PERGAMON

Vision Research 39 (1999) 3710–3720

Vision Research

www.elsevier.com/locate/visres

Discrimination of the speed and direction of global second-order motion in stochastic displays

Timothy Ledgeway *

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

Received 10 August 1998; received in revised form 22 February 1999

Abstract

The ability to integrate local second-order motion signals over space and time was examined using random-dot-kinematograms (RDKs) in which the dots were defined by spatial variation in the contrast, rather than luminance, of a random noise field. When either the speeds or the directions of the individual dots were selected at random from a range of possible values, globally the stimulus appeared to drift either in a single direction or at a single speed in a manner analogous to that reported previously for first-order (luminance-defined) RDKs. To quantify the precision with which observers could extract the global stimulus motion, speed- and direction-discrimination thresholds were measured using pairs of RDKs, one of which (the comparison) comprised dots whose speeds or directions were assigned stochastically and the other (the standard) comprised dots that all had the same drift direction and speed. Speed-discrimination thresholds were of the order of 8% and changed little as the range of dot speeds (bandwidth) of the comparison increased, in that performance was almost as good when the individual dot speeds were selected at random from a range spanning 3.84 deg/s as when all the dots moved at the same speed. There was a tendency for the perceived global speed of the comparison RDK to decrease as the speed bandwidth was increased and perceived speed tended to coincide with the geometric mean speed of the dots rather than the arithmetic mean speed. Direction-discrimination thresholds were lowest (~ 4°) when the range of dot directions was less than 90° but increased markedly thereafter. Observers were able to perform both discrimination tasks when the lifetimes of the dots comprising the RDKs was reduced from 25 to 2 frames, a manipulation that prevented observers from determining the overall speed or direction of image motion from the extended trajectories of individual dots within the display. Thresholds under these conditions were somewhat higher but were otherwise comparable to those obtained with a dot lifetime of 25 frames. The similarities between the present results and those of previous studies that have employed first-order RDKs suggest that the extraction of the global speed and direction of each type of motion is likely to be based on computationally similar principles. © 1999 Elsevier Science Ltd. All rights reserved.

Keywords: Stochastic displays; Second-order motion; Discrimination

1. General introduction

Drifting contours defined by either first-order image statistics (luminance) or second-order image statistics such as contrast and texture granularity can give rise to compelling impressions of motion (Chubb & Sperling, 1988; Cavanagh & Mather, 1989). The initial stages of first-order motion extraction have been studied extensively and on the basis of both physiological (e.g. Hubel & Wiesel, 1968) and psychophysical evidence (e.g. Anderson & Burr, 1987) are presumed to operate over restricted regions of space. Indeed first-order motion sensors have been successfully modelled as spatiotemporal filters that are localised in space (Adelson & Bergen, 1985). The principles underlying second-order motion detection have yet to be clearly elucidated, but the suggestion (Chubb & Sperling, 1988, 1989; Johnston, McOwan & Buxton, 1992; Wilson, Ferrera & Yo, 1992; Werkhoven, Sperling & Chubb, 1993) that local motion is extracted using analogous principles (but not necessarily the same mechanisms) to those used for first-order motion has received empirical support (Cavanagh & Mather, 1989; Smith, Hess & Baker, 1994; Nishida, Ledgeway & Edwards, 1997). In the case of second-order motion, however, some form of nonlinear processing (e.g. rectification) prior to motion analysis is typically deemed necessary.

* Present address: McGill Vision Research Unit, 687 Pine Avenue West, Rm. H4-14, Montreal, Quebec, Canada H3A 1A1.

E-mail address: tledgewa@vision.mcgill.ca (T. Ledgeway)

PII: S0042-6989(99)00082-6

The earliest stages of motion analysis are believed to involve detectors that only explicitly encode the local direction of image motion and not its speed, as neurones in mammalian cortex tend to be tuned for temporal frequency rather than speed (Holub & Morton-Gibson, 1981; Foster, Gaska, Nagler & Pollen, 1985). However, several computational schemes have been developed (Heeger, 1987; Grzywacz & Yuille, 1990; Simoncelli & Heeger, 1998), at least in the context of first-order motion, in which local speed is coded by grouping the responses of detectors, each with the same preferred velocity, that sample a similar region of space. In Heeger's model, for example, local speed is given by the peak response of several sets of such groupings having different velocity sensitivities. These schemes allow the local direction and speed of first-order (and in principle second-order) motion to be recovered, but suffer from the ambiguity inherent in any motion-detecting process that operates locally. This is characterised by the well known aperture problem: A one-dimensional (1-d) contour, of a translating two-dimensional (2-d) object, viewed through a spatial aperture has a velocity that is not necessarily indicative of the global (overall) speed and direction of the object of which it is part. To determine the global velocity of the object as a whole, local motion estimates must be subsequently combined (integrated). Relatively little is known about these integrative processes, particularly in the case of second-order motion, although some phenomena are difficult to explain without recourse to either spatial or temporal integration. For example, sensitivity to second-order motion is influenced by both the spatial area (Zanker, 1993) and temporal duration of the stimulus (Derrington, Badcock & Henning, 1993).

Random-dot-kinematograms (RDKs) have proved useful for studying the perception of global first-order motion. Williams and Sekuler (1984), for example, used RDKs in which the individual dot directions were selected from a uniform probability distribution such that each dot was assigned an independent, random walk in direction over time. They reported that provided that the range of directions was constrained, globally the stimulus appeared to drift in a direction close to the mean of the dot directions. The processes that serve to integrate local direction signals do not necessarily implement a strict averaging strategy, however, as the perceived direction of global motion can be biased, to a limited extent, towards the mode and away from the mean when the distribution of dot directions is skewed asymmetrically (Zohary, Scase & Braddick, 1996). Williams and Sekuler found that the probability of perceiving global motion increased with stimulus duration up to ~ 440 ms and depended only on the overall set of dot directions present rather than the particular paths (trajectories) traveled by individual dots. They proposed that individual dot motions were initially detected independently, combined across space and then over time to produce the percept of global motion. This notion has received support (Williams, Phillips & Sekuler, 1986; Williams & Phillips, 1987; Smith, Snowden & Milne, 1994) and models of global motion perception have been developed (Watamaniuk, Sekuler & Williams, 1989; Williams, Tweten & Sekuler, 1991).

