as dot density was decreased from 50 to 0.025% for high Michelson contrast (0.997) stimuli. As stimulus area was reduced from 645 deg², this trend gradually disappeared so that by a stimulus area of 2.56 deg², there was no effect of density upon d_{\max} . A further experiment investigated the effects of reducing Michelson contrast from 0.77 to 0.2 on d_{\max} over this same range of dot densities. It was found that at the highest densities, d_{\max} declined as contrast was reduced. Furthermore, for contrasts at and below 0.4, d_{\max} was invariant of density over the range 50–5%. These results can be accounted for by the fact that both reducing contrast, while keeping density fixed, and reducing density, while maintaining a fixed high contrast, reduce the stimulus mean luminance. For all contrasts, decreasing density below 5% led to an increase in d_{\max} . However, the rate of this increase was slower for the lower contrast stimuli. A two-stage model based on bandpass filtering followed by an informationally limited motion detection stage is proposed and shown to provide a good account of these data. © 1997 Elsevier Science Ltd.

Motion detection Dot density d_{\max} Random-dot kinematograms

INTRODUCTION

Random-dot kinematograms provide a useful tool for psychophysicists to investigate visual motion detection (Anstis, 1970; Julesz, 1971). In its simplest form, two identical but spatially displaced versions of the same random pattern of dark and light dots are exposed sequentially to an observer. Both frames are viewed through a stationary window so that the boundaries of the pattern are not visible. Braddick (1974) originally noted that when the spatial displacement between two frames of a random-dot kinematogram (RDK) was small (less than 15 arc min) subjects perceived coherent motion. However, for larger displacements, the percept was one of incoherent local motions—which he attributed to the spatial limits of the motion detecting mechanisms.

This maximum detectable displacement, or d_{\max} , has become a cornerstone, not only for researchers attempting to determine the nature of motion detecting processes, but also for a theoretical debate on the number of motion detection systems. In contrast to this 15 arc min

displacement limit, classical studies of apparent motion have shown that an impression of motion can still be gained when the displacement of a single element is as large as 18 arc deg (e.g. Zeeman & Roelofs, 1935). There are at least two possible explanations of the discrepancy between these findings. First, d_{\max} may reflect the *spatial* limit of a qualitatively distinct motion process, perhaps operationalized by the spatial complexity of a RDK. Alternatively, because RDKs present a far more difficult correspondence problem, it may be that a single motion system operates on all displays, in which d_{\max} is not constant in spatial units but varies with the number of “false targets” in the stimulus. In this case, different d_{\max} values for different stimuli are due to a constant *informational* limit of the system (see Fig. 1). Braddick (1974) sought to test these two hypotheses by varying the dot size of the elements—and so the spacing of false targets—in his RDK. The idea was that a constant value of d_{\max} under this manipulation would be evidence for an absolute spatial limit of a distinct process, whereas an increase in d_{\max} proportional to dot size would imply a single process with an informational limit. The actual finding was that d_{\max} remained constant for different dot sizes, which led Braddick to postulate a dichotomy of processes: (i) a “short-range” process that can detect the

*Department of Experimental Psychology, University of Oxford, South Parks Rd, Oxford OX1 3UD, U.K.

†To whom all correspondence should be addressed [Email: richard@psy.ox.ac.uk].




FIGURE 1. Illustration of how varying dot density leads to different predictions for a spatially limited and an informationally limited motion detector. (a) A spatially limited mechanism has a fixed spatial limit (the radius of the circle) regardless of the dot density. (b) An informationally limited mechanism is constrained only by the spacing of false targets, and so its spatial limit varies to keep the number of false targets within its search band here, the circle) constant.

motion of spatially complex patterns, such as RDKs, but with an upper spatial limit of approx. one-quarter of a degree; and (ii) a “long-range” process, that can only detect the motion of simpler patterns, such as single element stimuli, but which can operate over much greater spatial extents. While the short-range process was tentatively linked to the properties of directionally selective cortical cells, the long-range process was considered to be more central or cognitive (see also Braddick, 1980; Anstis, 1980; Petersik, 1989).

Subsequent studies by Lappin & Bell (1976), Baker & Braddick (1982), Chang & Julesz (1983) and Nakayama & Silverman (1984) have all shown that d_{\max} increases with increasing patch size. Baker & Braddick (1982) found that displaying a small window of motion at an eccentric location yielded similar direction discrimination error rates to a much larger window displayed at zero eccentricity. In contrast, a pattern of identical size yielded a significantly smaller d_{\max} in the fovea than in the periphery. Baker and Braddick's interpretation of these results was that d_{\max} is determined by a fixed spatial-limit mechanism, although the exact value varies (systematically) with eccentricity—an elaborated version of Braddick's (1974) original hypothesis.

The predictions of the spatial limit and informational limit hypotheses can also be tested by examining whether d_{\max} is affected by the dot density of a RDK. If d_{\max} is determined by dot spacing then it should increase as density is decreased, whereas if there is a fixed spatial limit then d_{\max} should not vary with density (see Fig. 1). Baker & Braddick (1982) measured d_{\max} in an experiment in which they varied the dot density over a range of 50% down to 1%, with a fixed patch size. They found that this manipulation had a negligible effect on d_{\max} and regarded this result as “...the most important argument against displacement being limited by the number of pixels shifted” (p. 1258).

In even more recent studies of this issue, Cavanagh *et*

al. (1985), Sato (1990) and, most thoroughly, Morgan (1992) have each manipulated dot size over a greater range than in Braddick's original experiment. Their collective results show that variations in dot size up to around 10–15 min have little effect on the magnitude of d_{\max} , but that increases in dot size beyond 15 min lead to proportional increases in d_{\max} . While these findings show that d_{\max} is not a constant spatial magnitude, it must be noted that because the largest dot size used by Braddick (1974) was 10.8 arc min, his data are perfectly compatible with these findings. In explaining their data, these investigators have all suggested that the visual system lowpass filters the stimulus *prior* to motion detection. One consequence of lowpass filtering a dense RDK is that the spacing of “blobs” in the filtered image is set by the dot size, for dots larger than the filter space constant, but by the filter size, for dots smaller than the filter space constant. Given this, it can be seen that the finding that d_{\max} is constant for small dot sizes is not inconsistent with either a spatially or an informationally limited process of motion detection. Similarly, the differences between a high and a low-density pattern will not be registered in the filter response if the mean dot spacing of the lower density pattern is less than the filter space constant. Thus, the fact that d_{\max} is constant for dot densities in the range 50–1% can also be predicted from coarse-scale filtering, and so does not, after all, distinguish between the two hypotheses.

