

What determines the maximum displacement limit for spatially broadband kinematograms?

Richard A. Eagle

Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford OX1 3UD UK

Received June 12, 1995; revised manuscript received September 26, 1995; accepted October 5, 1995

Two experiments are described that are designed to investigate what determines the maximum spatial displacement detectable (d_{\max}) for spatially broadband patterns exposed in a two-frame motion sequence. In experiment 1, d_{\max} was found to be 1.63 times greater for a two-dimensional (2-D) broadband random pattern with a $1/f$ Fourier amplitude spectrum (equal contrast in each octave) than for a 2-D binary-valued random-dot pattern with a flat spectrum (higher contrast in higher-frequency octaves). In experiment 2, d_{\max} was shown to vary in inverse proportion to the lowest stimulus frequency for random patterns with a one-octave bandwidth and normalized contrast. Furthermore, when these five one-octave patterns were summed together, d_{\max} for this new five-octave pattern was found to be only 1.46 times lower than d_{\max} for the lowest-frequency one-octave pattern presented alone. A model is described in which direction discrimination is based on the nearest-neighbor matching of zero crossings in the output of a single-spatial-filter bandpass in both spatial frequency and orientation. Data from the model show that the difference between d_{\max} for the five-octave and the lowest one-octave patterns can be accounted for by the same filter passing some of the additional higher frequencies in the former pattern. Furthermore, it is argued that all the data can be accounted for by assuming that d_{\max} is determined by the coarsest spatial filter activated by each stimulus. Modeling the results of both experiments suggests that the bandwidth of this filter is ~ 2.6 octaves and reaches peak sensitivity at ~ 0.47 c/deg. The model is shown to be capable of accounting for a wide range of other two-frame d_{\max} data. © 1996 Optical Society of America

1. INTRODUCTION

There is substantial psychophysical and physiological evidence to suggest that a range of filters or channels exist in the human visual system, each one independently tuned to analyze the retinal image at a particular spatial scale (see De Valois and De Valois¹ for a review). Several lines of evidence also indicate that this spatial-frequency selectivity is preserved to the level of motion detection. For instance, Keck *et al.*² and Cameron *et al.*³ have shown that the motion aftereffect, induced by prolonged exposure to a drifting sinusoid, is selective to the spatial frequency of the adapting grating. Similarly, Anderson and Burr,⁴ using the masking paradigm, have demonstrated that direction-discrimination thresholds are higher when the test and mask sinusoids have the same frequency. Physiologically, it has been reported that directionally selective cells in cat striate cortex are also tuned to the spatial frequency of the stimulus.⁵ In addition, most current computational models of motion detection specify that local motion sensors are bandpass in spatial frequency.⁶⁻⁸

Sinusoidally modulated luminance patterns have been used in many of these studies because such stimuli selectively activate a narrow range of spatial-frequency channels. A quite distinct tradition in the psychophysical study of motion detection has developed through the use of random-dot kinematograms^{9,10} (RDK's). In the minimal case, two frames of random dots are successively exposed, the second being a translated version of the first. Under appropriate conditions this stimulus elicits a strong sensation of smooth motion, despite the fact that the physical displacement is discrete and instantaneous. An important variable known to affect observers' ability to detect the direction of motion in RDK's is the magnitude of the

displacement. Braddick¹⁰ originally found that the maximum displacement detectable (d_{\max}) was near 15 arcmin. For larger displacements the detection process breaks down and the predominance of the false matches made by the motion system results in a percept of spatially incoherent motion.

Unlike sine waves, random-dot patterns (RDP's) have a broad orientation and spatial-frequency spectrum. Over the past decade, these two-stimulus-based approaches to motion detection have ostensibly been brought closer together through the use of spatially filtered RDK's that contain only a narrow band of frequencies. Using spatially bandpass RDK's, several investigators have shown that d_{\max} is inversely proportional to the lowest frequency in the stimulus.¹¹⁻¹⁶ This relationship is linked to the quasi-periodicity of these stimuli^{13,15} and can be interpreted in terms of a displacement limit determined by the spacing of false targets in the images.^{16,17}

This finding raises the intriguing question of how d_{\max} is determined for a spatially broadband pattern. A RDP composed of equiprobable bright and dark pixels has statistically flat spectrum such that components at all spatial frequencies and orientations are present at roughly equal amplitude (within limits set by the pixel and image size). Consider how such a stimulus would be analyzed by a set of motion filters, with each channel tuned to a narrow range of spatial frequencies. The aforementioned findings with bandpass RDK's imply that the upper displacement at which all channels can signal the correct direction of motion will be determined by the highest-frequency channel. Beyond this displacement, the finest filter will signal only the spurious motions that fall within its range, although coarser filters will still be able to detect the correct direction of motion. Clearly, then, in

such cases it would be advantageous if the response of the motion system was based on the output of the lowest-frequency filter. However, this strategy requires that the system be organized so that the decision mechanism has independent access to the motion signals stemming from each channel.

Some data relevant to this issue were gathered in an experiment by Chang and Julesz.¹⁸ They found that d_{\max} was generally greater for a low-pass-filtered RDK, containing no high frequencies, than for an unfiltered RDK. In a more systematic study, Cleary and Braddick¹⁹ found that the high-frequency cutoff of a rectangular low-pass filter could be reduced to 3.56 cycles/degree (c/deg) without affecting the magnitude of d_{\max} but that further filtering led to large increases in d_{\max} . As they pointed out, low-pass filtering does not add anything to the stimulus; obviously, it simply removes high frequencies.²⁰ On the basis of these results, Cleary and Braddick argued that d_{\max} is limited by the *highest* frequencies in the image that are passed by the motion system. Because for large displacements high-frequency channels will effectively signal only an incoherent pattern of motion, they suggested that these channels mask the coherent signal of the lower-frequency channels. The effect of low-pass filtering, then, is to release the low-spatial-frequency information from masking, which, in turn, leads to an increase in d_{\max} .

One problem with Cleary and Braddick's hypothesis concerns the fact that mild low-pass filtering does not elevate d_{\max} .¹⁹ To account for this, they argued that the visual system must low pass the stimulus before any motion detection, such that blurring out stimulus components beyond 3.56 c/deg has no discernible effect on the detection process. However, Cleary and Braddick¹³ showed that under similar experimental conditions, observers had no difficulty in discriminating the direction of motion for bandpass RDK's when the lowest frequency present in the stimulus was 7.1 c/deg. It is difficult to reconcile their low-pass-filtering hypothesis with this finding, which strongly implies that higher-frequency filters do exist for motion detection.

The experiments reported in the present paper were designed to address the issue of how d_{\max} is determined for spatially broadband patterns. In addition, a model of early visual filtering and motion detection was developed for testing particular hypotheses developed from the data.

2. GENERAL METHODS

A. Apparatus and Stimuli

A RDP of 50% density and measuring 512×512 pixels was generated by a spatially random assignment of high- and low-luminance dots, with equal probability. This pattern was then spatially filtered in the Fourier domain with the Heritable Image Processing System software package²¹ to produce a range of novel patterns. (Further details are given in the descriptions of methods for experiments 1 and 2.) Frequencies below 2 cycles per image (0.167 c/deg) were removed so that there would be no mean luminance variation for any of the subsampled stimulus displays (see below).

Before the experiment the filtered 512×512 image was split up into two smaller images of 512×256 pixels,

which were then contrast modulated by a vertically oriented one-dimensional (1-D) Gaussian with an extended plateau. This procedure aimed to minimize the contrast of artifactual spectral components introduced at the image edges parallel to the motion. The value of the contrast modulation was 1.0 for the central 512×209 pixels and fell off for pixels outside this region according to a Gaussian of space constant 16 pixels. The actual stimulus images, measuring 360×256 pixels, were extracted from these 512×256 images on line on each trial, with a random horizontal start coordinate and full wraparound. The major, horizontal, axis was parallel to the axis of motion.

