Pergamon

0042-6989(94)00301-7

Vision Res. Vol. 35, No. 15, pp. 2117–2125, 1995

0042-6989/95 $9.50 + 0.00

Motion Detection in Interleaved Random Dot Patterns: Evidence for a Rectifying Nonlinearity Preceding Motion Analysis

GEORGE MATHER,*† HILARY TUNLEY*

Received 13 May 1994; in revised form 17 November 1994

Three experiments examined direction discrimination in temporally interleaved random dot patterns. The stimulus consisted of two or more uncorrelated random patterns presented in a repeating temporal sequence, so that each pattern appeared only once every *n* frames, separated by uncorrelated patterns. Each pattern shifted either leftward or rightward at each re-appearance (all patterns shifted in the same direction in any one presentation). Subjects could specify shift direction correctly even when eight different patterns were interleaved, provided that the duration of each frame was brief. An explanation based on responses in first-order motion energy detectors tuned to low spatiotemporal frequencies (effectively summating the interleaved patterns over time) was tested using a stimulus in which each pattern inverted in contrast mid-way through each frame. Contrary to predictions based on temporal summation, performance with contrast-inverting patterns was only slightly lower than with non-inverting patterns. An alternative explanation was examined, based on responses in motion detectors that full-wave rectify image contrast before extracting motion energy. Computed responses from such detectors successfully predicted psychophysical performance with interleaved random patterns. Implications for models of motion analysis are discussed.

First-order motion   Second-order motion   Rectification   Random dot kinematograms

INTRODUCTION

Discontinuously moving random dot patterns have been used extensively to study human motion perception. The spatial arrangement of dots in each static view or *frame* is identical, except for the addition of a fixed frame-to-frame spatial displacement in one direction or its opposite. The observer's task is to discriminate between the two possible directions. Early studies cast the problem facing the observer in terms of the *correspondence problem*: how to match each dot in each frame with its corresponding partner in the next frame. In suitable experimental conditions the stimulus does evoke an impression of movement, and observers correctly identify the direction of the spatial displacement. Recent results suggest that it is not sufficient to consider random dot patterns in terms of the fate of individual dots. It may be more enlightening to consider the interaction between the spatiotemporal frequency content of the pattern and the filtering properties of the visual system (Mather & Tunley, 1993, 1995). The spatiotemporal frequency spectrum of discontinuously moving stimuli contains frequency components introduced by sampling (alias signals) at spatial and temporal frequencies determined by the spatial and temporal sampling rates used to generate the stimulus (Watson, Ahumada & Farrell, 1986). Performance limits seem to be governed by the proximity of this alias energy to the signal energy, since low-pass spatial or temporal filtering improves performance at lower sampling rates. We can infer that discrimination is mediated by visual processes whose response faithfully reflects the spatiotemporal Fourier content of the stimulus (so-called "Fourier-based" or "first-order" motion detectors).

In this paper we introduce a new random dot stimulus, and ask whether its psychophysical properties support the theoretical framework outlined above. In the stimuli described earlier, dots in immediately succeeding frames are always correlated. In the new stimulus two uncorrelated dot patterns are temporally interleaved, so that odd-numbered frames contain one set of dots, and even-numbered frames contain a different set of dots, with each dot field undergoing a fixed small displacement at each presentation (see Fig. 1). Pilot observations indicated that observers can correctly identify the direction of interleaved dot patterns. A straightforward explanation is that first-order detectors tuned to low spatiotemporal frequencies effectively summate the interleaved patterns over time, removing the problem of establishing correspondence across intervening patterns.

*Experimental Psychology, University of Sussex, Brighton BN1 9QG, England [*Email* georgem@epvax.sussex.ac.uk].
†To whom all correspondence should be addressed.

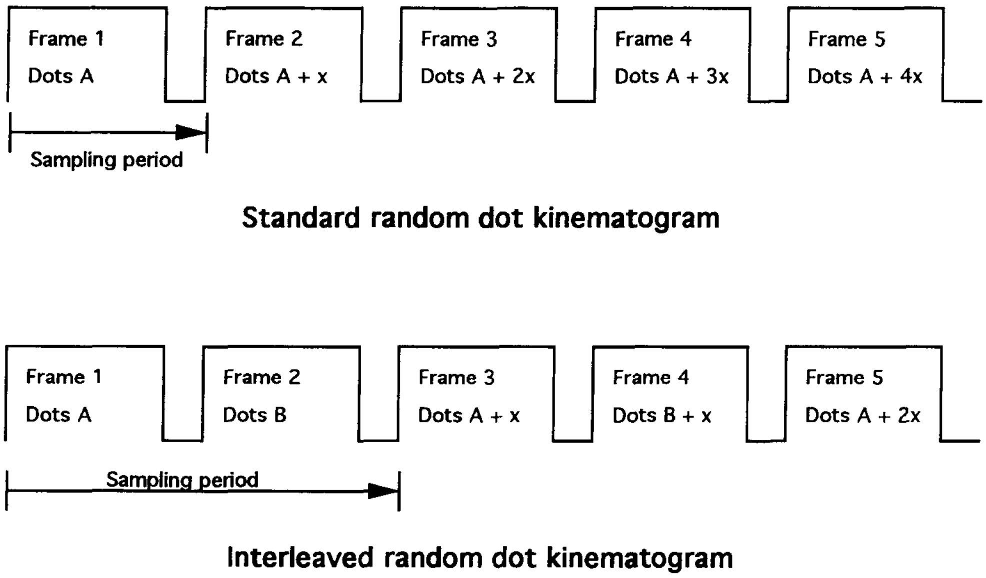

FIGURE 1. Schematic depiction of the experimental stimulus, representing the time-course of presentation. The upper row depicts a standard random dot kinematogram, in which the spatial arrangement of dots in each frame is identical except for the addition of a fixed frame-to-frame displacement (x). The lower row depicts an interleaved random dot kinematogram, in which successive frames contain uncorrelated dot patterns, but dots in odd-numbered frames are correlated (dots A), and dots in even numbered frames are correlated (dots B). At each re-appearance, a fixed spatial displacement (x) is added to each dot pattern. In the experimental stimulus, up to eight different patterns were interleaved in this way, and the full stimulus sequence contained five presentations of each pattern.