In order to quantify the precision with which observers can utilise the direction of global motion, studies (e.g. Watamaniuk & Sekuler, 1992; Watamaniuk, 1993) have measured direction-discrimination thresholds. That is, the minimum difference between the mean directions of two RDKs needed to determine reliably that they are different. Watamaniuk et al. (1989) found that for long presentations (~ 1.4 s) thresholds approximately doubled when the bandwidth (standard deviation) of the Gaussian probability distribution, from which the individual dot directions were selected, was increased from 0 to 51° and for shorter presentations greater increases occurred. This indicates that the efficacy with which the direction of global motion is extracted is determined in part by the directional bandwidth (variability) of the local motion signals present. However, even at the most extreme bandwidth used thresholds were $\sim 3°$. In a related study Watamaniuk and Duchon (1992) measured global speed-discrimination thresholds using RDKs in which the individual dots took a random walk in speed, rather than direction, over time. They presented evidence that observers tended to base discrimination on changes in the mean dot speed, rather than changes in either the median or mode speeds, and were able to discriminate reliably less than a 9% change in speed. Thresholds were largely unaffected as the width of the probability distribution increased and were similar whether the speeds assigned to individual dots were selected at random after each displacement or remained fixed. This suggests that judgements were indeed based on the global motion of the display rather than the speeds of, or distances traveled by, individual dots.

Unlike first-order motion, few studies have investigated the perception of global motion from stochastic second-order motion signals. Edwards and Badcock (1995) addressed this issue using RDKs in which the dots were defined by contrast differences with respect to the background, rather than luminance. The direction of motion could be reliably identified when as few as 10% of the dots ('signal' dots) were displaced in that direction and the remainder ('noise' dots) were displaced randomly over a 360° range. The particular dots that comprised the 'signal' and 'noise' were selected at random after each displacement so integration of direction information along the trajectories of individual dots could not be used to perform the task. Adding first-order 'noise' dots to the display increased the number of second-order 'signal' dots needed to determine the direction of global motion but not vice versa.

The authors took this as evidence that first-order and second-order motion are extracted by independent visual pathways (cf Chubb & Sperling, 1988, 1989; Wilson et al., 1992) that are segregated up to and including the stage at which global motion analysis occurs. To explain the partial interaction between first-order and second-order dot motions they suggested that local second-order motion sensors are also sensitive to first-order motion. This conclusion remains tentative, however, as it predicts that observers should be able to perceive motion reliably between first-order and second-order dots when they are interleaved in a RDK and this is not the case (Mather & West, 1993). Nevertheless if the mechanisms that extract the global motion of first-order and second-order patterns are at least partially distinct then it is important to establish the properties of the putative second-order motion system.

The accuracy with which changes in the speed of second-order motion can be discriminated has received little attention, although Johnston and Benton (1997) reported Weber fractions as low as 0.09, under optimal conditions, for discriminating the speed of a contrast-defined bar or edge. Discrimination performance for comparable first-order motion was similar, though better in most cases. Integration of local second-order speed signals was not the focus of this study, but it nonetheless demonstrates that observers can utilise the speed of some varieties of second-order motion with an efficacy similar to that of first-order motion.

It is apparent that little is known about the processes by which local second-order motion signals are combined within the visual system and several crucial issues have yet to be studied, such as the precision with which observers can extract global motion from stochastic patterns. A simple direction-identification task, involving a choice between two opposite directions of motion (Edwards & Badcock, 1995), could be performed on the basis of a relatively crude estimate of global direction that only needs to be accurate to within $\pm 90°$. Indeed when the bandwidth (variability) of the dot directions is large, precise encoding of the direction of global second-order motion may not be possible. This uncertainty applies equally to the encoding of the global speed of second-order motion of which little is known. The aim of the present study was to address these issues by measuring discrimination thresholds for second-order, random-walk RDKs in which the individual dot speeds or directions were drawn from a uniform probability distribution with a range of distribution widths. In line with previous studies (Smith et al., 1994; Zohary et al., 1996), the lifetimes of the dots were limited under some conditions to ensure that discrimination was based on global motion rather than the extended trajectories of individual dots in the display.

2. Experiment 1: speed discrimination

2.1. Methods

2.1.1. Observers

Two observers participated in the experiment and each had normal or corrected-to-normal acuity. Observer TL was the author and JGC was a naive observer who was unaware of the purpose of the experiment.

2.1.2. Apparatus and stimuli

Motion stimuli (RDKs) were computer generated and displayed on an APPLE monochrome monitor (with a frame rate of 67 Hz) which was carefully gamma-corrected with the aid of internal look up tables. As an additional precaution psychophysical procedures were used to ensure that any residual luminance nonlinearities were minimised (see Ledgeway & Smith, 1994; Nishida et al., 1997). The stimuli were presented within a square window at the centre of the display which subtended an angle of $1.94 \times 1.94°$ at the viewing distance of 2.05 m. The mean luminance of the remainder of the display (which was homogeneous) was approximately 18 cd/m^2. Viewing was binocular and a prominent fixation spot was located at the centre of the display in order to minimise eye movements and prevent ocular tracking of the motion stimuli.