While the finding that d_{\max} can vary with dot size is inconsistent with the idea that the short-range process has an absolute spatial limit, the hypothesis can be modified successfully by taking into consideration the spatial-frequency tuned channels that exist in early vision (e.g. De Valois & De Valois, 1988). Chang & Julesz (1985) and subsequently De Bruyn & Orban (1989), Cleary & Braddick (1990) and Bischof & Di Lollo (1990, 1991) have all measured d_{\max} for band-pass filtered RDKs i.e., patterns that contain only a narrow band of spatial frequencies. The consistent finding is that d_{\max} increases for the lower-frequency patterns, in inverse proportion to the centre frequency. The explanation offered by all these investigators was that the differently filtered stimuli are processed by separate channels, each with its own spatial limit set according to a fixed proportion of the period which it is maximally tuned to. This “phase-limit” hypothesis can account for the dependency of d_{\max} on dot size, as increasing dot size increases the energy at low frequencies while reducing it in the higher frequencies. However, it must be pointed out that these results are also entirely compatible with an informational limit hypothesis, since the effect of either increasing dot size or decreasing spatial frequency is to increase the mean spacing of elements in the image.

In sum, while the recent d_{\max} results have given rise to more sophisticated models of motion detection, they still do not allow the hypotheses of an informational limit, as opposed to a phase limit, in motion detection to be distinguished. Consequently, it is still unclear whether the distinction between short- and long-range motion can

be validated on the basis of d_{\max} . The aim of the present experiments is to investigate more thoroughly the effects of varying dot density on d_{\max} in an attempt to shed light on this issue. Lowering dot density increases the mean spacing of pattern elements, but leaves the relative energy across different spatial frequencies unaffected (although the absolute energy at all frequencies decreases). This means that any increase in d_{\max} at lower densities would be difficult to explain on the basis of either an absolute spatial-limit or a modified phase-limit hypothesis, but would be predicted on the basis of an informational-limit hypothesis.

GENERAL METHODS

Apparatus and stimuli

The stimuli used in all experiments were two-frame sequences of random-dot patterns (RDPs) in which the dots were displaced *en masse* between frames. The direction of motion could be either up or down and the kinematograms were viewed through a stationary window, to prevent pattern-edge displacement. These patterns were generated on a Commodore Amiga 2000 microcomputer and were displayed either on a Commodore 1084S RGB video monitor (Experiments 1 and 2) or on a Panasonic WV-5410 grey-scale monitor (Experiment 3), both of which ran at a refresh rate of 50 Hz. Viewing distance was 47.5 cm (Experiments 1 and 2) or 45.4 cm (Experiment 3) so that each dot (a single pixel) subtended 6×6 arc min in all cases. In general, the window or patch size was 25.4×25.4 arc deg, although Experiment 2 investigated the effects of varying this parameter.

Two methods were used for generating the RDPs, depending on the density. For high dot densities, five patterns of 360×256 pixels were generated prior to the experiment, and randomly sampled on each trial with full "wrap-around" (see Baker & Braddick, 1982). For cases where the expected number of dots in the display fell below 40 (as determined by the dot density and the patch size) the patterns were generated on-line for each trial. In these cases, the vertical co-ordinate of each dot that moved out of the patch window was wrapped around in order to keep the number of dots in each frame constant. However, to prevent observers perceiving reversed motion in these cases, each of these dots was given a new horizontal co-ordinate.

Photometric measurements were made with a Minolta Luminance Meter LS-110. The luminance of the bright dots under experimental conditions was 50 cd/m^2 , against a dark background of 0.06 cd/m^2 . This meant that the Michelson contrast of the displays, defined by $(\frac{L_{\max} - L_{\min}}{L_{\max} + L_{\min}})$, was 0.9976. The luminance of the screen area outside the stimuli window was also 0.06 cd/m^2 . These values were systematically varied in Experiment 3, which investigated the effects of contrast on d_{\max} .

Procedure

Subjects viewed a series of motion sequences each

containing a vertical displacement between the two frames and were required to indicate the perceived direction of the motion (up/down). The exposure duration of each frame was always 100 msec. Using only two frames meant that the stimuli contained the minimum amount of information required to make such a judgement. In addition, the brief exposures used ensured that eye movements did not confound the task. Viewing was binocular and the subject's head was supported in a chin-rest. The laboratory was dimly lit for all experiments.

Subjects fixated a central grey spot and initiated each trial with a key-press. The fixation spot was present throughout the trial, but disappeared with the removal of the second frame. Subjects indicated their decision by pressing one of two keys on a standard keyboard. Subsequently, the fixation spot re-appeared and signalled that the next trial was ready for them to initiate.

In a single sitting, subjects performed a block of 100 trials. This block was composed of five sets of 20 trials, each of which contained a different magnitude of pattern displacement. The range of displacements was chosen for each subject on the basis of a practice run to span the transition from errorless performance to around chance performance. In each set, there were ten downward displacements and ten upward displacements. Presentation order was randomized. For each condition (e.g., a particular patch size and dot density), subjects performed three blocks of trials so that each displacement was tested 60 times. The resulting data were then pooled across blocks. From the resulting psychometric function, d_{\max} was defined as the magnitude of displacement that elicited 80% correct direct discrimination following linear interpolation (Baker & Braddick, 1982).

Five subjects were tested, two females and three males, all between the ages of 20 and 29 years of age. Three were inexperienced psychophysical observers (LJS, BL and PAB) and were unaware of the purpose of the experiments.

EXPERIMENT 1: OPTIMAL INTER-STIMULUS INTERVAL (ISI) AS A FUNCTION OF DISPLACEMENT AND DOT DENSITY

The primary focus of this research was the spatial limits of motion detection. However, in pilot studies it was noticed that for very large displacements of a sparse RDK, observers had the impression of simultaneity of the two frames rather than of motion. This suggested that in order to elicit a convincing percept of motion, the temporal parameters of the display needed to be adjusted for the different magnitudes of displacements (see also Kolers, 1972). To this end, an experiment was designed to determine the optimum ISI duration for different dot density-displacement conditions. The data obtained from this experiment were then incorporated into the main experiments.

Methods

An immediate problem with measuring the optimal, rather than the maximal ISI durations (e.g. Baker &

Braddick, 1985) for motion detection is a floor effect. For a sparse pattern displacing a short distance, direction discrimination performance is close to perfect over a large range of ISIs. To counter this, a stimulus was developed in which the signal-to-noise ratio (SNR) could be varied, by altering the amount of correlation between the two frames.

Seven dot densities were used: 50, 20, 5, 1, 0.25 and 0.025% and all stimuli subtended 25.4×25.4 arc deg. The SNR was defined as the ratio of correlated to uncorrelated dots across the two frames. Two programs were employed to generate the stimuli. The first covered those conditions where the mean spacing of the dots within a single pattern was less than the spatial displacement being tested. In this case, the positions of each pair of noise dots were uncorrelated in the two frames.