Experiment 1 presents parametric data on performance, as a function of the number of interleaved patterns, and the time period for which each pattern is visible. Experiment 2 tested the temporal summation account.

EXPERIMENT 1

Method

Subjects. Five observers participated, both authors and three others who were naive as to the purpose of the experiment.

Apparatus and stimuli. Stimuli were generated by a PC-compatible computer equipped with a high-resolution raster-graphics sub-system, and displayed on a Hitachi 14MVX monitor (P22 phosphor) at a frame rate of 83 Hz (non-interlaced). The stimulus display consisted of a 256 × 256 array of 50% random black–white elements (16 and 147 cd m^{-2}) against a uniform grey background (82 cd m^{-2}). Each dot subtended 0.94 arc min, so the stimulus array subtended 4 × 4 arc deg at the 114 cm viewing distance. Motion displays of interleaved patterns were created, as depicted in Fig. 1. In different presentations either 1, 2, 4, or 8 uncorrelated patterns were temporally interleaved. Each pattern shifted in the same direction by the same amount (one dot width) in each frame, but shift direction varied randomly from trial to trial. The duration of each frame in the interleaved sequence was also manipulated. Frame duration in each presentation was selected from one of four values: 12, 24, 48, or 96 msec. In addition to the 16 possible experimental stimuli (all combinations of four interleaving levels and four frame durations), a corresponding set of 16 control stimuli was also created, identical to the experimental stimuli except that only one of the interleaved patterns was visible. Dots in all other patterns in the sequence were set to mean luminance so, for example, the control stimulus for a two-pattern interleaved stimulus consisted of a single visible pattern presented with a blank inter-stimulus interval equal to one pattern duration (of course experimental stimuli containing one pattern were identical to the corresponding control stimuli, since there were no intervening patterns). Control stimuli were included to assess the effect of uncorrelated patterns intervening between presentations of correlated patterns.

Procedure. The total of 32 stimuli was shown to each subject in random order 40 times over a number of experimental sessions. Each trial involved a motion sequence containing five frames of each of the interleaved patterns so, for example, a two-pattern trial contained 10 stimulus frames whereas an eight-pattern trial contained 40 stimulus frames. Stimulus direction varied randomly from trial to trial between rightward and leftward, and after each presentation the observer pressed one of two response keys to signify perceived direction. In between presentations the stimulus field was uniform at mean luminance. A small red fixation cross was visible continuously in the centre of the display, and subjects were instructed to stare at it.

Results

Figure 2 plots mean discrimination performance for the experimental stimuli as a function of the number of interleaved patterns, with frame duration as the parameter. Standard errors have been omitted for clarity, but were on average 3.77% (variance of SEs was 4.8%).

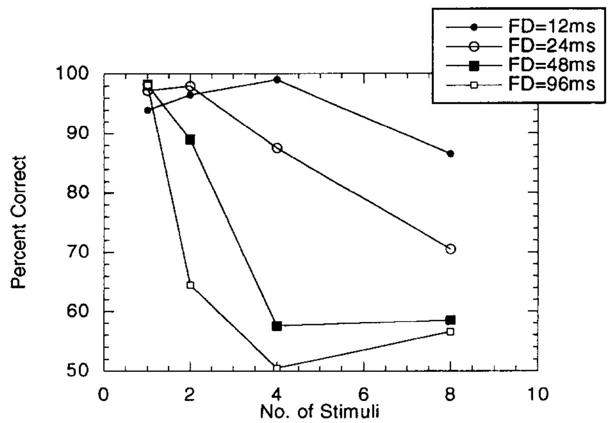

FIGURE 2. Results of Expt 1, showing mean percentage correct in a direction discrimination task as a function of the number of interleaved patterns in the stimulus, with frame duration as the parameter defining different curves.

At the shortest frame duration performance was near-perfect even when eight uncorrelated patterns were interleaved. As frame duration increased high levels of discrimination became confined to stimuli containing fewer interleaved patterns. The pattern of data suggested that the time interval between successive views of corresponding patterns was important, so Fig. 3 re-plots the data as a function of sampling period (time interval between the onset of corresponding patterns, traditionally known as stimulus onset asynchrony; see Fig. 1), with number of interleaved patterns as the parameter. This figure also plots results for the control stimuli containing only one visible pattern separated by blank intervals. Clearly the data now collapse onto a single function, indicating that sampling period is the crucial parameter. If we define the temporal limit of performance (T_{max}) as the sampling period yielding 80% correct discrimination, then T_{max} falls at about 100 msec. Since spatial displacement was fixed at 0.94 arc min per frame, performance limits can also be expressed in terms of velocity.

Discussion

The T_{max} value obtained from Expt 1 is typical of temporal limits reported using random dot patterns but cannot be treated directly as a reflection of motion detector properties, since detection limits result from an interaction between visual filter properties and pattern spatiotemporal frequency content. Discrimination in interleaved patterns could be based on responses in motion energy detectors tuned to low spatiotemporal frequencies, which effectively summate the patterns over time. As an illustration, Fig. 4 shows an *xt* plot (top left) of a stimulus containing four interleaved rows of random elements, which shift rightwards at each re-appearance. The spatiotemporal Fourier transform of this pattern contains signal energy (passing through the origin) which is tilted clockwise from vertical. Motion energy detectors with receptive fields tuned to rightwards motion at low spatiotemporal frequencies (i.e. near the origin in the top-right and bottom-left quadrants of the transform) would generate appropriate directional responses to this pattern. Note that the transform also contains replicas of the signal away from the origin, introduced by discrete spatial and temporal sampling. These replicates at inappropriate velocities move closer to the signal at longer sampling periods, and could account for the deterioration in performance evident in Fig. 3. Experiment 2 was designed as a direct test of the summation account.