Each RDK was generated anew immediately prior to its presentation (on any one trial) and was composed of a sequence of 25 images. Continuous apparent motion was produced by presenting the images consecutively at an update rate of 22.3 Hz, which is comparable to those used in previous studies of global motion (e.g. Williams & Sekuler, 1984; Watamaniuk et al., 1989; Watamaniuk & Duchon, 1992; Watamaniuk & Sekuler, 1992; Edwards & Badcock, 1995). Each image contained 128 dots and the diameter of each dot was 5.72 arc min. The dot density was approximately 34 dots/deg^2. Each dot was composed of two-dimensional, static noise produced by assigning individual screen pixels (0.57×0.57 arc min) within the area of each dot to be 'black' (0.03 cd/m^2) or 'white' (36.9 cd/m^2) with equal probability[1]. The remainder of the display area (back-

[1] Smith and Ledgeway (1997) have recently reported that under some circumstances the use of static noise may give rise to local first-order luminance artifacts in contrast-defined patterns at threshold, as a consequence of local clustering of noise pixels with the same polarity. However, such artifacts appear to be minimal, or absent, when there is no spatial variation in luminance within each noise pixel and the contrast profile of the image is displaced by integer numbers of pixels (Nishida et al., 1997) as they were in this study. Furthermore, several lines of empirical evidence (e.g., Ledgeway & Smith, 1994, 1997) are consistent with the view that at suprathreshold modulation depths such artifacts, if indeed present, are likely to be too small to influence greatly psychophysical performance.

ground) was a uniform mid 'grey'. On the first frame of each RDK the dot positions were determined randomly and on subsequent frames were shifted in the same direction. When a dot reached the edge of the square display window it was 'wrapped-around' so that it immediately reappeared on the opposite edge of the window.

Two basic types of RDK were employed in which the lifetimes of the dots were manipulated. For one type of RDK the lifetime of all dots was 25 frames (i.e. the duration of the entire motion sequence) and for the other type of RDK the lifetime of the dots was two frames (i.e. the dot positions were randomised after each successive displacement).

2.1.3. Procedure

On each trial two RDKs were presented sequentially. The duration of each RDK was 1.1 s and they were separated by a 1 s interval containing a homogenous blank field. One of the two RDKs, referred to as the *standard*, had the same characteristics on every trial of a given run. The dots of the standard RDK all moved in the same direction and at the same drift speed of 2.97 deg/s (i.e. all dots were displaced by $\sim 0.13°$ on each frame). The direction of motion was either upwards, downwards, leftwards or rightwards and was chosen at random at the beginning of each trial but remained fixed throughout the duration of the standard RDK. This was done to ensure that the perceived speeds of the RDKs were not affected by prolonged viewing, across trials, of second-order motion always in the same direction (Ledgeway & Smith, 1997). The *comparison* RDK again comprised dots which all had the same direction (same as the standard RDK) but each took a random walk in speed. The step size (displacement) of each dot on each frame was drawn, with replacement, from a uniform (rectangular) probability distribution and was independent of both its step size on previous updates and the step sizes of the other dots. As it was not possible to present a continuum of speeds, due to the need to displace each dot by an integer number of pixels, a uniform distribution of dot speeds was approximated by sampling at 0.21 deg/s intervals. Four distribution widths (0, 1.28, 2.56 and 3.84 deg/s) were used for the comparison RDK and the distribution width was constant for each run of trials. The arithmetic mean speed of the uniform probability distribution varied from trial to trial and was chosen at random from a set of nine speeds (selected on the basis of pilot studies) that differed from the standard RDK by about either 0, ± 7, ± 14, ± 21 or $\pm 28\%$. The largest usable step size for the dots of the comparison RDK was regarded as the step size which gave a clear percept of coherent motion when *all* of the dots moved in the same direction with that step size. This was $\sim 0.30°$ (6.70 deg/s) for the present stimuli. Beyond this speed, the percept of coherent motion became weaker even for zero distribution width and thus in the present experiment the step size of any individual dot never exceeded 0.27 (6.03 deg/s). The temporal order of presentation of the standard and comparison RDKs was randomised from trial to trial. The task of the observer was to compare the global speeds of the RDKs on each trial and decide, using one of two response buttons, which appeared to move faster. Observers completed six runs of 90 trials for each distribution width of the comparison RDK and the order in which the runs were completed was randomised for each observer.

Weibull (1951) functions were fitted to the resulting data expressed in terms of the percentage of trials on which each observer judged that the comparison RDK was drifting faster than the standard RDK, as a function of the percentage difference between the arithmetic mean speed of the comparison RDK and that of the standard RDK. The midpoint of each psychometric function was used to evaluate the physical speed of the comparison RDK at which the observer was unable to distinguish between the perceived speeds of the standard and comparison RDKs. This enabled any effect of distribution width on the perceived speed of the comparison RDK to be examined. Speed-discrimination thresholds were calculated as half the percentage difference in speed between the 75 and 25% response levels on the psychometric function for each distribution width examined.

3. Results and discussion

3.1. Random-walk RDKs with a dot lifetime of 25 frames

Fig. 1 shows typical psychometric functions for two observers (TL and JGC) for random-walk RDKs with a dot lifetime of 25 frames. It is clear that both observers were able to perform the speed-discrimination task at all four distribution widths of the comparison RDK examined. For each observer there was a moderate but clear tendency for the midpoints of the functions to shift laterally along the abscissa in the increasing speed direction as distribution width increased. This indicates that as the range of dot speeds in the comparison RDK increased, the perceived speeds of the two RDKs were judged equal when the mean physical speed of the comparison RDK was typically greater than that of the standard RDK. This is quantified in Fig. 2 (top) which shows the arithmetic mean physical speed of the comparison RDK that had the same perceived speed as the standard RDK, as a function of distribution width. As the distribution width increased to 3.84 deg/s the mean physical speed of the comparison RDK had to be 3.21 deg/s for TL and 3.19 deg/s for JGC in order to be indistinguishable from that of the standard RDK in which all dots drifted at speed of 2.97 deg/s. This suggests that the perceived speed of a second-order motion, random-walk stimulus does not necessarily coincide with the arithmetic mean of the distribution of dot speeds, especially when the range of dot speeds is large. As the step sizes of all of the individual dots comprising the comparison RDK were constrained to fall well within the range supporting the perception of coherent motion (see Section 2.1), it is clear that this finding is not simply an artifact resulting from a failure to perceive coherent motion of the dots with the highest step sizes.