Stimulus presentation was controlled by a Commodore Amiga 2000 microcomputer, which also collected subjects' responses. The stimuli were displayed on a Panasonic WV-5410 gray-scale monitor (white P4 phosphor) with a refresh rate of 50 Hz. The viewing distance was 195 cm, at which a single pixel subtended 1.4 arcmin and the stimulus window subtended 7.5×6.0 arcdeg.

The images were defined by 32 gray levels, and the display screen was linearized with the aid of a Minolta Luminance Meter LS-110 photometer. The maximum attainable screen luminance value was 76 cd/m², and the minimum attainable luminance value was 10 cd/m². All stimuli were exposed against a mean luminance background of 43 cd/m². The contrast of each stimulus is specified in the descriptions of methods for experiments 1 and 2. The contrast of the artifactual components introduced by the gray-level quantization process was assessed by considering the difference between a 32-integer valued version and a floating-point valued version of the same RDP passed through a one-octave rectangular filter (passband = 4–8 cycles per image). The Michelson contrast of the difference pattern, excluding components introduced within the filter passband, was only 0.014: 27.3 times lower than the signal contrast of the same pattern specified in floating-point values. The spectrum of this pattern of artifactual components was statistically flat. A control study showed that direction discrimination of a binary-valued RDK with a Michelson contrast of 0.02 around a mean luminance of 43 cd/m² was at chance over all displacements. This result is consistent with the finding of Morgan and Fahle²² that the motion of a RDK whose Michelson contrast is 0.03 was undetectable. Furthermore, it suggests that the artifactual components introduced into the present study's stimuli by the gray-level quantization process would have had no effect on subjects' performances.

B. Procedure

Subjects viewed a single motion sequence containing a discrete horizontal displacement and were required to indicate the perceived direction of the motion (left/right). The exposure duration of each frame was always 100 ms, and there was no interstimulus interval. The use of only two frames meant that the stimuli contained the minimum amount of information required for making such a judgment. In addition, the brief presentations 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 gray spot and initiated each trial with a key press. The fixation spot was present throughout the trial but disappeared as the second frame was removed. Subjects pressed one of two keys on a standard keyboard to record a decision. Subsequently the fixation spot reappeared, signaling that the next trial was ready to be initiated.

A block of trials consisted of 100 presentations. In a single block there were 5 sets of 20 trials, each of which contained a different magnitude of displacement. In each set there were 10 leftward displacements and 10 rightward displacements. Presentation order was randomized. Subjects performed three blocks of trials for each condition, making a total of 300 trials in all, 60 trials for each displacement value. From the resulting psychometric function, d_{\max} was defined as the displacement that produced 20% errors following linear interpolation of the data points.²³

Five subjects participated in the experiments: one was the author, and the other four were experienced psychophysical observers who were unaware of the purpose of the experiments.

3. EXPERIMENT 1: d_{\max} FOR BROADBAND KINEMATOGRAMS

This experiment was designed to measure d_{\max} for two different spatially broadband patterns. It is important to note that in previous experiments that measured d_{\max} with bandpass-filtered RDK's, the stimulus contrast was normalized after the filtering.¹¹⁻¹⁵ These investigators used rectangular filters that in the Fourier domain had a constant gain and bandwidth (measured in octaves) and varied only in the passband center frequency. Because RDP's have a flat spectrum, the Fourier energy contained in any two-dimensional (2-D) frequency band is proportional to $f_h^2 - f_l^2$, where f_h and f_l are the high and low cutoff frequencies. This means that the energy passed by a high-frequency filter will be four times that passed by the spectrally adjacent low-frequency filter. In turn, this means that without normalization the contrast of the higher-frequency bandpass pattern will be twice that of the low-frequency pattern.

Parseval's theorem²⁴ shows that for a spatially random 2-D broadband pattern to have the fractal property of yielding patterns of equal contrast when passed through such filters, it must originally contain equal energy in each octave band. This means that its energy spectrum needs to follow a $1/f^2$ slope.²⁵

Examples of a RDP and the same pattern scaled in the Fourier domain by a $1/f$ transfer function to yield what will be termed here a fractal-noise pattern (FNP) (with a $1/f^2$ energy spectrum) are shown in Fig. 1. The first experiment was designed to compare d_{\max} for a RDK and for a fractal-noise kinematogram (FNK).

A. Stimuli

A FNP was generated in the Fourier domain by passing the Fourier transform of a RDP through a $1/f$ transfer function and then taking the inverse Fourier transform. To normalize the contrast for the FNP and the RDP, the standard deviation of luminance was set to 11 cd/m^2 for both patterns, the frequency distribution of

luminance values being Gaussian for the FNP and binary valued for the RDK. In the latter case this meant that the maximum and minimum luminance values were 54 and 32 cd/m^2 , respectively.

B. Results and Discussion

Figure 2 plots the d_{\max} values for each subject for the two types of kinematogram. The graph shows that d_{\max} for the FNK was, on average, 1.63 times larger than that found for the RDK. Given the different spectra of the two patterns and the fact that low frequencies are clearly important in determining d_{\max} , it is important to consider whether this difference could be due to the low frequencies in the random-dot images falling below their detection threshold. As a control, subject RAE performed the same experiment with the binary-valued RDP at a Michelson contrast of 0.77, compared with 0.26 for the

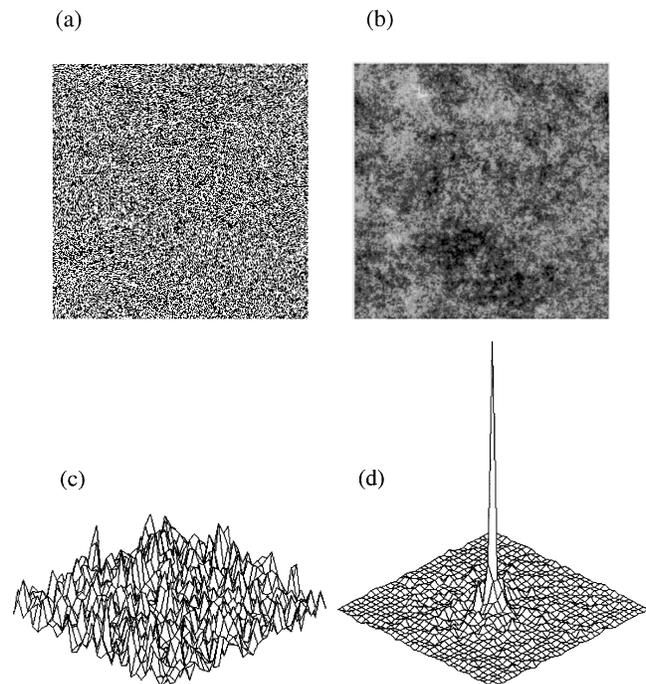


Fig. 1. (a) RDP whose spectrum is statistically flat; (b) the same pattern, whose amplitude spectrum has been scaled according to a $1/f$ function; (c), (d) 2-D amplitude spectra of the above images in polar plots. The x and y axes plot horizontal and vertical spatial frequencies in cycles per image sampled into 8×8 regions. To reflect the contrast of the FNP used in experiment 1 relative to that of the RDP, the spectrum of the FNP needs to be scaled by a factor of 3.