EXPERIMENT 2

Three different stimuli were created, containing either:

(1) four interleaved random patterns at a fixed frame duration of 24 msec (i.e. two TV frames), corresponding to one of the stimuli in Expt 1;
(2) eight interleaved random patterns at a fixed frame duration of 12 msec (i.e. one TV frame, same sampling period as the first stimulus);
(3) four interleaved random patterns at a fixed frame duration of 24 msec, as in the first stimulus, except that half way through each view of each pattern (i.e. after one TV frame of the presentation), the pattern reversed in contrast polarity so that bright dots became dark and vice versa.

On the basis of results in Expt 1, it is straightforward to predict that performance in the four-pattern and eight-pattern non-inverting stimuli (1 and 2) will be near-perfect. In the case of the four-pattern inverting stimulus, the explanation offered above based on energy detector responses at low spatiotemporal frequencies predicts that performance will be at chance levels. Intuitively, we can expect that contrast inversion will defeat any attempt to generate motion signals by temporal averaging, since the average of a pattern and its negative image is uniform. The effect of contrast inversion on spatiotemporal frequency content is illustrated in Fig. 4. The bottom left *xt* plot shows a four-pattern interleaved

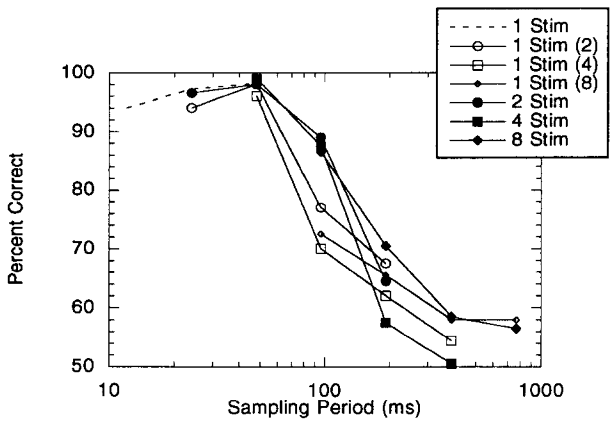

FIGURE 3. Results of Expt 1 re-plotted in terms of percentage correct as a function of sampling period (defined as in Fig. 1). The solid symbols identify stimuli containing interleaved patterns, and open symbols identify control stimuli containing only one pattern but with frame duration and sampling period matched to different interleaved stimuli. See text.