In terms of speed-discrimination performance for observer TL, the slopes of the psychometric functions shown in Fig. 1 (top) exhibit, if anything, a tendency to become shallower as the distribution width increased. The change in slope was relatively modest however given the high degree of variability introduced into the speeds of the dots across the four distribution widths. Observer JGC (Fig. 1, bottom) shows a similar pattern of results, in that the slopes of the psychometric functions changed remarkably little with distribution width, with the exception that for a distribution width of zero the psychometric function was markedly steeper than those for the other conditions. This is quantified in Fig. 2 (bottom) which shows, for each observer, speed-discrimination thresholds (which are inversely related to the slopes of the psychometric functions) as a function of the comparison RDK distribution width. The absolute threshold values were similar for the two observers and were lowest for a distribution width of 0 deg/s. The data are in close agreement with those of Johnston and Benton (1997) in terms of the range and magnitude of the speed-discrimination thresholds found. For example, averaged over all distribution widths, observers TL and JGC could both reliably discriminate a change in speed of $\sim 8\%$.

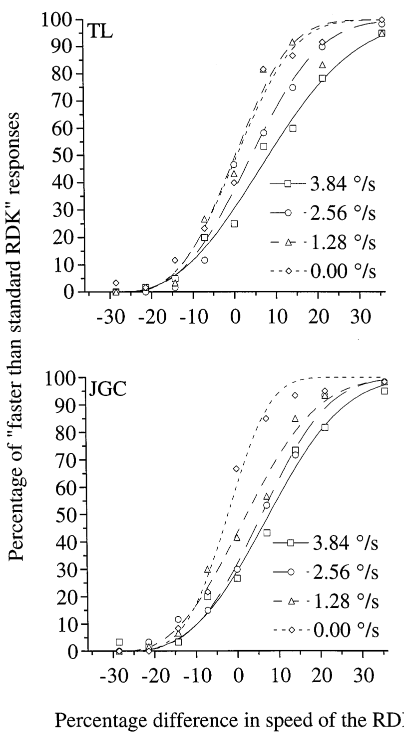

Fig. 1. Psychometric functions showing the percentage of trials on which observers TL (top) and JGC (bottom) perceived the comparison RDK to drift faster than the standard RDK as a function of the percentage difference between the arithmetic mean speed of the comparison RDK and that of the standard RDK. The speeds of the individual dots comprising the comparison RDK were drawn from a rectangular probability distribution with a bandwidth ranging from 0 deg/s to 3.84 deg/s (indicated by the different symbols). The dots of the standard RDK all drifted at the same speed of 2.97 deg/s. Positive values on the abscissae indicate that the arithmetic mean speed of the comparison RDK was greater than the standard RDK, negative values indicate the opposite and zero indicates a speed of 2.97 deg/s. The dots in each stimulus were defined by spatial variation in the contrast of a random noise field and had a lifetime of 25 frames.

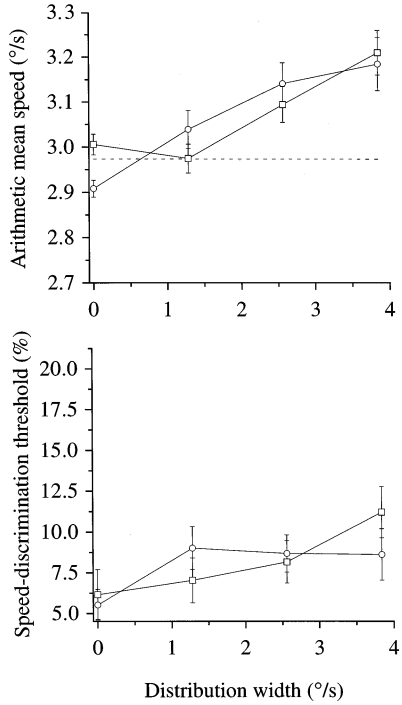

Fig. 2. The arithmetic mean physical speed of the comparison RDK that had the same perceived speed as the standard RDK, plotted as a function of the comparison RDK speed bandwidth (top). The dashed horizontal line represents the physical speed of the standard RDK (2.97 deg/s). Speed-discrimination thresholds (bottom) for second-order motion, expressed as the percentage change in arithmetic mean speed, as a function of the distribution width of the comparison RDK dot speeds. The dot lifetime was 25 frames. Squares and circles represent data for observers TL and JGC, respectively. The vertical bars above and below each data point (where visible) represent ± 1 S.E.M. based on variability between runs of trials.

3.2. Random-walk RDKs with a dot lifetime of two frames

For RDKs with a dot lifetime of two frames (Fig. 3, top) the mean physical speed of the comparison RDK that appeared to drift at the same rate as the standard RDK, shows a similar dependence on distribution width as the RDKs with a lifetime of 25 frames (Fig. 2, top) for both observers. The speed of the comparison RDK required to match that of the standard RDK, increased with increasing distribution width. Speed-discrimination thresholds (Fig. 3, bottom) for RDKs with a dot lifetime of two frames exhibit greater variability both between observers and between the different distribution widths than those for the 25 frame dot lifetime conditions. The absolute thresholds were somewhat higher (by a factor of 1.6 for TL and 1.8 for JGC) with a dot lifetime of two frames, but it is clearly evident that integration of speed signals along the trajectories of individual dots was not necessary to perform the task.