The second program was used when the mean spacing of dots within a single pattern was greater than the displacement being tested. In this program, the noise dots plotted in frame 1 were given a displacement direction drawn randomly from one of seven orientations equally stepped (every 45 deg) around the signal displacement direction. The reason for this was to ensure that the displacement of the noise dots was as salient to the motion system as that of the signal dots, for any given ISI. Where mean spacing was smaller than the displacement (program 1), deliberately shifting each noise dot by the displacement magnitude would have been inconsequential as the nearest noise dot to it (across frames) would most likely *not* have been its counterpart. Thus, assuming nearest-neighbour matching, the spatio-temporal characteristics of the matched noise dots would not have been altered by using the more complicated program 2.

Optimal ISIs were measured for up to eight spatial displacements at each density. These displacements were 0.4, 0.8, 1.6, 2.4, 3.2, 4.0, 5.6 and 7.2 arc deg. Conditions in which errors never fell below 20% at any ISI when no noise was present in the stimulus were excluded from the analysis.

Two observers initially performed a pilot experiment in order to find the SNR for each condition that elicited peak performance of about 90–95% correct. In the experiment proper, the procedure outlined in the General Methods section was followed, except that the range of five spatial displacements in a block of trials was replaced by a range of five ISIs, with the spatial displacement fixed. The five ISIs spanned a linear range of 80 msec, in 20 msec steps. For a particular density-displacement condition, the optimum ISI was defined as the one that yielded fewest direction discrimination errors. If two ISIs produced the same number of errors then the smaller ISI was taken as the optimum.

Results and discussion

Figure 2 shows the optimal ISIs for both subjects as a function of spatial displacement for the seven dot densities. For higher densities, testing was carried out

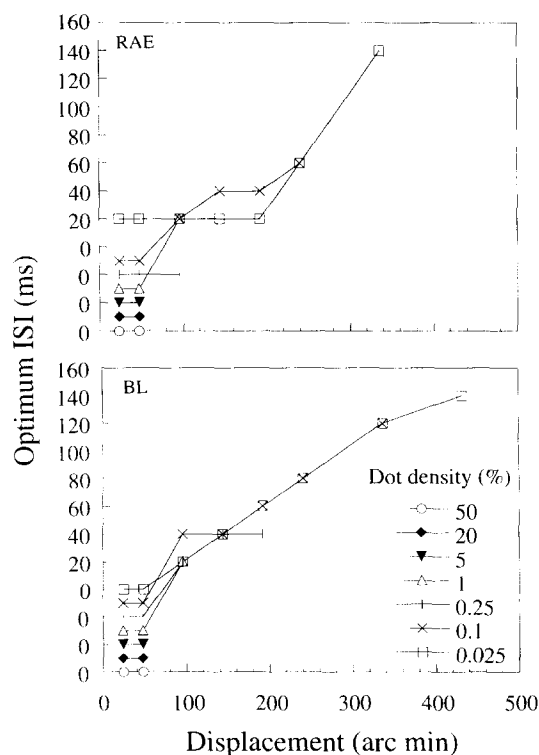


FIGURE 2. Optimal ISIs for two subjects over a range of dot densities and displacements. Note that the ordinate plots 0 msec at several positions for clarification of the data points.

over only a relatively small spatial range, as performance soon fell to chance for large displacements.

The data show that for displacements of 48 arc min and under, the optimum ISI was 0 msec for all conditions for subject RAE, and all conditions bar one for subject BL. These results fit well with those of a study by Baker & Braddick (1985). These authors measured d_{\max} for a low-density RDK as a function of exposure duration and ISI. For an exposure duration of 100 msec (as used here) they found that d_{\max} was highest when the ISI was 20 msec or less (their Fig. 4). Since d_{\max} was in the range of 50–60 arc min, this could be re-phrased by saying that the optimum ISI for a displacement of 50–60 arc min was found to be 20 msec or under. What the results of the present study add to Baker and Braddick's findings is that the optimum ISI appears to be dependent only on displacement and is unaffected by the dot density.

Figure 2 also shows that for displacements of about ≥ 1 deg and over, the optimum ISI began to increase. This increase was roughly linear for one subject, although more erratic for the other. No systematic dependency on dot density was observed. At the largest displacements tested, the optimum ISI reached approx. 140 msec. This observed trend is reminiscent of Korte's Third Law (Korte, 1915) which states that there is a proportional relationship between the optimum ISI and the magnitude of spatial displacement.

Interestingly, Baker & Braddick (1985) proposed that some of their data represented a breakdown in Korte's law. Using a 5 msec exposure duration, they showed that

larger displacements could be detected as the ISI was increased up to about 50 msec. However, they found that further increases in the ISI actually decreased, rather than increased, the displacement limit. They argued that this showed a separability of spatial and temporal tuning, rather than a trade-off between the two. They suggested that previous authors' data consistent with Korte's law (e.g. Kolers, 1972) were due to the fact that they had used isolated-element displays which had activated the long-range motion system. Baker and Braddick proposed that the results from their own experiments, based on RDK stimuli, reflected the properties of the short-range system.

In the present experiment, the data at high ISIs were obtained with low-density patterns. Thus, it could be argued that these stimuli activated the long-range process, and that only the dense patterns, where the optimum ISI was 20 msec or less, tapped the short-range process. However, as the optimum ISI was unaffected by dot density, and increases in the optimum ISI for larger displacements were not generally abrupt, it is more parsimonious to explain the results within the framework of a unitary motion system. Moreover, an alternative account of Baker and Braddick's data can be provided. Increases in ISI beyond 50 msec may not have led to increases in d_{\max} because detection was being limited by the spatial correspondence problem. If an ISI of about 50 msec was optimal for the 50 arc min d_{\max} value, as suggested by their data, then higher ISIs would have been sub-optimal for even the detectable displacements (under Korte's law), which might explain why d_{\max} actually declined. Further experiments are needed to decide which, if either, of these two hypotheses is correct.

EXPERIMENT 2: d_{\max} AS A FUNCTION OF DOT DENSITY AND PATCH SIZE

Stimuli

Five patch sizes were used with the length of each side equal to 25.4, 12.7, 6.4, 3.2, and 1.6 arc deg. Fixation was always in the centre of the patch. d_{\max} was measured at each patch size for up to seven dot densities: 50, 20, 5, 1, 0.25, 0.1 and 0.025%. All combinations of patch size and dot density where the mean expected number of dots in a frame exceeded 2 were investigated. The ISIs used in this experiment were derived from the data obtained in Experiment 1. For each condition, the ISI was set to the mean of the two subjects' individual optimal ISIs. When this was midway between realizable values (which were multiples of 20 msec) the lower value was used.