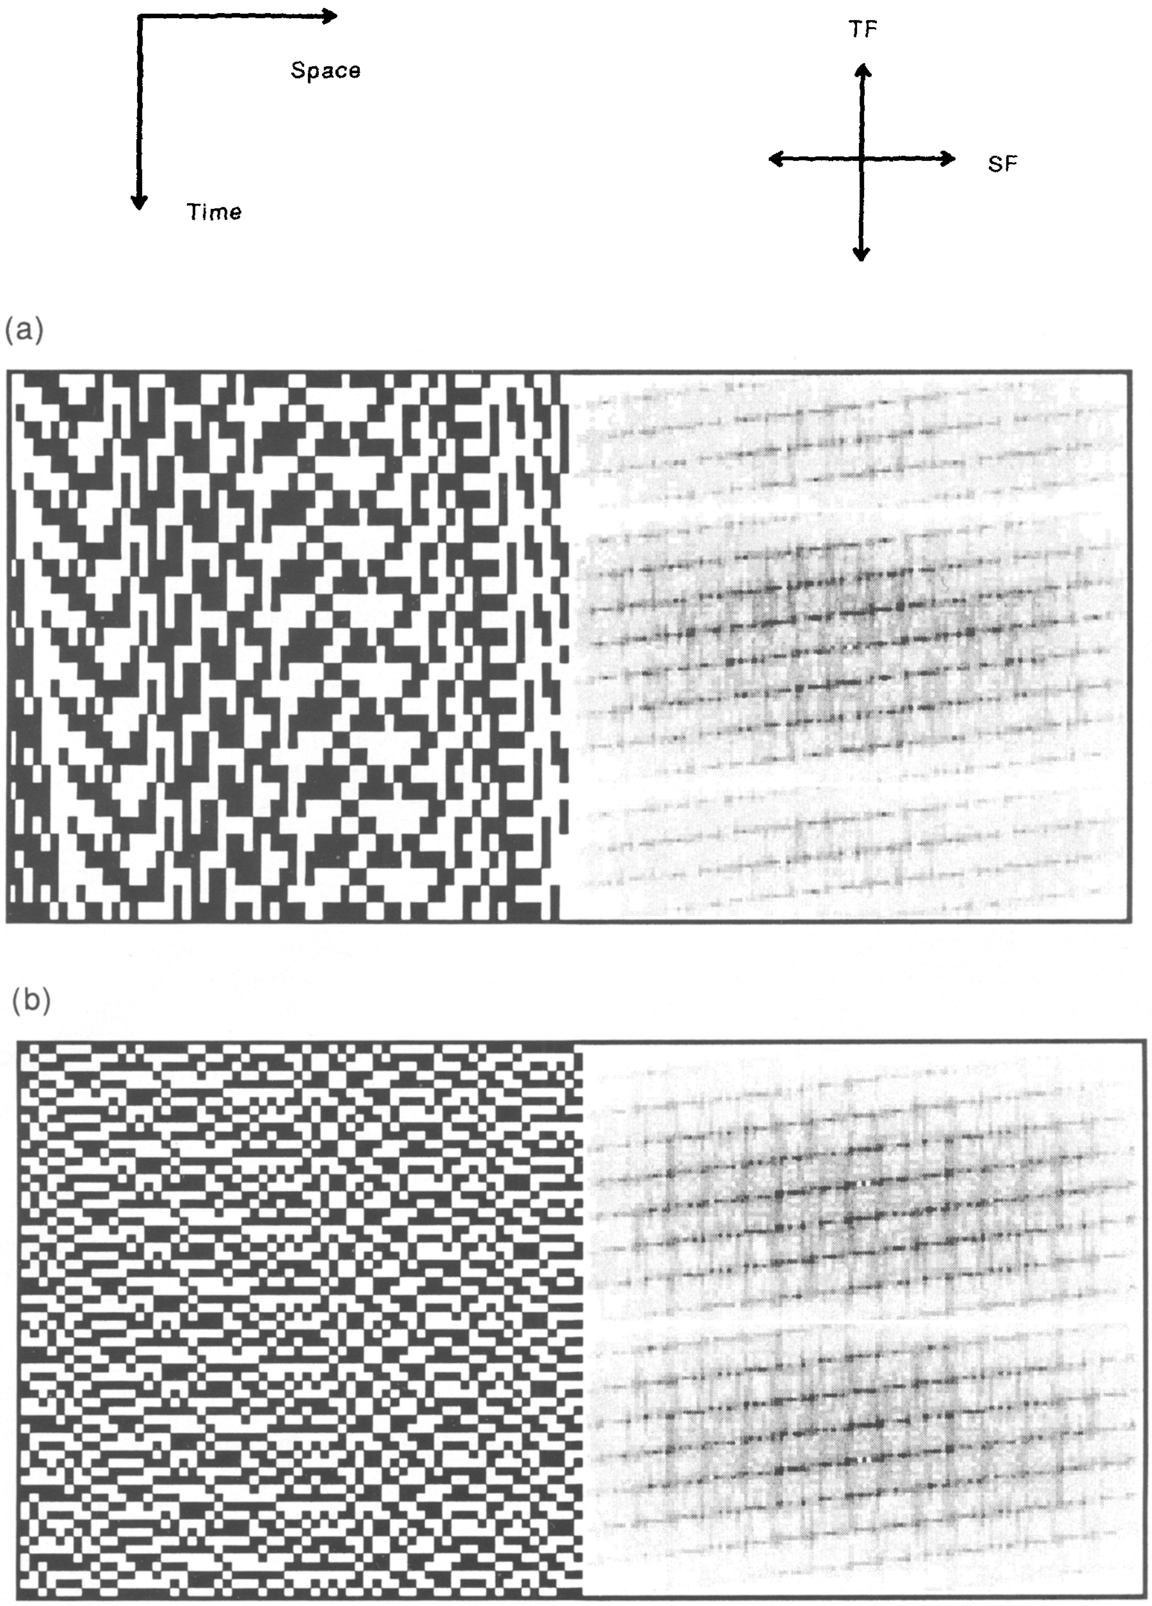

FIGURE 4. Space–time or *xt* plots of interleaved dot patterns, and their spatiotemporal Fourier transforms. In the left-hand plots space runs horizontally and time runs vertically. (a) The left plot shows four arrays of random black–white elements (64 elements in a row) interleaved over time, so that each pattern re-appears every four rows (time-frames), but is displaced rightward by one element width at each re-appearance (a total of 32 time-frames are shown, so each correlated array appears eight times). (b) The left plot is identical to the left plot of (a), except that each array of random elements reverses in contrast halfway through its time-frame. The right-hand plots show the spatiotemporal Fourier spectra of the adjacent *xt* plots. Each *xt* plot was held as a 128×128 pixel array, so each transform represents 64 spatial frequencies horizontally and 64 temporal frequencies vertically. Each pixel in the transform represents the Fourier amplitude at that frequency, $|F(u,v)|$, scaled conventionally as follows: $\log[1 + |F(u,v)|]$. Scaled amplitudes were quantized to 256 grey levels for display, with darker pixels representing higher amplitudes.

stimulus similar to that depicted in the top-left *xt* plot, except that each pattern reverses in contrast in the second half of each frame. The spatiotemporal Fourier transform of this pattern (bottom-right) reveals that contrast inversion neatly removes signal energy passing through the origin, but leaves intact alias energy away from the origin. The stimulus therefore offers no signal for Fourier-based detectors tuned to the spatiotemporal orientation of the pattern.

Method

Subjects. Five observers participated, both authors and three others naive as to the purpose of the experiment.

Apparatus and stimuli. The same equipment and stimuli were used as in Expt 1, except for variations necessary to define the three stimuli described above. In addition, the effect of displacement extent was examined, by generating versions of each stimulus at five displacements (5, 10, 15, 20 and 25 dots per frame).

Procedure. Forty trials were presented for each stimulus, in random order, over two experimental sessions. As before, each trial involved a motion sequence containing five frames of each of the interleaved patterns. Direction reversed randomly from trial to trial. After each presentation the subject pressed one of two keys to indicate perceived direction.

Results

Figure 5 plots mean discrimination performance as a function of displacement, for each stimulus condition. Standard errors have been omitted for clarity, but were on average 3.3% (variance 2.7%). Performance in the inverting condition was well above chance, and only slightly below performance in the other conditions.

Discussion

The results of Expt 2 are inconsistent with the proposal that performance is mediated by motion detectors

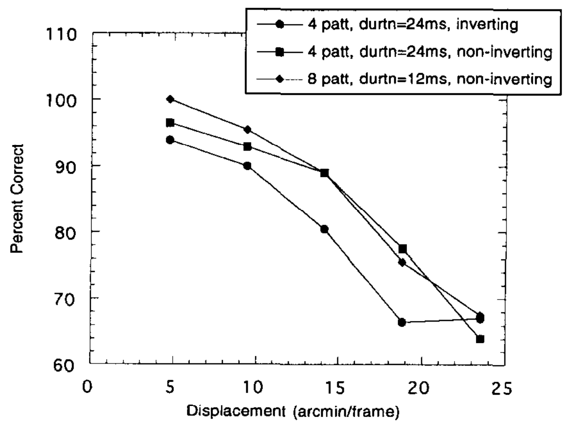

FIGURE 5. Results of Expt 2, showing mean percentage correct in a direction discrimination task as a function of frame-to-frame displacement. Different curves represent results from different stimuli, either: a four-pattern interleaved stimulus (■), an eight-pattern interleaved stimulus (◆), or a four-pattern interleaved stimulus in which each pattern inverted in contrast halfway through each frame (●).