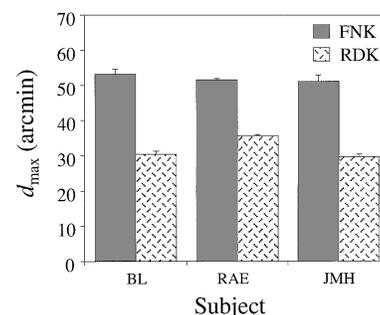


Fig. 2. d_{\max} values for a RDK and a FNK, as illustrated in Fig. 1, for three observers. The error bars show the standard error of the mean for d_{\max} across three runs of each condition.

patterns used in the main experiment. The mean luminance of the stimulus was not altered. As raising the contrast of this pattern uniformly increases the energy at all frequencies, a clear prediction is that if d_{\max} was lower only in the main experiment for the RDK because of the subthreshold low-frequency components, raising the contrast should increase d_{\max} . Experimentally this result was not observed: d_{\max} was 33.4 arcmin for the high-contrast condition, compared with 35.7 arcmin for the original, lower-contrast pattern. This shows that the difference in performance for the RDK and the FNK cannot be attributed simply to subthreshold low-frequency components in the former stimulus.

These data show that the visual system is able to exploit the increased *relative* energy at low frequencies in the FNK when discriminating the direction of large displacements. This finding raises the possibility that d_{\max} for the FNK might be comparable to that for a pattern containing only low frequencies. The aim of the following experiment was to compare d_{\max} for broadband and narrowband kinematograms that were matched for contrast within a specified spectral band, to make possible a quantitative analysis of this issue.

4. EXPERIMENT 2: d_{\max} FOR BROADBAND VERSUS NARROWBAND KINEMATOGRAMS

Broadband patterns were constructed by summing same-contrast one-octave band patterns together. d_{\max} was measured for both the one-octave patterns and the summed pattern. If d_{\max} is dependent on the highest frequencies in the stimulus, then we might expect its value for the broadband kinematogram to be similar to that for the highest-frequency one-octave kinematogram. In contrast, if d_{\max} is dependent on the lowest frequencies in the stimulus, then its value for the broadband kinematogram should be comparable to that for the lowest-frequency one-octave kinematogram.

A. Stimuli and Procedure

Five filters were set up in the Fourier domain, each with a sharp high- and low-frequency cutoff. Each filter had a one-octave bandwidth and a $1/f$ gain within its passband (see Fig. 3). The spectral range of each filter was 0.33–0.67, 0.67–1.33, 1.33–2.67, 2.67–5.33, and 5.33–10.67 c/deg. In the Fourier domain each filter was multiplied by the Fourier transform of a RDP. Because of the structure of the filters the inverse Fourier transforms of the five products had roughly the same contrast. Any small deviations were eliminated by scaling each pattern to have a mean luminance of 43 cd/m² with a standard deviation equal to 3.9 cd/m² (the luminance distributions were Gaussian). The Michelson contrast of the patterns was approximately 0.38. A five-octave pattern was produced by summing these bandpass patterns, with the mean luminance kept constant. The standard deviation of luminance for this image was 8.6 cd/m². Examples of all six patterns are shown in Fig. 4.

The five-octave pattern was used in preference to the broader-band $1/f$ noise pattern used in experiment 1, which contained energy over the range of 0.167–21.4 c/deg, to permit more quantitative comparisons to be made with data from the one-octave stimuli.

B. Results and Discussion

Figure 5 shows the d_{\max} values for three subjects for the six stimuli. d_{\max} is plotted against the lowest frequency (f_1) present in the stimulus. The three isolated data points represent d_{\max} for the five-octave kinematogram.

1. d_{\max} for One-Octave Stimuli

The one-octave data replicate the previously known result^{11–16} that d_{\max} rises as f_1 is decreased. These

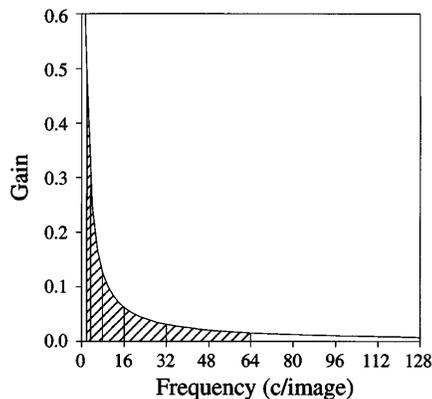


Fig. 3. 1-D slice through the $1/f$ transfer function of the filter used to produce the image shown in Fig. 1(a). The hatched region shows the five-octave band used to produce the images in Fig. 4. Each octave band is marked by a vertical line.

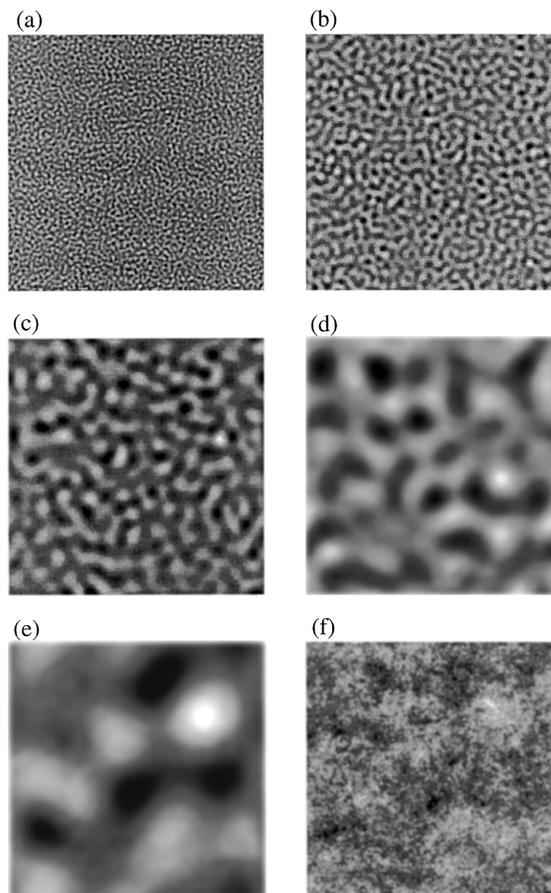


Fig. 4. (a)–(e) show the responses of one-octave-wide filters with sharp frequency cutoffs and a $1/f$ gain within their passbands to a RDP, with f_1 progressively halving. (f) shows the pattern obtained by simply summing these five patterns while the mean luminance is kept constant.

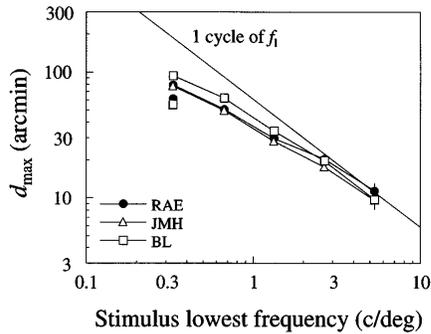


Fig. 5. d_{\max} values for three subjects for a range of bandpass kinematograms, plotted against the pattern's lowest frequency. The symbols linked by lines represent the data for one-octave stimuli. The isolated data points at $f_1 = 0.33$ c/deg represent the data for the five-octave stimulus generated by summing the single-octave patterns together. The error bars show the standard error of the mean for d_{\max} across three runs of each condition.

data may reflect the activity of a set of spatial-frequency channels tuned for motion, although the possibility that a single more broadly tuned channel can account for all of the data cannot be ruled out. A closer look at the data reveals that d_{\max} was largest for the high-frequency stimuli when expressed in cycles of f_1 . This can be seen in Fig. 5 by a comparison of the data to the solid line, which shows a power-law function of exponent -1 . A similar effect has been observed in other experiments that measure d_{\max} for narrowband stimuli.^{13,15,16} If d_{\max} is based solely on either the spacing of image features (e.g., the blobs) or on the spectral content of the patterns, then it should follow a slope of -1 .