3.3. Statistical analysis

To investigate whether or not the observed differences in performance were statistically significant, analysis of variance (ANOVA) was performed on the data. This revealed that the arithmetic mean speed of the comparison RDK that had the same perceived speed as

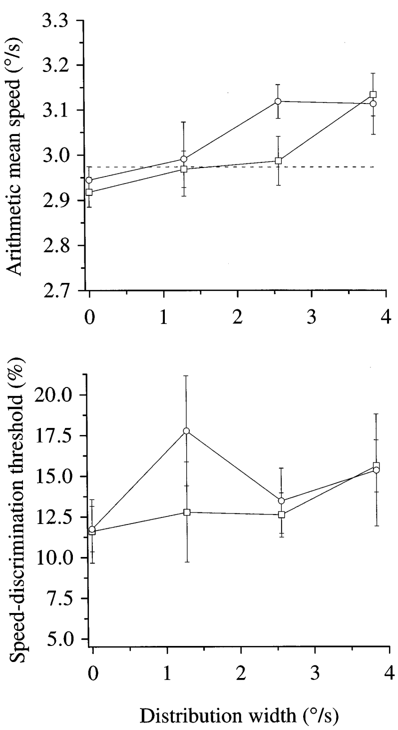

Fig. 3. Legend as for Fig. 2 with the exception that the dot lifetime was reduced from 25 to 2 frames (i.e. the dot positions were randomised after each successive displacement).

the standard RDK increased significantly with distribution width ($F_{(3,30)} = 18.48$; $P = 0.00001$) but did not differ significantly between the two observers or across the dot lifetime conditions. Speed-discrimination thresholds were significantly higher for the two frame than the 25 frame lifetime condition ($F_{(1,10)} = 58.17$; $P = 0.00001$) and increased significantly with distribution width ($F_{(3,30)} = 5.02$; $P = 0.01$). The pattern of thresholds obtained across all conditions was not significantly different for the two observers.

4. Experiment 2: direction discrimination

4.1. Methods

4.1.1. Observers

Observers TL and JGC were the same observers that participated in experiment 1.

4.1.2. Apparatus and stimuli

The apparatus and stimuli were identical to those used in experiment 1

4.1.3. Procedure

Discrimination thresholds were measured using a similar procedure to that described for experiment 1 with the following exceptions. The dots of the standard and comparison RDKs all drifted at the same speed which was 2.97 deg/s. The dots of the standard RDK were all displaced in the same direction and this direction was chosen at random from trial to trial from a range spanning 360° (it was necessary to constrain the directions to a set spaced at intervals of about 4° as a consequence of the need to move each dot an integer number of pixels in the vertical and horizontal directions). Each dot comprising the comparison RDK took a random walk in direction drawn from a uniform (rectangular) probability distribution. Five distribution widths (0, 30, 60, 90 and 120°) were used for the comparison RDK but the distribution width was fixed for each run of trials. The arithmetic mean direction of this uniform probability distribution varied from trial to trial and was chosen at random from a set of nine directions (selected on the basis of pilot studies) that differed from the standard RDK by about either 0, ± 8, ± 16, ± 24 or ± 32°. The task of the observer was to compare the directions of the two RDKs on each trial and decide which had a direction rotated clockwise relative to the other. Observers completed six runs of 90 trials for each distribution width of the comparison RDK and the order in which the runs were completed was randomised for each observer.

Weibull (1951) functions were fitted to the resulting data expressed in terms of the percentage of trials on which each observer judged that the comparison RDK

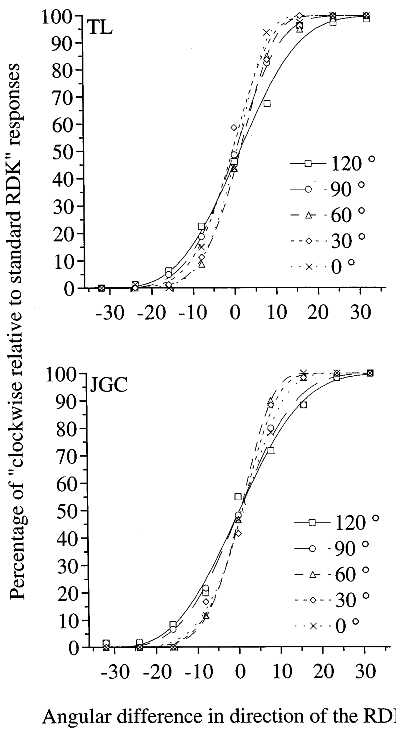

Fig. 4. Psychometric functions showing the percentage of trials on which observers TL (top) and JGC (bottom) perceived the comparison RDK to drift in a direction rotated clockwise relative to the standard RDK as a function of the angular difference between the arithmetic mean direction of the comparison RDK and that of the standard RDK. The directions of the individual dots comprising the comparison RDK were drawn from a rectangular probability distribution with a bandwidth ranging from 0 to 120° (indicated by the different symbols). The dots of the standard RDK all drifted in the same direction and this direction was chosen at random on each trial. Positive values on the abscissae indicate that the arithmetic mean direction of the comparison RDK was clockwise relative to the standard RDK, negative values indicate the opposite and zero indicates that the drift directions were the same. The dots in each stimulus had a lifetime of 25 frames.

direction was rotated clockwise relative to that of the standard RDK, as a function of the angular difference between the arithmetic mean direction of the comparison RDK and that of the standard RDK. Direction-discrimination thresholds were calculated as half the angular difference in direction between the 75 and 25% response levels on the psychometric function for each distribution width examined.

5. Results and discussion

5.1. Random-walk RDKs with a dot lifetime of 25 frames

Fig. 4 shows typical psychometric functions for the two observers. It is evident that the observers show the same pattern of results in that there was a tendency for the slope of the psychometric functions to become shallower as the distribution width increased particularly for the two largest distribution widths examined. This is quantified in Fig. 5 (top) which shows direction-discrimination thresholds as a function of the distribution width of the comparison RDK. For observer JGC thresholds for the three narrowest distribution widths are similar (~ 4–5°) and then increase markedly up to ~ 7° for the 90° and 120° conditions. Observer TL shows the same deterioration in performance as distribution width was increased but the decrement is, if anything, more gradual and his thresholds are typically lower than those of observer JGC.