tuned to low spatiotemporal frequencies, since contrast inversion had little effect. However, the fact that performance is limited by sampling period and displacement *is* consistent with a system operating on the basis of stimulus spatiotemporal frequency content. This apparent paradox can be removed by assuming that a non-linear transform is applied to the stimulus before Fourier-based motion analysis. Full-wave rectification should make the system immune to the effects of contrast inversion. As a quantitative test of this account, we applied the following sequence of computations to *xt* plots similar to those depicted in Fig. 4:

(1) *xt* "source" plots were generated corresponding to the three stimuli used in Expt 2. Each plot consisted of 64 random black–white elements along the *x* (horizontal) axis, in 64 time frames along the *t* (vertical) axis, as in Fig. 4 (left).
(2) A Gaussian spatiotemporal filter was applied to each source plot (space constant equal to the width of one dot and time constant equal to the duration of one-half of a time-frame). These values were chosen arbitrarily to provide a small degree of low-pass filtering, to simulate effects of image formation and early visual filtering.
(3) The filtered plots were full-wave rectified.
(4) The energy available for motion detection was estimated by taking the spatiotemporal Fourier transform of each rectified plot, and then computing directional power (DP), defined as the ratio of summed rightward power in the transform to summed leftward power. An absence of directional information would yield a DP ratio of 1. Rightward power would be indicated by DP values >1. and leftward power would be indicated by DP values <1.

A number of recent papers have used DP measures to estimate Fourier energy available for motion detection in simple visual displays (e.g. Dosher, Landy & Sperling, 1989; Nishida & Sato, 1992; Boulton & Baker, 1993; Mather & Tunley, 1993). Figure 6 illustrates the sequence of operations. We actually computed DP at three stages in the sequence. First, DP available in the source image, then DP in the low-pass filtered version of the source, and finally DP in the filtered and rectified image.

Figure 7 plots DP values computed for each of the stimuli used in Expt 2, along with data replotted from Fig. 5 (at the shortest displacement). Each DP value is the average obtained from six different random patterns. Turning first to DP in the source image, weak directional signals are available in the four-pattern and eight-pattern non-inverting stimuli, but no directional information is available in the four-pattern inverting stimulus, as expected given the absence of signal energy in Fig. 4. Low-pass filtering the source improves DP for the non-inverting patterns (presumably due to the removal of sampling artefacts), but introduces a reversed signal for the inverting pattern, i.e. DP values <1. This is the familiar reversed apparent motion effect caused by contrast reversal (e.g. Anstis & Rogers, 1975, 1986).