Results and discussion

Figure 3 plots d_{\max} as a function of dot density for the five different patch sizes. The broken data line in the top graph replots the data obtained by Baker & Braddick (1982). If d_{\max} was not dependent on dot density (but was dependent on patch size) the data would fall along vertically separated horizontal lines. This is clearly not the case and demonstrates that there is a dependency of d_{\max} on dot density. Three main aspects of the data can be noted.

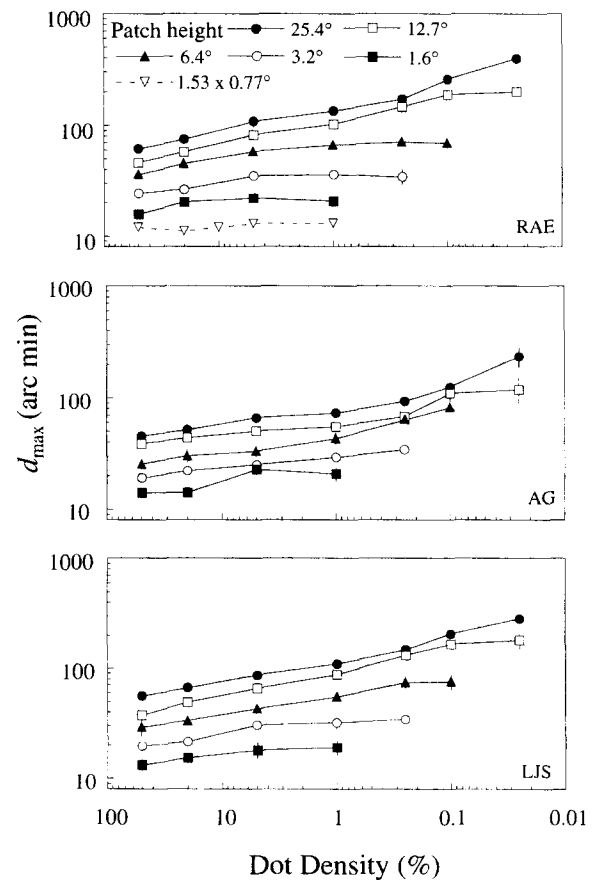


FIGURE 3. d_{\max} for three subjects for a range of dot densities and patch sizes. Note the reversed co-ordinates along the abscissa. The top graph also replots data from Baker & Braddick (1982) for comparison.

d_{\max} increases with increasing patch size. This is shown by the vertical offset of the five unbroken data lines. This finding is consistent with those of Baker & Braddick (1982), Chang & Julesz (1983) and Nakayama & Silverman (1984). In the Chang and Julesz and the Nakayama and Silverman studies, the RDK dot density was 50%. In the Baker and Braddick study, the absolute number of dots in the pattern was kept constant, while dot spacing was manipulated. The results presented here extend these findings by showing that this trend holds over a large range of dot densities from 50% down to 0.025%.

The effect is probably accounted for largely by the activation of mechanisms tuned to lower spatial frequencies as the patch extends further into the peripheral field. However, as discussed in the Introduction, how the motion system exploits these blurred parts of the image to extend d_{\max} is at present unclear. One possibility, as Baker & Braddick (1982) have suggested, is that there may be a "phase-based" limit whereby motion detectors can only register motion that occurs within displacements less than a constant fractional cycle (e.g., one-half) of the tuning of the detector. Alternatively, filtering the stimulus at a coarser scale may lead to larger values of d_{\max} simply because such filtering reduces the number of false targets in the display.

d_{\max} increases with decreasing dot density. This effect is illustrated by the slope of the data lines (note the reversed co-ordinates on the abscissa), a trend especially apparent for the larger patch sizes where d_{\max} rose from 1.0 to 6.6 deg for one subject (RAE). While this result can be contrasted with the findings of Baker & Braddick (1982), four other reports have described similar effects. Firstly, Ramachandran & Anstis (1983) used an oscillating RDK where the dot density was either 1 or 0.5%. They found that d_{\max} was higher for the 0.5% case over a range of stimulus-onset asynchronies. Secondly, Bischof & Groner (1985) measured d_{\max} with a 1-D strip of random dots placed around the circumference of a circle that rotated either clockwise or counter-clockwise. They found a small but consistent improvement in performance as dot density was reduced from 25 to 6.25%. Thirdly, Sato (1990) has reported how d_{\max} varies for two-frame kinematograms over a range of dot densities, with dots of 3 arc min. He found that d_{\max} was constant for dot density changes from 50–1% but rose rapidly as density was reduced further to 0.1%. Interestingly, all three sets of authors argued that their lower-density patterns probably stimulated the long-range process and, as such, were not inconsistent with Baker and Braddick's (1982) findings. Fourthly, Morgan & Fahle (1992) carried out an experiment with two-frame RDKs similar to the one reported here, in which they varied the dot density from 50% down to 5%. They found that d_{\max} tended to rise with decreasing density, for dot sizes above 4.5 arc min.

The rate of increase in d_{\max} with decreasing density decreases at smaller patches. This result is embodied in the flattening-off of the data lines at low densities, and is most marked for subjects RAE and LJS. It is clear that the data lines for the smallest patch used in the present study are quite similar to Baker and Braddick's results (shown on the top graph by the broken line), with further decrements in dot density below 5% having no effect on d_{\max} . One possible reason why density plays a lesser role for smaller displays is that large displacement values result in an increasing percentage of dots being displaced

out of the patch window. These uncorrelated dots are effectively noise, and cannot aid in the detection of the actual displacement. As patch size is reduced, this noise becomes an increasing proportion of the total number of dots in the display. It is possible that motion detection is limited to displacements that involve the presence of a significant percentage of correlated dots in both frames.

In order to examine this possibility, Fig. 4 shows the maximum values of d_{\max} for each patch size (taken across all densities) as a proportion of the patch size. It can be seen that the data are in fact remarkably consistent, both across subjects and patch size, with a mean of 20%, and a standard deviation, taken across subjects and patch size, of 3%. One implication of this finding is that Baker & Braddick (1982) may have found no effect of density upon d_{\max} because of their restricted patch size.

MODELLING d_{\max} WITH A SINGLE BANDPASS FILTER

The increase in d_{\max} with decreasing density is not consistent with the spatial-limit hypothesis in its original form or its modified phase-limit version. Is it quantitatively consistent with the predictions of the informational hypothesis? Within an image, as the dot density is halved the 1-D mean spacing of elements is doubled. However, masking data from Ball & Sekuler (1979) suggests that the search space is a sector emanating from the location of the frame 1 element. In this case, halving element density leads to only a $\sqrt{2}$ increase in the mean 2-D spacing. This is the result of the fact that while the area of the sector that contains the nearest neighbour will double when density is halved, the radius will increase by $\sqrt{2}$.