5.2. Random-walk RDKs with a dot lifetime of two frames

Although the absolute threshold values for all distribution widths were uniformly larger (Fig. 5 bottom) by about a factor of 1.6 for each observer, the shape of the threshold versus distribution width function does not differ from that obtained when the dot lifetime was 25 frames. Thresholds for both observers changed little as the distribution width increased from 0 to 60° and then rapidly increased thereafter. Direction-discrimination

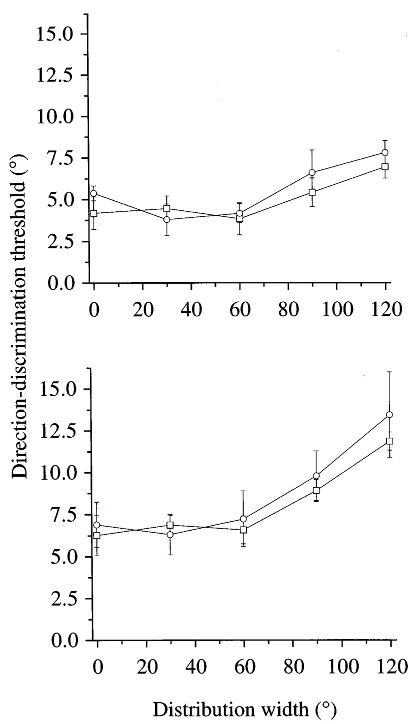

Fig. 5. Direction-discrimination thresholds for second-order motion as a function of the distribution width of the comparison RDK dot directions. The dot lifetime was either 25 frames (top) or two frames (bottom). Squares and circles represent data for observers TL and JGC, respectively. The vertical bars above and below each data point (where visible) represent ±1 S.E.M. based on variability between runs of trials.

for the limited-lifetime dots must have been based upon the global directions of motion of the stimuli and under these conditions it is clear that even at the most extreme distribution width used (120°) observers were able to discriminate reliably a difference in global direction of the order of 13°.

5.3. Statistical analysis

ANOVA confirmed that direction-discrimination thresholds were significantly higher for the two frame than the 25 frame lifetime condition ($F_{(1,10)} = 52.08$; $P = 0.00001$) and increased significantly with increasing distribution width ($F_{(4,40)} = 13.72$; $P = 0.00001$). The pattern of results obtained across all conditions did not differ significantly between the two observers.

6. General discussion

The primary aim of the present study was to explore the accuracy with which observers can extract the global speed and direction of second-order motion in stochastic displays and thus probe the mechanisms that serve to integrate local velocity signals in second-order motion patterns. The precision with which observers could utilise the speed of second-order motion, as measured by speed-discrimination thresholds, was dependent on whether the individual dots existed for the entire duration of the display (25 frames) or were limited in terms of their lifetime (2 frames). In the case of the extended lifetime dots observers could discriminate a change in speed of about 8% and thresholds increased only marginally as the bandwidth of the distribution of dot speeds increased from 0 to 3.84 deg/s. When the dot lifetime was two frames, absolute thresholds increased to about 14% but again changed little with distribution width. In terms of direction discrimination, thresholds for the extended lifetime dots were lowest ($\sim 4°$) for the narrowest distribution widths (0 to 90°) and increased to about 7° at the widest distribution widths. Reducing the dot lifetime to two frames elevated thresholds by a factor of about 1.6 but performance showed the same dependence on distribution width as in the extended lifetime condition.

The results of the present study show that although information from individual dot trajectories might, in principle, be used for discriminating the direction and speed of second-order motion when the lifetime of each dot spans the entire duration of the display, the performance for the limited lifetime dots shows that such information is not necessary. Under these latter conditions observers had to use the globally-defined speed and direction of the random-walk RDKs to perform the task. This clearly implies that the outputs of mechanisms sensitive to the local speeds and directions of second-order motion are indeed combined in order to compute the global velocity of image motion. It is interesting that the effects of varying the bandwidth of the probability distribution from which the individual dot speeds or directions were drawn were qualitatively similar for both the extended and the restricted dot lifetime conditions. The most parsimonious explanation of this result is that observers always utilised the global motion of the RDKs and that differences in the absolute thresholds between the different dot lifetime conditions reflect the additional noise introduced into the two frame lifetime RDKs by the successive randomisation of the dot positions after each displacement. In the case of first-order motion, comparable differences in sensitivity have been reported using analogous displays (Smith et al., 1994) and there is much evidence that discrimination is not dependent on the paths traveled by individual dots over time.

Previous studies that have used random-walk RDKs (e.g. Williams & Sekuler, 1984) have shown when the distribution of dot directions are symmetrically distributed (but see Zohary et al., 1996), as they were in the present study, the overall direction of perceived global motion coincides approximately with the mean of the distribution. Similarly, Watamaniuk and Duchon (1992) demonstrated that for a range of distribution widths speed discrimination appears to be determined by the difference between the arithmetic mean speeds of two random-walk RDKs. However, unlike the present study, they did not actually measure if perceived speed corresponds to the arithmetic mean of the dot speeds present. It is clear from Figs. 2 and 3 (top) that although the arithmetic mean physical speed of the comparison RDK that had the same perceived speed as the standard RDK was close to 2.97 deg/s, there was a small but consistent tendency for the perceived speeds of the comparison RDKs to decrease as the distribution width increased.

One explanation for the observed reduction in perceived speed of the comparison RDK with increasing distribution width is that it is a consequence of presenting the dots within a spatial window. As the distribution width of the comparison RDK increased, the likelihood that any given dot would fall outside of the square display area on its next displacement was greatest for those dots that were randomly assigned the largest displacements. Displacements that cause dots to fall outside of the display area will be effectively invisible and consequently the arithmetic mean speed of the remaining dots in the display will tend to decrease as the distribution width increases. This will be true of any experiment that presents a range of dot speeds within a limited display area. The magnitude of this effect was evaluated for each distribution width using the same displacement algorithm as was used for the comparison RDK with a dot lifetime of 25 frames in experiment 1.