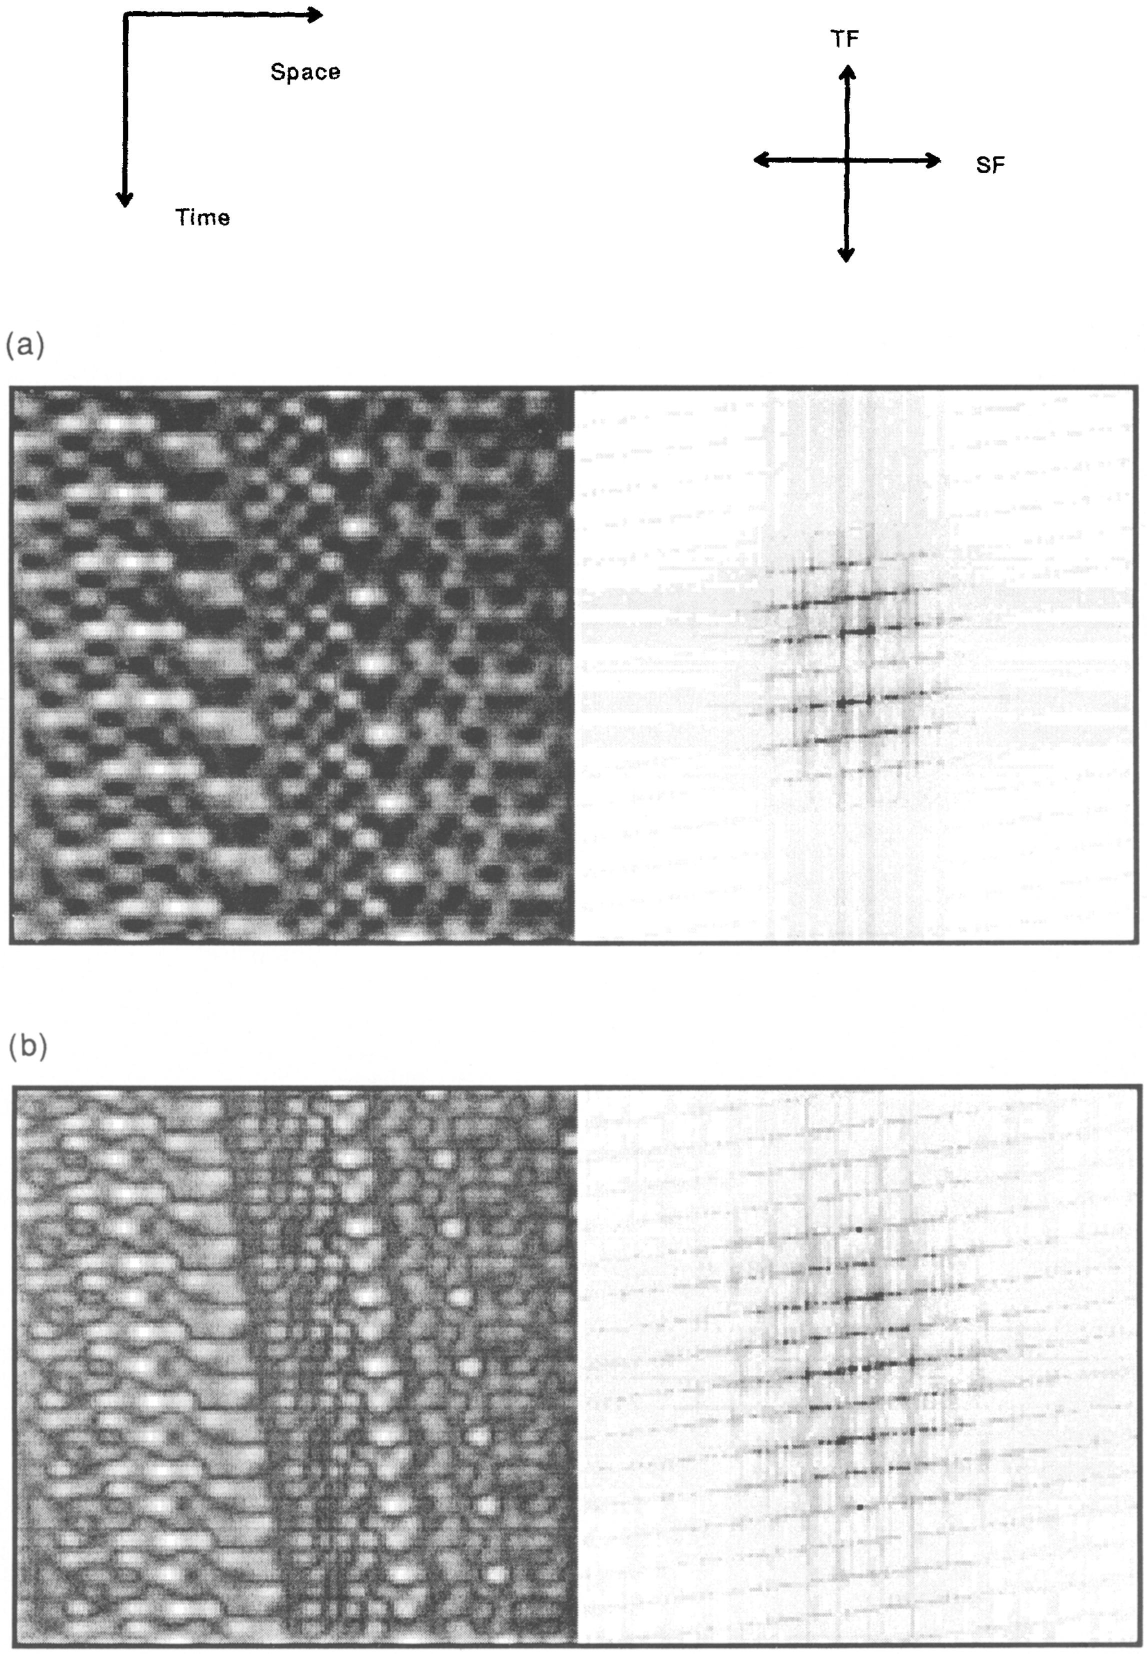

FIGURE 6. Computations required to reveal that motion of contrast-inverting interleaved patterns. (a) The left *xt* plot depicts a four-pattern inverting sequence similar to that show in Fig. 4, after the application of a spatiotemporal Gaussian filter (space constant equal to one dot width, and time constant equal to half a time-frame). (b) The left *xt* plot depicts the same pattern after full-wave rectification. Spatiotemporal Fourier transforms corresponding to the plots are shown on the right (conventions as in Fig. 4). Contrast-inverting patterns do not contain concentrations of Fourier energy which pass through the origin (upper transform), but do contain such energy concentrations after rectification (lower transform). Energy passing through the origin can be detected by energy-based motion detectors with receptive fields placed anti-symmetrically about the origin.

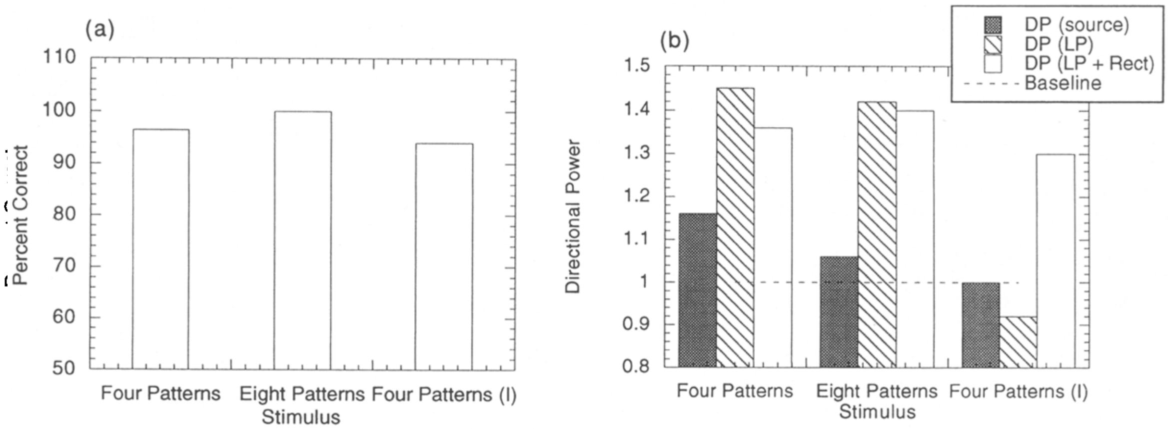

FIGURE 7. Comparison of data and predictions for Expt 2. (a) Re-plots data from Fig. 5 for each of the three patterns, at the shortest displacement. (b) Plots DP values computed from Fourier transforms similar to those in Figs 4 and 6. Grey bars give DP values based on untransformed images (e.g. lower transform in Fig. 4), hatched bars give DP values after images were low-pass filtered with a spatiotemporal Gaussian function (e.g. upper transform in Fig. 6), and plain bars give DP values for images after filtering and rectification (e.g. lower transform in Fig. 6). DP values >1 signify net rightward energy in the transform, and DP values <1 signify net leftward energy. All DP values were means computed from six different patterns (see text for details).

However, DP measures taken after filtering and rectification are well above unity for all three stimuli, in agreement with the psychophysical data. Of course, we do not know what DP value would be necessary for reliable direction discrimination in the experiment, but the computations indicate at least 30% more power in the correct direction than in the reverse direction for all three stimuli, a high signal-to-noise ratio.