If it is the 2-D spacing of dots that determines d_{\max} , the exponent for the data functions should be -0.5 . In fact, the mean exponent for the best fitting slopes to four subjects' data at the largest patch size is -0.2 . Moreover, the slopes tend not to be straight, but steepen for densities below 1%.

Spatial primitives for motion detection

One explanation for this exponent of -0.2 might be that dots are not the primitives for the motion system. It is generally believed that multiple spatial-frequency tuned channels filter the incoming retinal image early on in the visual pathway (e.g. De Valois & De Valois, 1988). A simple revision to the informational hypothesis is that motion detection is based on the spacing of elements in the filtered stimulus, rather than the spacing of elements in the original stimulus.

Eagle & Rogers (1996) have described such a model in detail, the outline of which is given here. Based on psychophysical and physiological investigations of the motion system (e.g. Anderson & Burr, 1989; Baker & Cynader, 1986; Keck *et al.*, 1980; Watson & Turano, 1995), the first stage of the model is to pass the two frames of the RDK stimulus through a spatial filter bandpass in both frequency and orientation. The filters were constructed from difference-of-Gaussians that were balanced, and so passed no d.c. They were also




FIGURE 4. A bar-graph illustrating the maximum d_{\max} values for each patch size shown in Fig. 3 above, plotted as a proportion of the patch height for each of the three subjects.




FIGURE 5. (a) A 50% random-dot pattern. (b) A difference-of-gaussian filtered version of the pattern shown in (a), where the filter $f_{\text{peak}} = 21.8$ cycles per image. The peak of the filter's orientation tuning is orthogonal to the vertical axis of motion. The 2-D peaks of this image have been superimposed.

orientation-tuned, according to a Gaussian function. In the Fourier domain, these filters are defined in polar coordinates by Eq. (1).

$$F(f, \theta) = \left[e^{-\pi^2 f^2 2\sigma_c^2} - e^{-\pi^2 f^2 2\sigma_s^2} \right] \times e^{\left[\frac{-0.5(\theta - \theta_{\text{peak}})^2}{\beta} \right]}, \quad (1)$$

where f is the frequency, the radial distance from the origin, σ_c is the standard deviation of the centre Gaussian and σ_s is the standard deviation of the surround Gaussian, θ is orientation, θ_{peak} is the mean of the Gaussian filter (its peak orientation tuning) and β is the standard deviation of this Gaussian (half the bandwidth). σ_c and σ_s were always kept in a fixed ratio of 1:4.5 such that in the Fourier domain the half-gain full bandwidth was 2.6 octaves. This value was based on estimates derived from psychophysical experiments measuring d_{max} for different broad-band two-frame kinematograms (Eagle, 1996). θ_{peak} was set to 0 deg, orthogonal to the axis of motion. β was set to 15, giving a half-height orientation bandwidth of 35.25 deg.

A measure of element density in the images subsequent to any additional filtering clearly requires a definition of what constitutes an element. In this implementation, 2-D peaks were used as the primitive, defined at each pixel in the image whose luminance is higher than that of its eight immediately surrounding neighbours. An example of how the 2-D peaks relate to the image structure is shown in Fig. 5. The simple proposal is that it is the nearest-neighbour spacing of elements within an image that determines d_{max} . In the present analysis, the mean spacing of 2-D peaks was measured along the direction orthogonal to the filter θ_{peak} (parallel to the direction of

motion), within a sector of angle $\pm 2\beta$ —equal to 30 deg in the present case.

The model predicts that

$$d_{\text{max}} = km, \quad (2)$$

where m is the mean spacing of elements for the particular image and k is a constant scale factor. While in the most simple implementation of this model d_{max} would be limited to half the mean spacing of elements (the point at which 50% of the matches will be false) this free parameter allows for the possibility that global attributes of the system, such as probability summation or co-operativity (e.g. Chang & Julesz, 1984; Ullman, 1979) might extend d_{max} .

In Fig. 6, the functions show the predictions of how d_{max} varies with dot density according to the informa-




FIGURE 6. d_{max} as a function of dot density for the model described in the text for a range of different bandpass difference-of-gaussian filters. The absolute values of d_{max} can be varied by applying a constant scale factor, k , to each data point. In this plot, the value of k is set, arbitrarily, to 1.0 so that d_{max} is equal simply to the mean spacing of 2-D peaks in the filtered output.

tional model, for four different regimes of bandpass filtering. It can be seen from this figure that each of these functions essentially has two parts: a region at high densities where d_{\max} is independent of dot density and a region at low densities where there is a power-law relationship between density and d_{\max} , with the exponent asymptoting at -0.5 . These two regions can be explained intuitively as follows. When the filter is coarse relative to mean dot spacing, lowering the density has no effect upon 2-D peak density because the filter blurs the additional dots in the higher density pattern together. When the dot density is lowered, such that the mean spacing increases beyond the filter size, then gaps begin to appear in the filter output leading to a sparser distribution of 2-D peaks. The exponent of the function asymptotes at -0.5 as the dot density reaches a point where the probability of two dots being blurred together by the filter becomes vanishingly small. It can be seen that the critical dot density between these two stages is at lower dot densities for the lower-frequency filters.

A comparison with the psychophysical data in Fig. 3 shows that none of the functions provides a good fit. An increase in the slopes of the psychophysical data appears to occur around 0.25–1%, but the exponent never reaches -0.5 . This might be explained by the dependence of d_{\max} upon a filter so large that it has not begun to asymptote. The problem then is to explain the fact that there are increases in d_{\max} for changes at even the highest dot densities, which by this model implies the involvement of a very small filter.

One possible reason for why the slopes of the human data lines fall short of the theoretically predicted slope of -0.5 is that as density is decreased, the effective contrast also reduces (Eagle, 1992; Morgan & Fahle, 1992). Experiment 3 was designed to investigate this possibility.

EXPERIMENT 3: EFFECTS OF STIMULUS CONTRAST ON d_{\max}

There are several possible measures of contrast that might be applied to RDPs of varying dot density. The following observations are based on patterns comprising bright dots on a dark background. Michelson contrast is often defined as $(\frac{L_{\max} - L_{\min}}{L_{\max} + L_{\min}})$. Under this definition, contrast will be unaffected by changes in dot density since neither L_{\max} nor L_{\min} change. However, Michelson contrast can also be defined as $(\frac{L_{\max} - L_{\min}}{2L_{\text{mean}}})$. As dot density is decreased, contrast will actually increase according to this definition because L_{mean} will decrease. A different formulation, which takes into consideration all image points, is the r.m.s. contrast. On this basis, contrast decreases with decreasing density as the mean square luminance decreases. For these three measurement systems, the effect on contrast would be reversed if the patterns contained dark dots on a bright background.