Bischof and Di Lollo¹⁵ have argued that this trend in the data is due to the increased contrast sensitivity at higher frequencies. Their model of motion detection, which incorporates weightings based on the contrast-sensitivity function for direction discrimination of displacing sinusoids, fits this aspect of the data well. However, it may be inappropriate to invoke the contrast-sensitivity function, which is measured at threshold, to explain these data obtained with suprathreshold stimuli.²⁶

Another possible reason for this trend is that probability summation or some form of spatial co-operativity may be involved in motion detection.^{27,28} Because the patch size was held constant in angular size rather than in cycles of f_1 , there would have been a greater number of motion detectors activated by the higher-frequency stimuli. It may be that keeping patch size constant in cycles of f_1 would yield comparable magnitudes of d_{\max} (also in cycles of f_1).

2. d_{\max} for Five-Octave versus One-Octave Stimuli

In experiment 1 the mean value of d_{\max} for the RDK was 32 acrimin—well below d_{\max} for the lowest-frequency-bandpass kinematogram used in this experiment (mean value of 83.26 acrimin for $f_1 = 0.33$ c/deg). As in previous studies that have found this effect,^{12,14,18} d_{\max} for the unfiltered RDK in the present experiment was similar to that for the bandpass kinematogram whose $f_1 = 1.33$ c/deg.

More interesting are the data for the five-octave patterns, shown in Fig. 5 by the isolated symbols, as this

stimulus can be more easily related to the individual one-octave stimuli. As can be seen, the d_{\max} value for this stimulus was roughly equal to d_{\max} for the single-octave stimulus, whose $f_1 = 0.67$ c/deg, and is only 1.46 times smaller than that for the lowest-frequency single-octave stimulus.

The lowest-frequency one-octave pattern can be considered a low-pass version of the five-octave pattern, and so these results can be compared with data from Cleary and Braddick,¹⁹ who measured d_{\max} for low-pass-filtered RDK's. In their experiment the ratio of d_{\max} values for two stimuli where $f_h = 0.89$ and 7.1 c/deg was 2.63. In the present experiment the d_{\max} ratio for two stimuli where $f_h = 0.67$ and 10.67 c/deg (the single-octave and five-octave stimuli, respectively) was only 1.46. That is, compared with Cleary and Braddick's findings, the present results show a much smaller effect of low-pass filtering on d_{\max} . This discrepancy is almost certainly due to the spectral differences between the RDK stimuli used by Cleary and Braddick¹⁹ and the FNK used here and is discussed further in Section 5.

That d_{\max} was at all higher for the lowest-frequency one-octave stimulus than for the five-octave stimulus could be interpreted as support for Cleary and Braddick's¹⁹ hypothesis. However, such an interpretation is not necessary. A simpler alternative explanation is that a single, relative broadband filter was responsible for performance on both the lowest-frequency one-octave stimulus and the five-octave stimulus. In this case, even if the filter's peak frequency tuning were centered on the lowest-frequency octave, it would have passed higher frequencies in the five-octave pattern that were absent in the one-octave pattern. In turn, this would have increased the number of false targets in the filtered five-octave kinematogram and so lowered d_{\max} relative to that for the one-octave pattern. This argument has some of the flavor of that of Cleary and Braddick, but the crucial difference is that any masking proposed by the scheme suggested here occurs within a *single* channel rather than between channels. That is, there is no need to posit interactions between channels to explain the differences between d_{\max} for the broadband and narrowband stimuli found here; all data can be accounted for by considering the activity within a single filter. This does not necessarily mean that there is only a single channel in existence but merely that for a given stimulus only the signal from one channel determines d_{\max} . This channel, if there is more than one, may differ from stimulus to stimulus, and the simplest rule would be that d_{\max} is determined by the signal from the lowest-frequency channel activated.

5. MODELING d_{\max} WITH A SINGLE-BANDPASS FILTER

A. Model

This section describes a four-step computational model of direction-discrimination performance for two-frame kinematograms. Figure 6 summarizes the model pictorially. The steps are as follows:

1. To model the tuning properties of cortical simple cells (see, e.g., Hawken and Parker²⁹), the two frames of the kinematogram were convolved with a filter bandpass

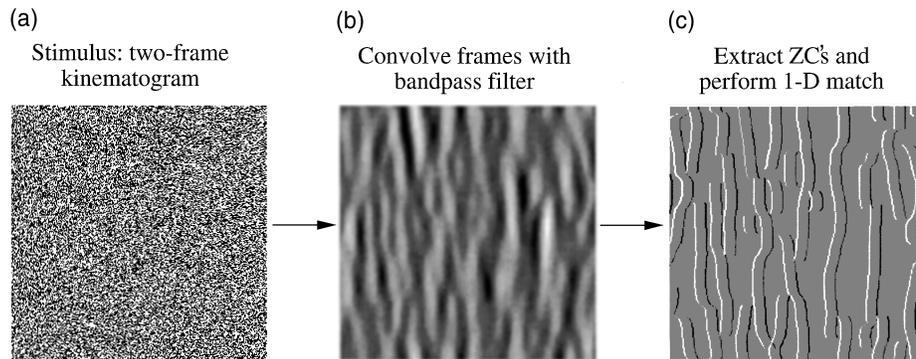


Fig. 6. Stages of an algorithm designed to discriminate the direction of motion in two-frame kinematograms. Two frames of a rightward-displacing image, such as a RDP [(a)], are initially passed through a filter bandpass in both spatial frequency and orientation, with the orientation of the filter orthogonal to the axis of displacement [(b)]. The upward-going zero crossings are then extracted from the two filtered images and superimposed [(c)], with white showing the first frame and black the second frame. Single-pixel zero crossings (ZC's) are then matched along the axis of displacement. d_{\max} is taken as the displacement at which 60% of the matches are in the correct direction.

in both orientation and spatial frequency. The spatial-frequency tuning was achieved by taking the difference of [two] Gaussians (DOG). In the spatial domain the two Gaussians were isotropic and balanced, so that no dc was passed. The ratio of the two space constants, and so the filter bandwidth, was variable. The peak frequency (f_{peak}) of the filter was determined by the absolute values of the space constants, as described for individual cases below. The orientation tuning was achieved in the Fourier domain by multiplying the Fourier transform of the DOG by a function that was Gaussian bandpass in orientation and with a peak tuning orthogonal to the horizontal axis of motion. In polar coordinates, the 2-D Fourier transform of this filter is defined as

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

where f is frequency, θ is orientation, σ_c and σ_s are the standard deviations of the frequency-tuned center and surround Gaussians, respectively, and θ_{peak} is the mean of the orientation-tuned Gaussian (its peak tuning) and β its standard deviation (set at 15° , giving a half-height bandwidth of 35.25°).

2. The second step was to extract the zero crossings of the two filtered images (see also Marr and Ullman⁸ and Morgan¹⁷). For ease of computation, only negative-to-positive zero crossings, going from left to right across the image, were taken (notice that this has no consequences for the model's performance).

3. In the third step, single-pixel zero-crossing segments were matched to the nearest segment in the second frame, along the axis of displacement. No other constraints, such as enforcing that a segment have only one match, or be matched only to a segment of similar slope (effectively of similar contrast) were employed. It is important to note that the choice of zero crossings is not essential to any of the modeling results that follow; other primitives such as 1-D centroids or 2-D peaks could also have been used with similar results.³⁰

4. The estimation of the overall direction of motion was computed by separately summing the number of

matched zero crossings displaced to the left and to the right. d_{\max} was taken as that displacement at which 60% of the directional judgments (disregarding magnitude) were correct. This choice of the 60% point was essentially arbitrary, and increasing or decreasing this value would affect the absolute values of d_{\max} arrived at by the model.