In summary, computational modelling confirms that the contrast inverting pattern used in Expt 2 can support a directional response from Fourier-based detectors only when a rectifying non-linearity precedes energy analysis. The similar displacement functions for inverting and non-inverting patterns in Fig. 5 suggests a common underlying detection process. As a further test for consistency between data from inverting and non-inverting patterns, Expt 3 measured direction discrimination in contrast inverting patterns as a function of sampling period, for comparison with corresponding results of Expt 1.

EXPERIMENT 3

Method

Subjects. Four subjects participated, both authors and two naive observers.

Apparatus and stimuli. Equipment and stimulus generation techniques were identical to previous experiments. A four-pattern inverting stimulus was used, at a fixed displacement of one dot width per frame. Frame duration varied in different presentations between 24, 48, 72 and 96 msec. Pattern contrast inverted half way through each frame of each pattern. The four frame durations correspond to the following sampling periods, 96, 192, 384 and 768 msec.

Procedure. A total of 40 trials were presented at each frame duration, in random order, during a single experimental session. As before, each trial involved a motion sequence containing five frames of each of the four interleaved patterns. Direction reversed randomly from trial to trial. After each presentation the subject pressed one of two keys to indicate perceived direction.

Results

The solid symbols in Fig. 8(a) represent mean discrimination performance as a function of sampling period (± 1 SE). For comparison, the open symbols represent results for non-inverting patterns, replotted from Fig. 3. Data from the two experiments coincide almost exactly.

Discussion

The close similarity between results with inverting patterns and results with non-inverting patterns bears out the earlier suggestion that they are mediated by a common underlying mechanism. There is a hint of improved performance at the longest sampling periods [rightmost data point in Fig. 8(a)] perhaps reflecting the system's biphasic temporal impulse response (Shioiri & Cavanagh, 1990). Next we considered whether the computations described earlier would actually predict the decline in performance at longer sampling periods shown in Fig. 8(a). DP values were computed for inverting four-pattern stimuli at different sampling intervals (note that the minimum sampling period available was two frames, required to represent the positive- and negative-contrast versions of each pattern). Computational procedures were identical to those described earlier. DP values were computed before low-pass filtering the patterns, after low-pass filtering, and after full-wave rectification. Computational results are shown in Fig. 8(b). Open symbols depict DP values computed from source images, crosses depict DP values after source images were low-pass filtered, and solid symbols depict DP values after low-pass filtered patterns were full-wave rectified. As in the earlier computations, DP values computed without rectification predict either no

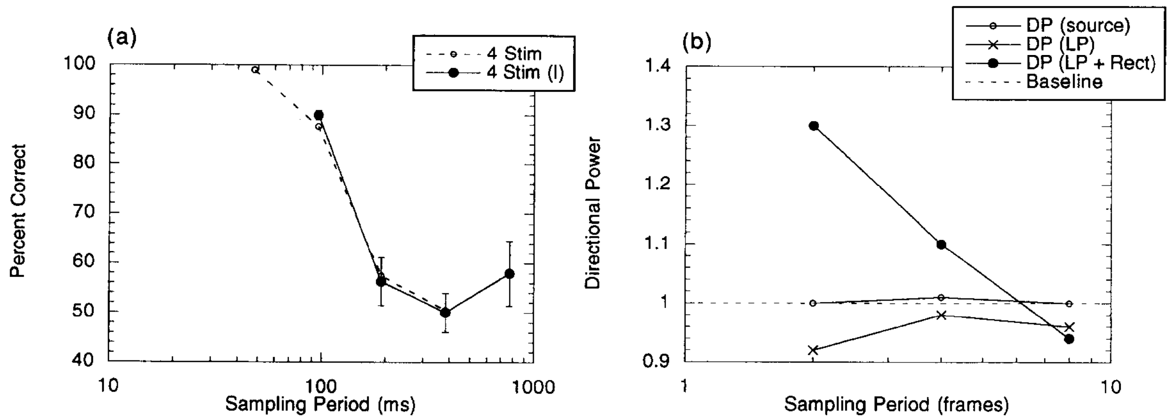

FIGURE 8. (a) Comparison of data from Expt 3 on contrast-inverting interleaved patterns (●) with data from Expt 1 on non-inverting patterns (○), re-plotted from Fig. 3. (b) Computed DP values for inverting patterns, as a function of sampling period. (○) Values computed from untransformed images, × show values computed from low-pass filtered images, and ● show values computed from images after filtering and rectification. Details as for previous figure.

consistent directional signal, or a reversed directional signal. DP values computed from rectified images predict an appropriate directional response, but only at short sampling periods, in agreement with psychophysical data. We repeated DP computations on the same patterns at the most effective sampling period (two frames) using half-wave rectification instead of full-wave rectification. No consistent signal was available: mean DP with positive half-wave rectification was 1.05; mean DP with negative half-wave rectification was 0.96; mean DP with full-wave rectification was 1.30, as plotted.

GENERAL DISCUSSION

The psychophysical and computational results presented here allow us to conclude that:

(1) observers can extract reliable directional information from temporally interleaved random dot patterns;
(2) direction discrimination is mediated by motion processes that extract Fourier energy from the image after full-wave rectification (or an equivalent transform).

Early centre–surround receptive fields in the visual system approximate half-wave rectifiers, and it is likely that so-called first-order motion energy analysis is based directly on signals from these receptive fields. Several lines of evidence point to motion analysis based on half-wave signals (e.g. polarity-specific adaptation effects (see Moulden & Begg, 1986; Anstis, 1990; Mather, Moulden & O'Halloran, 1991). However, our computations showed that half-wave rectification is not sufficient to expose the contrast-reversing stimulus to motion analysis. Instead, we must assume that half-wave rectified signals are summed to create a full-wave signal for motion analysis. Such a sequence of operations constitutes "second-order" motion analysis (Cavanagh & Mather, 1989), and forms part of the model proposed by Wilson, Ferrera and Yo (1992).