If there is a range of spatial-frequency channels for motion detection, then it may be more valid to consider the contrast of the patterns at particular spatial scales, and methods exist for such a calculation (e.g., Peli, 1990). In order to understand what happens to the response of a

bandpass filter as dot density is varied, one approach is to consider the spectra of the different patterns. In fact, the only effect of changing dot density is that the mean energy level is lower for the sparser patterns—all are statistically flat across all orientations and spatial frequency (within limits set by the patch size and the dot size). At a given spatial scale, the number of “blobs” in the filtered images will remain relatively stable across density (down to a critical density, which can be identified in Fig. 6 as the “knee-point” of the function), but as the number of dots falling within the receptive field of the filter decreases, so too will the filter response. As density is reduced beyond this critical point, a different pattern emerges. As each dot is resolved by a single filter, the number of active mechanisms will decline with decreasing density, but the response of each active mechanism will remain stable.

Several psychophysical studies have demonstrated that motion-sensitive mechanisms saturate at Michelson contrasts as low as 0.03 for drifting sinusoids (e.g. Nakayama & Silverman, 1985; Keck *et al.*, 1980). In Experiment 2, although the Michelson contrast of the stimulus was high, it was distributed evenly across all spatial frequencies and orientations. Thus, it may have been the case that as dot density was decreased, a point was reached at which the pattern contrast failed to saturate the bandlimited mechanisms used to detect the motion. This would be particularly likely if a low-frequency filter was responsible for d_{\max} , owing to the fact that a flat stimulus spectrum contains lower energy in lower-frequency octaves (Field, 1987; Eagle, 1996). If this were the case, then further decrements in dot density would have led to increasingly lower effective contrasts. In turn, this could have caused the data functions to fall short of the predicted slope of -0.5 .

Stimuli

d_{\max} was measured for two subjects over a range of dot densities under three contrast conditions. One of the subjects, RAE, had performed the earlier experiments while the other, PAB, was inexperienced in these tasks and was naïve to their aims. The luminance value of the background screen, the inter-trial screen and the ISI was 10.0 cd/m² in all conditions. The luminance of the dots was either 77, 39 or 15 cd/m², depending on the contrast conditions. Defined as $(\frac{L_{\max} - L_{\min}}{L_{\max} + L_{\min}})$, the Michelson contrasts of these three conditions were: 0.77, 0.39 and 0.2. All the definitions of contrast outlined above yield the same ranking of contrast magnitude for these three conditions.

If d_{\max} was attenuated in the low-density conditions in Experiment 2 because of the decline in effective contrast, then this manipulation of stimulus contrast generates the simple prediction that d_{\max} should be lower in the low-contrast conditions.

Results and discussion

The results for two subjects are shown in Fig. 7. At densities below 0.25%, d_{\max} was lower for both subjects in the lower-contrast conditions. This finding suggests




FIGURE 7. d_{\max} for two subjects for a range of dot densities at three levels of Michelson contrast. The patch size was 25.4×25.4 deg, the largest used in Experiment 2.

that stimulus contrast does play a role in determining d_{\max} . It is consistent with the assumption that at the lower densities the motion-sensitive mechanisms responsible for determining d_{\max} are not saturated over the whole range of densities. That lower contrast levels have no detrimental effect on d_{\max} until density reaches 0.1% suggests that these mechanisms were saturated at higher densities. These results are in agreement with data from Morgan & Fahle (1992) who found that d_{\max} saturated at a lower contrast for a 50% density RDK than for a 5% density pattern.

Effect of mean luminance on d_{\max} . It is also noticeable that the tendency of d_{\max} to increase as dot density is decreased from 50 to 5% found in Experiment 2 has largely disappeared. The one exception to this is the high-contrast condition for RAE. The explanation for these differences may lie in how the stimulus mean luminance varies across dot density for the different contrast conditions. It is known that d_{\max} is inversely related to stimulus mean luminance, independent of contrast (Dawson & Di Lollo, 1989). For the high-contrast condition in Experiment 2, the stimulus mean luminance fell by a factor of 9.8 as density was decreased from 50 to 5%, compared with factors of 4.6, 2.1 and 1.2 for the three contrast conditions in the present experiment over the same density range. Thus, the increase in d_{\max} as density is decreased from 50 to 5% in the highest-contrast conditions may be an artefact of changes in stimulus mean luminance.

Interestingly, Baker & Braddick (1982) used a RDK with a Michelson contrast of 0.44, and so their stimulus was comparable to the mid-contrast condition here. Thus, it appears that while the failure to find an increase in d_{\max} as density was decreased from 5 to 1% may be an artefact

of their small patch size, a similar invariance at higher densities was found in the present experiment using larger patch sizes.

Further support for this hypothesis comes from a study by Morgan & Fahle (1992) who found that for dot sizes of 4.5 arc min and lower, d_{\max} decreased slightly when dot density was reduced from 50 to 5%. Since Morgan and Fahle used dark dots on a bright background, the mean luminance actually increased as density was reduced, and so this decline in d_{\max} would be predicted. Also using dark dots (of size 6 arc min) against a bright background, Eagle (1992) obtained a similar pattern of results. Interestingly, for larger dot sizes, Morgan & Fahle (1992) found a gentle increase in d_{\max} over this same density reduction. Of course, the effects of mean luminance change cannot explain this interaction of dot size and density upon d_{\max} . The explanation lies in the fact that d_{\max} should only be invariant of density (mean luminance effects apart) while the mean dot spacing is smaller than the extent of the visual system's blurring (see Introduction). As dot size is increased, the mean spacing of elements also increases and so the critical dot density at which d_{\max} starts to rise will be higher for larger dot sizes. Two points in Morgan and Fahle's (1992) informational-limit modelling of their own data are worth mentioning in these respects. Firstly, the model, which has no mechanism for considering mean luminance, predicts a slight increase in d_{\max} with decreasing density in the range of 50–5% for the small dot sizes—confirming that the basic informational model cannot account for the actual trend found. Secondly, the model predicts that the increase in d_{\max} with decreasing




FIGURE 8. d_{\max} replotted from Fig. 7 (high-contrast condition) against the model prediction for a difference-of-gaussian filter whose $f_{\text{peak}} = 0.6$ c/deg. The values of the scale factor k in Eq. (2), which determine the vertical positions of the model functions, have been individually set for each subject according to a least-squares fit to the data in the density region 50–1%.

dot density should be steeper when dot size is larger—testifying to the fact that the informational model can account for the interaction between dot size and density upon d_{\max} .

Re-examining the single filter model. It is instructive to compare these data with those from a study by Eagle & Rogers (1996). They also measured d_{\max} for two-frame kinematograms but now the stimuli were bandpass filtered RDKs of different initial dot densities. Both the luminance and the contrast of all stimuli were normalized and the consequence of the filtering was that only a narrow range of spatial-frequency channels were activated. In this simpler case, d_{\max} did reach an asymptotic slope of -0.5 . Indeed, the functions were well-fitted by the theoretical predictions of the single-filter model shown in Fig. 6. It is noteworthy that the data functions obtained in the present experiment have both a flat and sloping part to them which are qualitatively similar to those of the single-filter predictions, unlike those from Experiment 2. This lends support to the notion that when luminance and contrast are taken into account, d_{\max} is determined by an informational limit.