In summary, there are three crucial aspects of this model. First, direction discrimination is spatial frequency tuned. Second, the upper displacement limit is determined by the spacing of false targets: there is no absolute spatial limit determining the range over which matches are sought. It is important to recognize that this aspect of the model, while being supported by recent psychophysical data,^{16,17} is not essential to the main computational results reported in the following sections. For instance, d_{\max} would be limited similarly in the well-known model of the elaborated Reichardt detector.^{6,7} For elaborated Reichardt detectors two sensors generally have a spatial offset that is inversely proportional to their spatial-frequency tuning (in order to avoid aliasing). In the present model, there is no hard-wired spatial limit, but the quasi-periodicity introduced by the bandpass-filtering properties on the front-end filter coupled with the fact that the images are spatially dense means that aliasing will set in beyond displacements of near half a cycle of the peak frequency of the input image. Thus the values of d_{\max} in both models will be directly proportional to one another and will be similarly affected by changes in the spectral content of the stimulus as well as in the filter's bandwidth and f_{peak} .

The third important aspect of the model is that direction discrimination is a local process, although the final output of the decision process is determined by an equally weighted global function. Little significance is attached to the absolute d_{\max} values produced by the model: there are various ways of increasing or decreasing these, such as setting another criterion value for d_{\max} or using off-axis-orientation filters.¹⁵ Instead, attention is directed toward the *relative* values of d_{\max} , which are robust to the effects of parameters, such as those mentioned above, that linearly scale d_{\max} at all frequencies by a constant factor.

B. Modeling d_{\max} for Low-Pass-Filtered Random-Dot Kinematograms

Cleary and Braddick¹⁹ found that low-pass filtering RDK's led to increasingly large values of d_{\max} when the high-frequency cutoff (f_h) was reduced below some critical value. In order to test whether such data can be accounted for by considering the signal from a single motion-sensitive filter, the algorithm described above was run with Cleary and Braddick's stimuli. Before the results of the modeling are given, it is important to describe one further finding of Cleary and Braddick's study. It has already been mentioned that when patch size is increased, there is a corresponding increase in d_{\max} . Cleary and Braddick observed that the high-frequency cutoff at which d_{\max} began to increase was increasingly lower for larger patches; i.e., increasing the patch size meant that the blurring had to be increasingly severe before any effect on d_{\max} was registered. Following Baker and Braddick,²³ they suggested that this was because the filters in the visual system tend to be increasingly coarse moving away from the fovea.

The model was run with a range of filters, each with a different f_{peak} . The ratio of the two Gaussians in the spatial domain was fixed at 1.5, such that the bandwidth of each filter was a constant 1.8 octaves. The results of both Cleary and Braddick's study and the modeling are shown in Fig. 7. It can be seen that modeling with a single filter captures the three trends in the human data well. First, for both the human and the model data, d_{\max} is affected only by relatively severe low-pass filtering. This supports Cleary and Braddick's argument that mild filtering does not affect d_{\max} , because frequencies above the critical value of f_h are not passed with a gain sufficient to alter the filter response significantly. Second, the critical value of f_h at which d_{\max} starts to rise varies as a function of patch size for the human data and filter size in the model. This similarity supports Cleary and Braddick's proposal that increases in d_{\max} with patch size are due to the activation of coarser filters in peripheral retinal regions.

The critical third trend that the model captures is the steep increase in d_{\max} as f_h is reduced past its critical value. For the model this is due to the relatively broad passband of the DOG filter, with the consequence that increases in the severity of the filtering continue to affect its response over a large frequency range. This shows that it is not necessary to propose that d_{\max} at high and low values of f_h are determined by high- and low-frequency filters, respectively. Rather, the results of the modeling support the hypothesis that, for a given stimulus, d_{\max} is determined by the response of the single lowest-frequency filter activated.

C. Estimating the Filter Bandwidth: d_{\max} for the Fractal-Noise versus the Random-Dot Kinematogram

The results from experiment 2 and the modeling data presented above offer strong support for the hypothesis that the upper spatial limit for motion detection is based on the response of the lowest-frequency channel activated. Using the data obtained in experiment 1, one can estimate the bandwidth of this filter independently of its peak spatial-frequency tuning. The basis for this is that the amplitude spectra for both a FNP and a RDP tend to fol-

low simple power-law functions, f^{-1} and f^0 , respectively (within limits set by the display size and pixel size). On logarithmic coordinates the amplitude spectrum of a RDP will fall close to a line of zero slope, whereas that of the FNP will tend to a slope of -1 . What will happen if the Fourier transforms of these patterns are passed through two filters, centered on different frequencies but with the same form and logarithmic bandwidth? For the RDP the two spectra will be related simply by a shift along the frequency axis, and their shapes will match that of the filter itself. For the FNP the two spectra will be the same up to a shift plus a linear scaling of amplitude, although now the shapes will not match that of the filter. However, this amplitude scaling does not affect the model, as it is only the relative strengths of the components that determine the location of the zero crossings. Thus, for both a RDP and a FNP, the effective consequence of applying a filter with a different f_{peak} is to shift the spectrum by the same distance along the frequency axis. In turn, this means that any effect on d_{\max} of changing the filter f_{peak} will be the same for the RDK and the FNK. Thus, whatever the filter f_{peak} , the ratio of d_{\max} values for the FNK and the RDK will remain constant. Since the ratio of values for the human observers is known from experiment 1, one can estimate the bandwidth of the filter by varying it to fit the empirically determined ratio. It should be noted that this reasoning relies on the assumption that

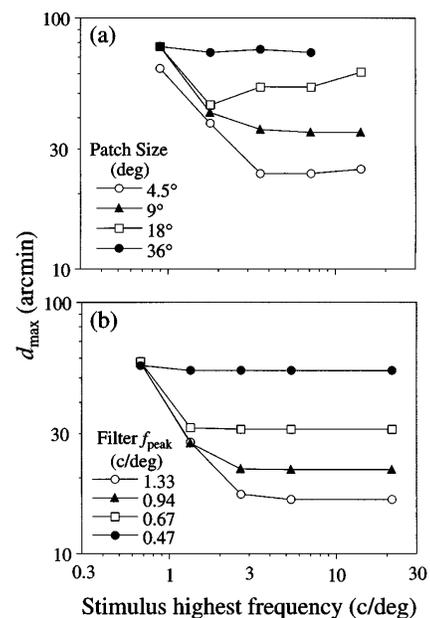


Fig. 7. (a) Data redrawn from Fig. 4 of Cleary and Braddick¹⁹ (two subjects pooled), where d_{\max} was measured for a range of low-pass-filtered RDK's at four patch sizes. Two effects are clear: (1) as the patch size increases so too does d_{\max} ; and (2) as the stimulus filtering removes increasingly lower frequencies, d_{\max} is initially unaffected; but then, at some critical value of f_1 (lower for larger patches), further filtering leads to large increases in d_{\max} . (b) Data from the model described above when it is presented with Cleary and Braddick's stimuli. Now, rather than patch size being varied, the filter f_{peak} is varied. Note that using a single filter captures the regions where d_{\max} is affected and unaffected by the stimulus filtering. Furthermore, decreasing the filter f_{peak} in the model has the same effect as increasing the patch size for the human data. These results offer strong support for the single-filter model.

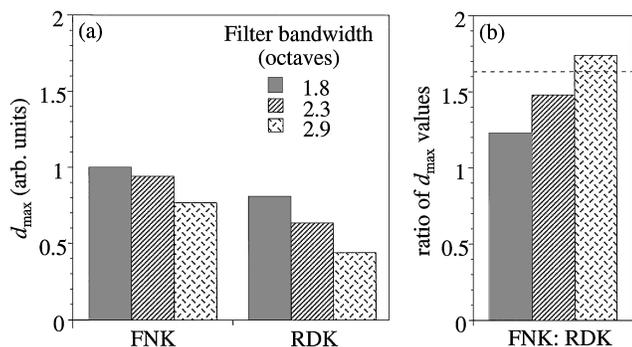


Fig. 8. Modeling data for the RDK and FNK used in experiment 1. (a) d_{\max} for three DOG filters with variable bandwidths but a fixed f_{peak} . The filter f_{peak} was set arbitrarily at 1.33 c/deg, and the absolute values of d_{\max} have no significance. This is made explicit if we normalize d_{\max} to 1.0 for the 1.8-octave filter applied to the FNK and linearly scale the other values accordingly. (b) The ratios of each pair of d_{\max} values shown in (a). The dashed line shows the ratio of the d_{\max} values in experiment 1. From this, the best-fitting filter bandwidth can be estimated as being between 2.3 and 2.9 octaves.

the values of d_{\max} for the FNK and the RDK obtained in experiment 1 were based on the responses of the same filter.