The arithmetic mean of all dot displacements that did not exceed the display area was computed for each motion sequence and the entire process was repeated sixty times (the same number of times each motion sequence was presented in experiment 1). This revealed that even for the widest distribution width used (3.84 deg/s) the predicted reduction in the arithmetic mean dot speed was very small (on average ~ 0.03 deg/s) and was similar for each of the nine comparison RDK speeds used. It is doubtful that this phenomenon considered in isolation could account for the largest reductions in perceived speed found in experiment 1, which would require a decrease in the effective arithmetic mean dot speed of up to 0.24 deg/s. An alternative, albeit more tentative, hypothesis is that as the physical speed of a second-order motion stimulus increases the response to it exhibits a compressive nonlinearity. If the form of this nonlinearity were (say) logarithmic, then the perceived speed of a random-walk RDK may be better characterised by the *geometric* mean of the individual dot speeds rather than the arithmetic mean. On this basis one would expect the perceived speed to decrease as the range of dots speeds is increased symmetrically about a given physical speed, as was found in experiment 1. A RDK centred on an arithmetic mean speed of say 2.97 deg/s and a distribution width of 0 deg/s has a geometric mean speed that is also 2.97 deg/s, but if the distribution width is increased to 3.84 deg/s the geometric mean speed will decrease to about 2.72 deg/s. In Fig. 6 the data from Figs 2 and 3 (top) are combined and plotted together with the predicted arithmetic mean speeds of the comparison RDK that would all have a geometric mean speed equal to 2.97 deg/s (the same as the standard RDK) for each distribution width examined. It is apparent that overall the empirical data are reasonably well described by assuming a nonlinear (in this case logarithmic) relationship between perceived speed and physical speed. Further research is needed to address the reliability of this finding but at present it is sufficient to conclude that the perceived global speed of a second-order motion stimulus composed of a range of dot speeds coincides closely with the geometric mean dot speed.

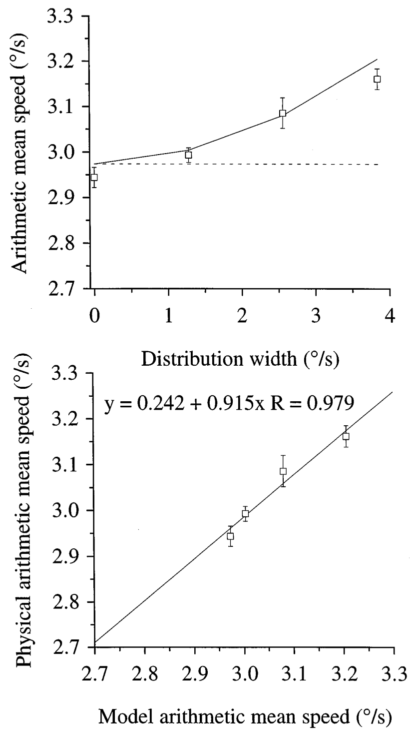

Fig. 6. The overall arithmetic mean speeds of the comparison RDK that had the same perceived speed as the standard RDK plotted as a function of the comparison RDK speed bandwidth (top). The dashed horizontal line represents the physical speed of the standard RDK (2.97 deg/s) and the continuous line shows the arithmetic mean speeds of the comparison RDK that would all have a corresponding geometric mean speed of 2.97 deg/s. Each square represents the average of the data points shown in Figs. 2 and 3 (top), for the two observers and the two dot lifetime conditions, at each comparison RDK bandwidth. Plotting the empirical data against the predictions of a model based on geometric mean speed (bottom) illustrates that there is a highly linear correlation between the two sets of values. The continuous line is the best fitting straight line through the data points and has a slope close to unity indicating that the model provides a good fit to the empirical data. The vertical bars above and below each data point (where visible) represent ± 1 S.E.M.

The marked similarities between the present results and those of previous studies that have employed comparable first-order motion stimuli suggest that the extraction of the global speed and direction of each type of motion is likely to be based on computationally similar principles. Although several possible schemes have been proposed, at least in the context of first-order motion, one model that has proved successful in accounting for the encoding of global direction in stochastic patterns is the line-element model of Watamaniuk et al. (1989) and Williams et al. (1991). The model is composed of a number of band-limited mechanisms covering all 360° of motion direction, with each having a Gaussian sensitivity profile for direction. That is, each mechanism exhibits directional tuning in that it responds maximally to a particular direction of motion. An individual mechanism's response is computed by multiplying its sensitivity to each local direction signal by the frequency with which each dot direction occurs in the stimulus and summing the results (Space is not treated explicitly in this model, but given that local motion sensors are themselves directionally tuned this scheme is equivalent computationally to summing linearly the outputs of local motion sensors tiling visual space that have the same preferred direction of motion). Importantly, the overall response of each mechanism depends only on the set of directions in the stimulus and not on the spatial distribution of those directions or the paths (trajectories) of individual dots. When a stimulus contains local motion signals that are close to the preferred direction of a particular mechanism, most of the model's activity will be centred about that mechanism. In this manner the visual system could compute global direction by selecting (say) the mechanism with the maximum response. Although not mentioned explicitly by the original authors, as the responses of the model's mechanisms take into account the frequency of each direction of local motion, the direction of global motion computed in this way will tend to coincide with the mean of the dot directions when all directions occur with equal frequency but will also be influenced by the mode when the distribution of dot directions is asymmetric (Zohary et al., 1996). If direction discrimination depends on the precision with which the global directions of motion of two random-walk RDKs can be compared (Watamaniuk et al., 1989) then it is apparent that increasing the distribution width (variability) will cause the resulting activity in the model to become more diffuse (less localised about a particular direction-selective mechanism) and consequently one might expect discrimination performance to eventually deteriorate. Similarly, randomising the dot positions after each displacement, as was done in the two frame dot lifetime displays, will introduce a large degree of directional noise into the model and again one would expect performance to be worse under these conditions.

Watamaniuk and Duchon (1992) proposed a similar line-element model for encoding global speed information. In this model, the band-limited mechanisms that serve to integrate local dot motions are selectively sensitive to a particular range of speeds, rather than directions, but the output of the model is computed in a manner analogous to that proposed for global direction. A nonlinear relationship between perceived (encoded) and physical speed could be accommodated within such a model if it is assumed that each band-limited mechanism has a sensitivity profile that is asymmetric about its preferred speed. If, for example, each mechanism was more sensitive to speeds higher than its preferred speed than to speeds comparably lower than its preferred speed, then the mechanisms activated most strongly by a RDK stimulus containing a broad range of speeds will tend to be those that have preferred speeds less than the arithmetic mean of the actual speeds present. Although line-element models were originally proposed to account for the perception of first-order global motion it is clear that the results of the present experiments could also be readily interpreted within the framework of these models if one assumes that the underlying mechanisms are sensitive to second-order motion.