Figure 8 shows the data for the two subjects in the high-contrast condition together with the predictions of a filter whose $f_{\text{peak}} = 0.6$ c/deg (solid line). This filter was chosen so that the upturn in the human and model d_{\max} functions occurred at the same dot density. The divergence between the model and data for subject PAB at low densities has been explained by the loss in effective contrast at low densities. It is not clear why the slope for subject RAE is much closer to the theoretical prediction. Interestingly, the best-fitting power-law functions over the density range 1–0.025% for subject RAE in Experiment 2 (largest patch) and the high-contrast condition of the present experiment are -0.29 and -0.41 , respectively. This suggests that some learning may have taken place at the lowest densities.

The value of the constant k [see Eq. (2)], which determines the vertical position of the model function, was set using a least-squares technique to best-fit the human data in the density range 50–1%. The values of k were 0.95 for RAE and 0.97 for PAB. If d_{\max} was limited simply by nearest-neighbour matching, within the orientation band specified by the model, k should be no more than 0.5. It is possible that spatially global processes that pool the responses from individual motion detectors may act to increase d_{\max} beyond this point.

SUMMARY AND CONCLUSIONS

1. In Experiment 2, it was found that for large stimulus areas, d_{\max} increased by a factor of 5.5 (mean of three subjects) as dot density was decreased from 50% down to 0.025%. As stimulus area was reduced, this trend gradually disappeared so that by a stimulus area of 2.56 square deg, there was no effect of density (in the tested range of 50–1%) on d_{\max} . Baker & Braddick (1982) found no effect of dot density in the range of 50–1%, but as their

stimulus area was 1.18 deg², their results are compatible with the present findings. Since increasing displacement magnitude leads to an increased proportion of dots displacing out of the stimulus window, the size of the patch effectively imposes an upper displacement limit for motion detection. An analysis of the co-variation of d_{\max} with patch size in Experiment 2 suggests that this imposed limit is fixed at about 20% of the patch height.

2. The energy spectrum of a dense RDP is flat. As density is reduced, the effect on the spectrum is a uniform decrease in energy across frequency. This means that both the spatial-limit model of d_{\max} (Braddick, 1974) and the modified phase-limit hypothesis model, wherein d_{\max} is a fixed number of cycles of a spatial-frequency tuned sensor (e.g. Cleary & Braddick, 1990) predict that d_{\max} will not be affected by dot density. Thus, these data are clearly inconsistent with both these models. A simple informational-limit model, in which d_{\max} is proportional to the mean 2-D spacing of dots in the image, also fails quantitatively to account for the data, as the slope of the data function is much flatter than the predicted slope of -0.5 .
3. A modified informational hypothesis, in which d_{\max} is based on the 2-D spacing of 2-D peaks subsequent to a bandpass filtering stage, was then described. For all but the finest filters, the model predicts that d_{\max} should not be affected by changes in density in the high-density region. For all but the coarsest filters, the model predicts that d_{\max} should rise with a slope asymptoting to -0.5 at the lowest dot densities. As the empirical data increased across the whole range of decreasing densities, but never with this magnitude of slope, there appeared to be little support for this model.
4. Experiment 3 investigated the effects of contrast on d_{\max} . It was found that decreasing the Michelson contrast from 0.77 to 0.2 led to a decline in d_{\max} at dot densities below 0.25%. This suggests that as dot density is decreased, the reduction in stimulus energy means that the pattern becomes less effective in activating motion sensors.

A further consequence of reducing contrast was that the trend found in Experiment 2 for d_{\max} to rise as density fell from 50 to 5% disappeared. This is probably attributable to the fact that at lower contrasts the mean luminance did not decline as rapidly when dot density was lowered. In support of this, Dawson & Di Lollo (1989) have shown that lowering stimulus mean luminance can lead to an increase in d_{\max} , independently of changes in dot density.

In summary, the lower-contrast data functions remained flat in the density range 50–5% and then rose as density was reduced further. This pattern matched qualitatively the theoretical predictions of the single-filter informational model. The filter f_{peak} that best predicts the density at which the upturn in

the human data function occurs was ~ 0.6 c/deg. The model and data tend to diverge at lower densities, where the model predicts a slope of -0.47 over the range $1-0.025\%$. It is likely that this is due to the increasingly reduced effective contrast of the lower-density stimuli.

Interestingly, studies by Morgan (1992) and Eagle (1996) also suggest that d_{\max} for spatially broad-band patterns is based on the output of a single bandpass filter. Collating the results from various studies, Eagle (1996) suggested that the f_{peak} of this filter was ~ 0.47 c/deg, within half an octave of the present estimate.

5. The plausibility of a model in which d_{\max} is determined by a single informationally limited direction discrimination mechanism, subsequent to a stage in which the stimulus is bandpass filtered, has been assessed. It was found that such a model can account for the data reported here, as long as changes in the effective contrast at low densities and in the stimulus mean luminance at high densities are taken into consideration. This model is consistent with the results of an experiment by Eagle & Rogers (1996). They measured d_{\max} using RDKs of different dot densities spatially filtered with difference-of-gaussians at a range of scales, ensuring that both stimulus mean luminance and Michelson contrast were held constant throughout. With this much simpler stimulus, they showed that the pattern of results was entirely accounted for by the informational model described above.
6. It is not clear exactly what the predictions of a two-process model of motion detection would be for d_{\max} as dot density is reduced, other than that at some point d_{\max} should begin to increase. The possibility that the flat part of the data functions shown in Fig. 7 reflects the operation of a short-range process (fixed spatial limit, or fixed phase limit) while the rising part reflects the operation of the long-range process cannot be ruled out. However, if low-density RDKs are assumed to stimulate the long-range process because of the small number of features present, the same would have to be said for coarse-filtered high-density RDKs, which also contain few features. In turn, this would imply that the proportional relationship found between d_{\max} and the scale of bandpass filtering (Chang & Julesz, 1985; De Bruyn & Orban, 1989; Cleary & Braddick, 1990; Bischof & Di Lollo, 1990, 1991; Eagle & Rogers, 1996) is due to the increasing activation of the long-range process with lower-frequency stimuli. However, proponents of the two-process distinction have used this relationship as strong evidence for phase-limited short-range motion detectors (e.g. Cleary & Braddick, 1990). Thus, it is not clear how the range of data on d_{\max} can be treated consistently within a two-process framework.