In general, it can be expected that as the bandwidth is increased, the ratio of d_{\max} for the FNK to d_{\max} for the RDK will also increase. In the limit, as the bandwidth approaches zero, the filter will not be sensitive to the spectral differences between the two stimuli, with the consequence that d_{\max} will be identical for the two kinematograms. As the bandwidth is increased, the differences in the images' spectra will become increasingly apparent in the filter response. In particular, the centroid of the filter response to the RDK will move toward higher frequencies relative to the centroid of the filter response to the FNK. In turn, this will have the consequence that the ratio of d_{\max} values for the FNK to the RDK will increase. By use of Eq. (1), the ratio of $\sigma_c : \sigma_s$ was set to 1.5, 3.0, or 6.0, giving half-height, full-width bandwidths of 1.8, 2.3, or 2.9 octaves, respectively. f_{peak} was held constant for all cases at 1.33 c/deg, although, as argued above, this value was arbitrarily chosen and has no effect on the estimate of the bandwidth. Figure 8 shows the relative values of d_{\max} for the three filters applied to the two kinematograms as well as the ratios of these values.

Figure 8(a) shows that as the bandwidth of the filter is increased, d_{\max} decreases for both the FNK and the RDK. However, the rate at which d_{\max} decreases is higher for the RDK. This is shown in Fig. 8(b), which illustrates the ratios of the d_{\max} values for each filter size. The dashed line shows the empirically determined d_{\max} ratio for the two kinematograms obtained in experiment 1 (1.63). The plot shows that a DOG filter with a bandwidth of ~ 2.6 octaves would provide a good fit to the data. Although this bandwidth is quite large, it must be remembered that the results from experiment 2 suggest that the filter used by the motion system for determining d_{\max} in broadband patterns is tuned to low frequencies (see also modeling below). Other estimates of filter tuning characteristics obtained psychophysically for form perception³¹

and motion detection,^{4,32} as well as estimates from physiological studies of primate cortical cells,³³ all indicate that the bandwidth increases to more than 2.0 octaves as the peak tuning moves to frequencies below 1 c/deg.

D. Estimating the Filter f_{peak} : d_{\max} for the Five-Octave versus the One-Octave Stimuli

In experiment 2 it was found that d_{\max} for the lowest-frequency one-octave stimulus was 1.46 times that of d_{\max} for the five-octave stimulus ($f_1 = 0.33$ c/deg for both stimuli). This finding was used to model the filter's tuning with a range of filters, each with a fixed bandwidth of 2.6 octaves, and the filter f_{peak} was varied to indicate which produced the best fit to this ratio. Figure 9 illustrates the findings for four filter sizes, spaced in half-octave steps.

Figure 9(a) shows that varying the filter f_{peak} over a range of 1.5 octaves has little effect on d_{\max} for the one-octave stimulus. This is because the filter bandwidth is much broader than the stimulus so that the spectrum of the filtered stimulus is determined primarily by the stimulus in all cases. In contrast, d_{\max} for the five-octave kinematogram can be seen to be highly dependent on the filter f_{peak} , decreasing as the filter is moved to higher frequencies. Figure 9(b) shows the ratios of these d_{\max} values. When the filter is centered half an octave below f_1 , the ratio of the two d_{\max} values is close to 1, but when the filter is centered one octave above f_1 , d_{\max} is more than twice as large for the single-octave pattern. A ratio of 1.0 indicates that the filter is not passing the higher frequencies in the five-octave stimulus at significant enough strength to affect the filtered output. The dashed line indicates the 1.46 ratio of the d_{\max} values obtained for the three human observers. It can be seen that the best-fitting f_{peak} to this value is near 0.33–0.47 c/deg.

E. Estimating the Filter f_{peak} : d_{\max} for the FNK versus the Five-Octave-Band FNK

One limitation of using bandpass-filtered stimuli to estimate the size of the filter that determines d_{\max} for broad-

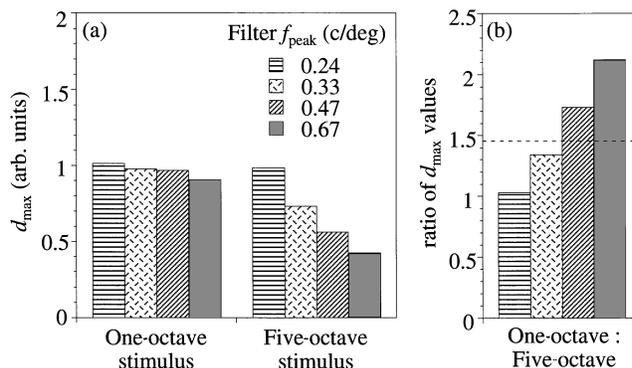


Fig. 9. Modeling data for the one- and five-octave stimuli used in experiment 2. (a) d_{\max} for four DOG filters with a variable f_{peak} but a fixed bandwidth of 2.6 octaves. Again, the absolute values of d_{\max} are not significant (see Section 5A), and so d_{\max} has been set to 1.0 for the one-octave stimulus filtered with the DOG whose $f_{\text{peak}} = 0.24$ c/deg, with the other values appropriately scaled. (b) The ratio of the d_{\max} values shown in (a). The dashed line shows the ratio of the d_{\max} values in experiment 2. It can be seen that the filter f_{peak} that best fits this ratio lies between 0.33 and 0.47 c/deg.

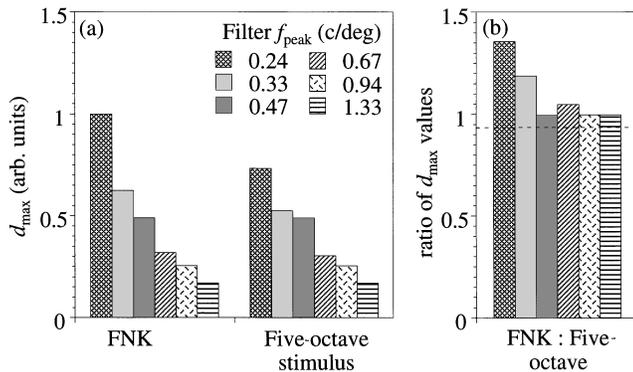


Fig. 10. Modeling data for the FNK used in experiment 1 and the five-octave version of the same stimulus used in experiment 2. (a) d_{\max} for six filters with a variable f_{peak} but a fixed bandwidth of 2.6 octaves. d_{\max} has been normalized to 1.0 for the condition in which the 0.24-c/deg filter was applied to the one-octave stimulus and linearly scaled for the other conditions. (b) The ratio of the d_{\max} values shown in (a). The dashed line shows the ratio of the two d_{\max} values for the human observers. Low-frequency filters produce a ratio above 1.0, because the additional low-frequency components in the FNK can be harnessed. As the filter f_{peak} is increased, these additional components are passed with an increasingly reduced gain until, at an f_{peak} of ~ 0.47 c/deg, the filter outputs to the two stimuli become indistinguishable. From this point on, the ratio of d_{\max} values is close to 1.0

band stimuli is that a lower-frequency filter may exist but simply not be activated by such patterns. It is possible to obtain an independent estimate of the lower limit on the filter f_{peak} by comparing d_{\max} for the FNK used in experiment 1 and the five-octave-bandpass version of the same stimulus used in experiment 2. The ratio of d_{\max} for the FNK to d_{\max} for the five-octave stimulus was 0.93 (mean of five subjects). The fact that this ratio is close to 1.0 suggests that the filter is not able to harness the extra octave of low frequencies in the FNK to increase d_{\max} relative to the five-octave kinematogram.³⁴ Rather, it suggests that the filter is centered within the five-octave band of the latter stimulus, where the responses of the filter to the two stimuli would be more similar and the ratio of 1.0 would be expected. To determine the lower limit on the filter f_{peak} , the model was applied to each stimulus with use of a range of filter f_{peak} values, again with the bandwidth kept constant at 2.6 octave. Figure 10 plots the results for six filters separated by half-octave steps.