Since the putative visual mechanism underlying performance is second-order, it follows that the four-pattern inverting stimulus used in Expts 2 and 3 should be viewed as a second-order motion stimulus, though superficially it is unlike other second-order stimuli (e.g. beat patterns, moving texture borders etc.). Its spatiotemporal frequency spectrum (shown in Fig. 4) contains no energy passing through the origin, but does contain concentrations of energy along lines in frequency space that do not pass through the origin. The spatiotemporal orientation of these lines corresponds to the velocity of the contrast inverting pattern. Fleet and Langley (1994) have shown that this kind of transform is typical of many second-order motion displays. In the frequency domain, full-wave rectification of the contrast-inverting pattern serves to introduce energy passing through the origin of the spatiotemporal frequency transform, as illustrated in Fig. 6. In the space–time domain, full-wave rectification registers the presence and absence of texture regions, and conveys information about the properties of the texture [Werkoven, Sperling and Chubb (1993) describe second-order detectors as "texture-grabbers"]. The interleaved patterns employed in the present experiments can be viewed as spatiotemporal texture patterns—texture varies randomly both across space and across time, but this spatiotemporal variation (e.g. defined by a collection of four interleaved patterns, whether contrast inverting or not) shifts coherently over time. According to our experiments, second-order detectors that "grab" this spatiotemporal texture are subject to sampling limits just as first-order detectors are, so coherent responses arise only at higher sampling rates (i.e. short sampling periods, as defined in Fig. 1).

Given the success of second-order detectors in predicting our data, and the similar performance obtained using contrast-inverting and non-inverting stimuli, is there any need to propose the existence first-order detectors at all? Polarity-selective motion adaptation is inconsistent with full-wave rectification, and suggests the presence of half-wave rectifying detectors. Other experiments also indicate the presence of at least two parallel motion analyses based respectively on half-wave and full-wave rectifying transforms (e.g. Chubb & Sperling, 1989; Solomon & Sperling, 1994). Further, several psychophysical studies indicate that observers have difficulty in detecting motion when order changes from frame to frame (Mather & West, 1992; Ledgeway & Smith, 1994). A second-process employing full-wave rectification, as proposed here, would be unable to integrate coherent signals across frames that switch order, because full-wave rectification of luminance modulated patterns results in a doubling of the modulation frequency, whereas full-wave rectification of texture modulated patterns does not result in frequency doubling (see Ledgeway & Smith, 1994). In order switching displays the unequal modulation frequencies that follow rectification should confound directional signals.

In conclusion, the simplest way to reconcile results presented here with results from previous research is to assume that the visual system possesses at least two populations of motion detector. In one population, motion energy analysis is preceded by half-wave rectification, and in the other it is preceded by full-wave rectification. These two populations can be identified with first-order and second-order motion detectors respectively.

REFERENCES

Anstis, S. M. & Rogers, B. J. (1975). Illusory reversal of visual depth and movement during changes of contrast. *Vision Research*, *15*, 957–961.

Anstis, S. M. & Rogers, B. J. (1986). Illusory continuous motion from oscillating positive–negative patterns: Implications for motion perception. *Perception*, *15*, 627–640.

Boulton, J. C. & Baker, C. L. (1993). Different parameters control motion perception above and below a critical density. *Vision Research*, *33*, 1803–1811.

Cavanagh, P. & Mather, G. (1989). Motion: The long and short of it. *Spatial Vision*, *4*, 103–129.

Chubb, C. & Sperling, G. (1989). Two motion systems revealed through distance-driven reversal of apparent motion. *Proceedings of the National Academy of Science, U.S.A.*, *86*, 2985–2989.

Dosher, B. A., Landy, M. S. & Sperling, G. (1989). Kinetic depth effect from optic flow—I. 3D shape from Fourier motion. *Vision Research*, *29*, 1789–1813.

Fleet, D. J. & Langley, K. (1994). Computational analysis of non-Fourier motion. *Vision Research*, *34*, 3057–3079.

Ledgeway, T. & Smith, A. T. (1994). Evidence for separate motion-detecting mechanisms for first- and second-order motion in human vision. *Vision Research*, *34*, 2727–2740.

Mather, G. & West, S. (1993). Evidence for second-order motion detectors. *Vision Research*, *33*, 1109–1112.

Mather, G. & Tunley, H. (1993). Temporal filtering enhances motion detection. *Perception*, *22*, S31.

Mather, G. & Tunley, H. (1995). Temporal filtering enhances direction discrimination in random-dot patterns. *Vision Research*, *35*, 2105–2116.

Mather, G., Moulden, B. & O'Halloran, A. (1991). Polarity specific adaptation to motion in the human visual system. *Vision Research*, *31*, 1013–1019.

Moulden, B. P. & Begg, H. (1986). Some tests of the Marr–Ullman model of movement detection. *Perception*, *15*, 139–155.

Nishida, S. & Sato, T. (1992). Positive motion after-effect induced by bandpass-filtered random-dot kinematograms. *Vision Research*, *32*, 1635–1646.

Shioiri, S. & Cavanagh, P. (1990). ISI produces reverse apparent motion. *Vision Research*, *30*, 757–768.

Solomon, J. A. & Sperling, G. (1994). Full-wave and half-wave rectification in second-order motion perception. *Vision Research*, *34*, 2239–2257.

Watson, A. B., Ahumada, A. J. & Farrell, J. E. (1986). Window of visibility: A psychophysical theory of fidelity in time-sampled visual motion displays. *Journal of the Optical Society of America A*, *3*, 300–307.

Werkhoven, P., Sperling, G. & Chubb, C. (1993). The dimensionality of texture-defined motion: A single channel theory. *Vision Research*, *33*, 463–485.

Wilson, H. R., Ferrera, V. P. & Yo, C. (1992). A psychophysically motivated model for two-dimensional motion perception. *Visual Neuroscience*, *9*, 79–97.

Acknowledgements—This work was supported by grants from the MRC and BBSRC.