Furthermore, in recent years the bases of several other characteristics that were thought to distinguish

between short- and long-range motion have been disputed, both on experimental and theoretical grounds (e.g. Cavanagh & Mather, 1989; Stout *et al.*, 1994). In summary, the adoption of a dichotomy of processes operationally distinguished on the basis of the stimulus density would seem unwarranted.

REFERENCES

- Anderson, S. J. & Burr, D. C. (1989). Receptive field properties of human motion detector units inferred from spatial frequency masking. *Vision Research*, 29, 1342-1358.
- Anstis, S. M. (1970). Phi movement as a subtraction process. *Vision Research*, 10, 1411-1430.
- Anstis, S. M. (1980). The perception of apparent movement. *Philosophical Transactions of the Royal Society of London, Series B*, 290, 153-168.
- Baker, C. L. & Braddick, O. J. (1982). The basis of area and dot number effects in random dot motion perception. *Vision Research*, 22, 1253-1259.
- Baker, C. L. & Braddick, O. J. (1985). Temporal properties of the short-range process in apparent motion. *Perception*, 14, 181-192.
- Baker, C. L. & Cynader, M. S. (1986). Spatial receptive field properties of direction selective neurons in cat striate cortex. *Journal of Neurophysiology*, 55, 1136-1152.
- Ball, K. & Sekuler, R. (1979). Masking of motion by broadband and filtered directional noise. *Perception and Psychophysics*, 26, 206-214.
- 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*, 30, 1341-1362.
- Bischof, W. F. & Di Lollo, V. (1991). On the half-cycle displacement limit of sampled directional motion. *Vision Research*, 30, 1341-1362.
- Bischof, W. F. & Groner, M. (1985). Beyond the displacement limit: an analysis of short-range processes in apparent motion. *Vision Research*, 25, 839-848.
- Braddick, O. J. (1974). A short-range process in apparent motion. *Vision Research*, 14, 519-527.
- Braddick, O. J. (1980). Low-level and high-level processes in apparent motion. *Philosophical Transactions of the Royal Society of London, Series B*, 290, 137-151.
- Cavanagh, P., Boeglin, J. & Favreau, O. E. (1985). Perception of motion in equiluminous kinematograms. *Perception*, 14, 151-162.
- Cavanagh, P. & Mather, G. (1989). Motion: the long and short of it. *Spatial Vision*, 4, 103-129.
- Chang, J. J. & Julesz, B. (1983). Displacement limits for spatial frequency filtered random-dot cinematograms in apparent motion. *Vision Research*, 23, 1379-1385.
- Chang, J. J. & Julesz, B. (1984). Cooperative phenomena in apparent movement perception of random-dot cinematograms. *Vision Research*, 24, 1781-1788.
- Chang, J. J. & Julesz, B. (1985). Cooperative and non-cooperative processes of apparent movement of random-dot cinematograms. *Spatial Vision*, 1, 39-45.
- Cleary, R. & Braddick, O. J. (1990). Direction discrimination for band-pass filtered random dot kinematograms. *Vision Research*, 30, 303-316.
- Dawson, M. R. & Di Lollo, V. (1989). Effects of adapting luminance and stimulus contrast on the temporal and spatial limits of short range motion. *Vision Research*, 30, 415-429.
- De Bruyn, B. & Orban, G. A. (1989). Discrimination of opposite directions measured with stroboscopically illuminated random-dot patterns. *Journal of the Optical Society of America A*, 6, 323-328.
- De Valois, R. & De Valois, K. (1988). *Spatial vision*. Oxford University Press.
- Eagle, R. A. (1992). *Spatial characteristics of human visual motion detection*. Unpublished D.Phil. thesis. University of Oxford, U.K.

- Eagle, R. A. (1996). What determines the maximum displacement limit for spatially broad-band kinematograms? *Journal of the Optical Society of America A*, 13, 408–418.
- Eagle, R. A. & Rogers, B. J. (1996). Motion detection is limited by element density not spatial frequency. *Vision Research*, 36, 545–558.
- Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A*, 4, 2379–2394.
- Julesz, B. (1971). *Foundations of cyclopean perception*. Illinois: University of Chicago Press.
- Keck, M. J., Montague, F. W. & Burke, T. P. (1980). Influence of the spatial periodicity of moving gratings on motion responses. *Investigative Ophthalmology and Visual Science*, 19, 1364–1370.
- Kolers, P. A. (1972). *Aspects of motion perception*. Oxford: Pergamon Press.
- Korte, A. (1915). Kinematoskopische untersuchungen. *Zeitschrift für Psychologie*, 72, 193–296.
- Lappin, J. S. & Bell, H. H. (1976). The detection of coherence in moving random dot patterns. *Vision Research*, 16, 161–168.
- Morgan, M. J. (1992). Spatial filtering precedes motion detection. *Nature*, 355, 344–346.
- Morgan, M. J. & Fahle, M. (1992). Effects of pattern element density upon displacement limits for motion detection in random binary luminance patterns. *Proceedings of the Royal Society of London, Series B*, 248, 189–198.
- Nakayama, K. & Silverman, G. H. (1984). Temporal and spatial characteristics of the upper displacement limit for motion in random dots. *Vision Research*, 24, 293–299.
- Nakayama, K. & Silverman, G. H. (1985). Detection and discrimination of sinusoidal grating displacements. *Journal of the Optical Society of America A*, 2, 267–274.
- Peli, E. (1990). Contrast in complex images. *Journal of the Optical Society of America A*, 7, 2032–2040.
- Petersik, J. T. (1989). The two-process distinction in apparent motion. *Psychological Bulletin*, 106, 107–127.
- Ramachandran, V. S. & Anstis, S. M. (1983). Displacement thresholds for coherent apparent motion in random dot patterns. *Vision Research*, 23, 1719–1724.
- Sato, T. (1990). Effects of dot size and dot density on motion perception with random dot kinematograms. *Perception (Suppl.)*, 19, 329.
- Stout, J., Pantle, A. & Mills, S. L. (1994). An energy model of interframe interval effects in single-step apparent motion. *Vision Research*, 34, 3223–3240.
- Ullman, S. (1979). *The interpretation of visual motion*. Cambridge, MA: MIT Press.
- Watson, A. B. & Turano, K. (1995). The optimal motion stimulus. *Vision Research*, 35, 325–336.
- Zeeman, W. P. C. & Roelofs, C. O. (1935). Some aspects of apparent motion. *Acta Psychologica*, 9, 159–181.

Acknowledgements—Parts of this work were presented in a paper given at the ARVO 1991 conference held in Sarasota, Florida. This work was funded by a UK SERC postgraduate studentship and a Wellcome Research Training Fellowship both awarded to RAE under the sponsorship of BJR. We thank Andrew Glennerster and Oliver Braddick for stimulating discussion of the project.