Figure 10(a) illustrates that as the filter f_{peak} is increased, d_{\max} decreases for both stimuli. Figure 10(b) plots the ratios of these d_{\max} values. It can be seen that for filters whose $f_{\text{peak}} \geq 0.47$ c/deg, the ratio of d_{\max} for the FNK to d_{\max} for the five-octave stimulus is approximately 1.0. For coarser filters this ratio begins to rise as the filter passes the lower frequencies in the FNK. Because the ratio of the d_{\max} values for the human observers was close to 1.0 (dashed line) these data suggest that the filter f_{peak} is no lower than ~ 0.47 c/deg

6. GENERAL DISCUSSION

The experimental results and computer simulations reported in this paper suggest that the d_{\max} for broadband

kinematograms is determined by a single spatially band-pass filter. Cleary and Braddick¹⁹ found that low-pass filtering a RDK increased d_{\max} and that as patch size was increased, the stimulus filtering had to become increasingly severe before this improvement in d_{\max} was registered. It has been shown in the present paper that although a single filter can capture this whole range of behavior for any one patch size, the best-fitting filter f_{peak} decreases with increasing patch size, consistent with Cleary and Braddick's¹⁹ own explanation that coarser filters operate in eccentric vision.

Cleary and Braddick's hypothesis that the highest-frequency channel limits d_{\max} led them to argue that for a patch size of 4.5 arcdeg, the filter used by the motion system passes no frequencies beyond 3.56 c/deg. In Section 1 it was noted that the same investigators¹³ also showed, under similar experimental conditions, that observers had no difficulty in discriminating the direction of motion for bandpass RDK's when the lowest frequency present in the stimulus was 7.1 c/deg. These two findings can be reconciled if it is assumed that a range of motion filters exists and that it is the *lowest*- rather than the *highest*-frequency channel activated by the stimulus that determines d_{\max} . This suggestion is supported by the finding, in experiment 2, that d_{\max} for a five-octave stimulus was only 1.46 lower than d_{\max} for a one-octave stimulus with the same value of f_1 . From modeling, it was found that the single filter that best fitted this factor of 1.46 difference has an f_{peak} in the vicinity of 0.33–0.47 c/deg.

The fact that d_{\max} was similar for the broadband FNK used in experiment 1 and the five-octave kinematogram used in experiment 2 allowed an independent measure of the low-frequency limit to f_{peak} to be made. Consistent with the above estimate, the results of modeling this finding suggest that the low-frequency limit to f_{peak} was ~ 0.47 c/deg.

A third estimate of the filter f_{peak} can be gleaned from the modeling of Cleary and Braddick's data for low-pass-filtered kinematograms¹⁹ by matching the critical values of f_h at which d_{\max} for the human observers and the model start to rise. For the patch size used in the present experiments (7.5×6.0 arcdeg), the result of this comparison suggests that the filter f_{peak} is ~ 0.94 c/deg (see Fig. 7). The fact that this estimate is 0.5–1.5 octaves higher than the estimates derived from modeling the stimuli used in the present study is probably due to the differences in the bandwidths of the filters used in the two cases: 1.8 octaves and 2.6 octaves. Using a broader-band filter to model Cleary and Braddick's data would mean that the critical high-frequency cutoff at which d_{\max} started to rise would move to higher frequencies. In turn, this would mean that the estimate of the best-fitting filter f_{peak} would decrease, bringing it into line with the estimates of f_{peak} gleaned from modeling the present data.

It is possible that a modified version of Cleary and Braddick's¹⁹ hypothesis could also account for the data reported in this paper. This revised argument would be that rather than the highest-frequency channel activated determining d_{\max} in all cases, the amount of high-frequency masking is affected by the amount of stimulus energy at different frequencies. Thus it could be that d_{\max} is smaller for a RDK than for a FNK be-

cause there is more energy at high frequencies in the former stimulus, which produces greater masking. In this scheme, d_{\max} is determined by the spread of activity across a range of filters, with the weighting skewed toward the higher-frequency channels for a RDK. For a FNK it may be that there is no masking, which would explain why d_{\max} is so similar to the lowest-frequency single-octave stimulus. Although this account is not refuted by the results described in this paper, until a quantitative model of this masking is offered, it will be difficult to provide conclusive evidence either way. In contrast, the hypothesis presented in the present paper incorporates a quantitative model of direction discrimination at large displacements that can be tested with a wide range of stimuli. In addition, it is more parsimonious than the revised Cleary and Braddick hypothesis described above: all the data reported here and in Ref. 19 can be explained by considering the output from only a single channel, without the need to consider complex between-channel interactions.

Several authors have shown that d_{\max} can be affected by the dot size of RDKs. Changes in dot size below $\sim 10\text{--}15$ arcmin have no effect on d_{\max} , but any further increases lead to proportional changes in d_{\max} .^{10,17,35–37} Morgan¹⁷ has shown that with a patch size of 5.0 arcdeg, d_{\max} can rise to as large as 90 arcmin for a dot size to 72 arcmin—significantly larger than d_{\max} for the FNK used in experiment 1 (52 arcmin). How can these findings be reconciled with the hypothesis that stimuli with a $1/f$ amplitude spectrum yield a larger d_{\max} than do stimuli with a flat spectrum? The answer lies in the fact that the spectrum of a RDP is statistically flat only for components whose wavelengths are smaller than the patch size and larger than twice the dot width. For higher frequencies outside this band, the energy declines steeply. As dot size is increased, this band becomes increasingly low pass, and the energy of components within it increases. In other words, as dot size increases, the pattern spectrum ultimately becomes even more skewed toward low frequencies than is a $1/f$ amplitude spectrum. Thus the hypothesis that d_{\max} is greater for broadband stimuli whose spectra are dominated by low frequencies is not challenged by these findings.

The dependence of d_{\max} on dot size has been modeled by Morgan.¹⁷ The model is similar to the one proposed here, in that d_{\max} is based on the spacings of zero crossings in the output of a bandpass (Laplacian of a Gaussian) filter. Morgan found that the filter f_{peak} that best fitted his dot size data was between 0.85 and 1.7 c/deg ($\sigma = 8\text{--}16$ arcmin). The likely explanation of why his estimated value of the filter f_{peak} is higher than the estimate reported here again hinges on the fact that the bandwidth of a Laplacian of a Gaussian filter (1.8 octaves) is narrower than the 2.6-octave filter used in the present model. Due to the fact that increasing dot size has an effect on the spectrum similar to that of low-pass filtering, the arguments presented for that case above apply here. Thus using a broader-band filter to model Morgan's data should lead to a lower-frequency estimate of the best-fitting filter f_{peak} . Another factor that may have contributed to the difference in these estimates is that Morgan's patch size was 5.0×5.0 arcdeg compared with the larger patch (7.5×6.0 arcdeg) used

in the present study. As shown in Fig. 7, it is likely that coarser filters are activated by more-eccentric stimuli.