An alternative model for combining local direction signals has been proposed by Wilson et al. (1992). In this model the outputs of separate populations of motion detectors that are sensitive to either first-order motion or second-order motion are subsequently combined according to a cosine weighted sum by 'pattern units' in order to compute a strictly local measure of the resultant direction of motion at a given point in the retinal image. Although the output of the model is a local estimate of image motion, a process such as that embodied within line-element models, that appropriately combined the outputs of pattern units sampling different locations in space, could in principle compute the global direction of image motion. However, at present the model does not offer an explicit mechanism by which either local or global speed could be encoded and it remains to be seen whether or not all stages of global motion processing involve mechanisms that are sensitive to both first-order and second-order motion (Edwards & Badcock, 1995).

References

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

Anderson, S. J., & Burr, D. C. (1987). Receptive field size of human motion detection units. *Vision Research*, *27*, 621–635.

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

Chubb, C., & Sperling, G. (1988). Drift-balanced random stimuli: a general basis for studying non-Fourier motion perception. *Journal of the Optical Society of America A*, *5*, 1986–2007.

Chubb, C., & Sperling, G., 1989. Second-order motion perception: space/time separable mechanisms. In *Proceedings of the IEEE workshop on visual motion* (pp. 126–138). Washington, DC: IEEE Computer Society Press.

Derrington, A. M., Badcock, D. R., & Henning, G. B. (1993). Discriminating the direction of second-order motion at short stimulus durations. *Vision Research*, *33*, 1785–1794.

Edwards, M., & Badcock, D. R. (1995). Global motion perception: no interaction between the first- and second-order pathways. *Vision Research*, *35*, 2589–2602.

Foster, K. H., Gaska, J. P., Nagler, M., & Pollen, D. A. (1985). Spatial and temporal frequency selectivity of neurones in visual cortical areas V1 and V2 of the macaque monkey. *Journal of Physiology*, *365*, 331–363.

Grzywacz, N. M., & Yuille, A. L. (1990). A model for the estimate of local image velocity by cells in the visual cortex. *Proceedings of the Royal Society of London*, *B239*, 129–161.

Heeger, D. J. (1987). Model for the extraction of image flow. *Journal of the Optical Society of America A*, *4*, 1455–1471.

Holub, R. A., & Morton-Gibson, M. (1981). Response of visual cortical neurons of the cat to moving sinusoidal gratings: response contrast functions and spatiotemporal interactions. *Journal of Neurophysiology*, *46*, 1244–1259.

Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, *195*, 215–243.

Johnston, A., & Benton, C. P. (1997). Speed discrimination thresholds for first- and second-order bars and edges. *Vision Research*, *37*, 2217–2226.

Johnston, A., McOwan, P. W., & Buxton, H. (1992). A computational model of the analysis of some first-order and second-order motion patterns by simple and complex cells. *Proceedings of the Royal Society of London*, *B250*, 297–306.

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.

Ledgeway, T., & Smith, A. T. (1997). Changes in perceived speed following adaptation to first-order and second-order motion. *Vision Research*, *37*, 215–224.

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

Nishida, S., Ledgeway, T., & Edwards, M. (1997). Dual multiple-scale processing for motion in the human visual system. *Vision Research*, *37*, 2685–2698.

Simoncelli, E. P., & Heeger, D. J. (1998). A model of neuronal responses in visual area MT. *Vision Research*, *38*, 743–761.

Smith, A. T., Hess, R. F., & Baker, C. L. (1994). Direction identification thresholds for second-order motion in central and peripheral vision. *Journal of the Optical Society of America A*, *11*, 506–514.

Smith, A. T., & Ledgeway, T. (1997). Separate detection of moving luminance and contrast modulations: Fact or artifact? *Vision Research*, *37*, 45–62.

Smith, A. T., Snowden, R. J., & Milne, A. B. (1994). Is global motion really based on spatial integration of local motion signals? *Vision Research*, *34*, 2425–2430.

Watamaniuk, S. N. J. (1993). Ideal observer for discrimination of the global direction of dynamic random-rot stimuli. *Journal of the Optical Society of America A*, *10*, 16–28.

Watamaniuk, S. N. J., & Duchon, A. (1992). The human visual system averages speed information. *Vision Research*, *32*, 931–941.

Watamaniuk, S. N. J., & Sekuler, R. (1992). Temporal and spatial integration in dynamic random-dot stimuli. *Vision Research*, *32*, 2341–2347.

Watamaniuk, S. N. J., Sekuler, R., & Williams, D. W. (1989). Direction perception in complex dynamic displays: the integration of direction information. *Vision Research*, *29*, 47–59.

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

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

Williams, P., & Phillips, P. (1987). Cooperative phenomena in the perception of motion direction. *Journal of the Optical Society of America A*, *4*, 878–885.

Williams, D., Phillips, G., & Sekuler, R. (1986). Hysteresis in the perception of motion direction as evidence for neural cooperativity. *Nature*, *324*, 253–254.

Williams, D. W., & Sekuler, R. (1984). Coherent global motion percepts from stochastic local motions. *Vision Research*, *24*, 55–62.

Williams, D, Tweten, S., & Sekuler, R. (1991). Using metamers to explore motion perception. *Vision Research*, *31*, 275–286.

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

Zanker, J. M. (1993). Theta motion: a paradoxical stimulus to explore higher order motion extraction. *Vision Research*, *33*, 553–569.

Zohary, E., Scase, M. O., & Braddick, O. J. (1996). Integration across directions in dynamic random dot displays: Vector summation or winner take all? *Vision Research*, *36*, 2321–2331.