7. SUMMARY

This paper has reported the results of two experiments measuring d_{\max} with a range of broadband and narrow-band two-frame kinematograms. A quantitative model has been developed, in which direction discrimination is based on matching elements across the two outputs of a single-bandpass filter. Taking both the current set of results and those from other d_{\max} experiments into account, the most parsimonious explanation of human performance is that d_{\max} is dependent on the lowest-frequency filter activated by the stimulus. As argued in Section 1, this is also the most computationally advantageous strategy. Thus it appears that the visual system has access to the output of this channel while ignoring the incoherent motion signaled from higher-frequency channels.

ACKNOWLEDGMENTS

This work was funded by a Science and Engineering Research Council (UK) postgraduate studentship and also a Wellcome Trust Vision Training Fellowship awarded to the author, under the sponsorship of B. J. Rogers. Parts of this work were presented in a paper given in Pisa, Italy, at the European Conference on Visual Perception, August 30–September 3, 1992. Thanks are due to Mark Bradshaw and Andrew Glennerster for useful discussions and help in preparing the manuscript.

The author's telephone number is 44-1865-271309; e-mail address, richard@psy.ox.ac.uk.

REFERENCES AND NOTES

1. R. L. De Valois and K. De Valois, *Spatial Vision* (Oxford U. Press, 1988).
2. M. J. Keck, F. W. Montague, and T. P. Burke, "Influence of the spatial periodicity of moving gratings on motion responses," *Invest. Ophthalmol. Vis. Sci.* **19**, 1364–1370 (1980).
3. E. L. Cameron, C. L. Baker, and J. C. Boulton, "Spatial frequency selective mechanisms underlying the motion aftereffect," *Vision Res.* **32**, 561–568 (1992).
4. S. J. Anderson and D. C. Burr, "Receptive field properties of human motion detector units inferred from spatial frequency masking," *Vision Res.* **29**, 1342–1358 (1989).
5. C. L. Baker and M. S. Cynader, "Spatial receptive field properties of direction selective neurons in cat striate cortex," *J. Neurophys.* **55**, 1136–1152 (1986).
6. E. H. Adelson and J. R. Bergen, "Spatiotemporal energy models for the perception of motion," *J. Opt. Soc. Am. A* **2**, 284–299 (1985).
7. J. P. H. van Santen and G. Sperling, "Elaborated Reichardt detectors," *J. Opt. Soc. Am. A* **2**, 300–321 (1985).
8. D. Marr and S. Ullman, "Directional selectivity and its use in early visual processing," *Proc. R. Soc. London B* **211**, 151–180 (1981).
9. B. Julesz, *Foundations of Cyclopean Perception* (U. Chicago Press, Chicago, Ill., 1971).
10. O. J. Braddick, "A short-range process in apparent motion," *Vision Res.* **14**, 519–527 (1974).
11. J. J. Chang and B. Julesz, "Cooperative and non-cooperative processes of apparent movement of random-dot cinematograms," *Spatial Vis.* **1**, 39–45 (1985).
12. B. De Bruyn and G. A. Orban, "Discrimination of oppo-

- site directions measured with stroboscopically illuminated random-dot patterns," *J. Opt. Soc. Am. A* **6**, 323–328 (1989).
13. R. Cleary and O. J. Braddick, "Direction discrimination for band-pass filtered random dot kinematograms," *Vision Res.* **30**, 303–316 (1990).
 14. W. F. Bischof and V. Di Lollo, "Perception of directional sampled motion in relation to displacement and spatial frequency: evidence for a unitary motion system," *Vision Res.* **30**, 1341–1362 (1990).
 15. W. F. Bischof and V. Di Lollo, "On the half-cycle displacement limit of sampled directional motion," *Vision Res.* **30**, 1341–1362 (1991).
 16. R. A. Eagle and B. J. Rogers, "Motion detection is limited by element density not spatial frequency," *Vision Res.* **36**, 545–558 (1996).
 17. M. J. Morgan, "Spatial filtering precedes motion detection," *Nature (London)* **355**, 344–346 (1992).
 18. J. J. Chang and B. Julesz, "Displacement limits for spatial frequency filtered random-dot cinematograms in apparent motion," *Vision Res.* **23**, 1379–1385 (1983).
 19. R. Cleary and O. J. Braddick, "Masking of low frequency information in short-range apparent motion," *Vision Res.* **30**, 317–327 (1990).
 20. In the experiment reported in Ref. 19, contrast was not normalized across stimuli, and so it decreased as the filtering became more severe.
 21. Sharpimage Software, *Heritable Image Processing System* (New York, 1984).
 22. M. J. Morgan and M. Fahle, "Effects of pattern element density upon displacement limits for motion detection in random binary luminance patterns," *Proc. R. Soc. London B* **248**, 189–198 (1992).
 23. C. L. Baker and O. J. Braddick, "The basis of area and dot number effects in random dot motion perception," *Vision Res.* **22**, 1253–1259 (1982).
 24. R. N. Bracewell, *The Fourier Transform and Its Applications* (McGraw-Hill International Editions, New York, 1986).
 25. N. Brady and D. J. Field, "What's constant in contrast constancy? The effects of scaling on perceived contrast of band-pass patterns," *Vision Res.* **35**, 739–756 (1995).
 26. O. J. Braddick and R. Cleary, "On the half-cycle displacement limit of sampled directional motion," *Vision Res.* **31**, 649–660 (1991).
 27. J. J. Chang and B. Julesz, "Cooperative phenomena in apparent movement perception of random-dot cinematograms," *Vision Res.* **24**, 1781–1788 (1984).
 28. S. J. Anderson and D. C. Burr, "Receptive field size of human motion detection units," *Vision Res.* **27**, 621–635 (1987).
 29. M. J. Hawken and A. J. Parker, "Spatial properties of neurons in monkey striate cortex," *Proc. R. Soc. London B* **231**, 251–288 (1991).
 30. Eagle and Rogers¹⁶ described a model in which the 2-D spacing of 2-D peaks was used in the matching process. In fact, both the 2-D spacing of 2-D peaks and the 1-D spacing of zero crossings are similarly affected by the application of different-sized isotropic spatial filters. For instance, if the filter f_{peak} is halved, both of these spacings double.
 31. H. R. Wilson, D. K. McFarlane, and G. C. Phillips, "Spatial frequency tuning of orientation selective units estimated by oblique masking," *Vision Res.* **23**, 873–882 (1983).
 32. S. J. Anderson and D. C. Burr, "Spatial and temporal selectivity of the human motion detection system," *Vision Res.* **25**, 1147–1154 (1985).
 33. R. L. De Valois, D. G. Albrecht, and L. G. Thorell, "Spatial frequency selectivity of cells in macaque visual cortex," *Vision Res.* **22**, 549–559 (1982).
 34. It is not clear why this value is slightly below 1.0. The FNK has both lower and higher frequencies than are present in the five-octave pattern, but it is extremely unlikely that the additional higher frequencies had any detrimental effect on performance with the FNK. Cleary and Braddick¹⁹ found that frequencies beyond 3.56 c/deg had no effect on d_{max} for a broadband stimulus, and the additional high frequencies present in the FNK are all beyond 10.33 c/deg.
 35. P. Cavanagh, J. Boeglin, and O. E. Favreau, "Perception of motion in equiluminous kinematograms," *Perception* **14**, 151–162 (1985).
 36. T. Sato, "Effects of dot size and dot density on motion perception with random dot kinematograms," *Perception* **19**, 329 (1990).
 37. This figure of 10–15 arcmin is probably dependent on the patch size. The effects on the image energy spectrum due to increasing dot size and low-pass filtering are similar. So, as Cleary and Braddick¹⁹ have found for low-pass-filtered RDK's (see Fig. 7), increases in patch size would be expected to lead to increases in the critical dot size at which d_{max} begins to change.