PERGAMON

Vision Research 40 (2000) 833–841

Vision Research

www.elsevier.com/locate/visres

The role of relative motion computation in 'direction repulsion'

Steven C. Dakin [a,*], Isabelle Mareschal [b]

[a] *Department of Visual Science, Institute of Ophthalmology, 11-43 Bath Street, London EC1V 9EL, UK*
[b] *Center for Neural Science, New York University, 4 Washington Place, New York, NY 10003, USA*

Received 11 March 1997; received in revised form 3 December 1997

Abstract

When two sets of intermixed dots move in different directions the perceived direction of each is considerably shifted [Marshak & Sekuler (1979). *Science*, *205*, 1399–1401; Mather & Moulden, (1980). *Quarterly Journal of Experimental Psychology*, *32*, 325–333)]. This phenomenon has been attributed to 'repulsive' interactions between channels tuned to different directions of motion. However, we report that it is not only the relative direction, but also the density and speed of the sets, which determines the magnitude of the apparent shift. These results are difficult to reconcile with the notion of 'repulsive' interactions, and we describe an alternative, functionally motivated explanation. In the natural environment, observed motion results from objects moving over background surfaces that may themselves be mobile. Disentanglement of motion signals therefore necessitates a computation of *relative motion*. We propose that the phenomenon of 'direction repulsion' results from a deliberate adjustment of observed motion to compensate for an inferred source of 'background' motion. A simple scheme to do this subtracts the weighted vector-sum of all motion signals from observed motion. This relative motion computation quantitatively predicts the observed effects of the density of dot sets on perceived direction. The effects of speed cannot be reconciled with the scheme as it stands, but this could be due to the model's failure to consider the effect of temporal frequency on the effective contrast of the sets.
© 2000 Elsevier Science Ltd. All rights reserved.

Keywords: Motion; Random dot displays; Direction repulsion; Vector-sum; Relative motion

1. Introduction

There is both psychophysical and physiological evidence that the visual system processes moving images using a bank of filters tuned to specific spatial and temporal frequencies (Campbell & Robson, 1968; Tolhurst & Movshon, 1975). Opponent models posit that interactions between filters tuned to opposite directions can explain many visual illusions based on selective adaptation (Snowden, 1989). Non-opponent models propose that there are no specific interactions between channels tuned to opposite directions (Levinson & Sekuler, 1976). An example where non-opponent interactions are believed to underlie perceptual distortions, of simultaneously presented moving targets, is 'direction repulsion' (Marshak & Sekuler, 1979; Mather & Moulden, 1980; Kim & Wilson, 1996). When two intermingled sets of dots move in different directions the perceived direction of each is shifted away from the direction of the other set. The size of this shift has been shown to be inversely proportional to the angular difference between the two sets of dots, and can be as large as 25° (Marshak & Sekuler, 1979).

Previous studies have explained this illusion in terms of channels that are broadly tuned to direction of motion (Marshak & Sekuler, 1979; Mather & Moulden, 1980). These models propose that the distribution of responses over the channels is shifted away from a veridical representation of motion as a consequence of channels systematically reducing the activity of their neighbours in proportion to the similarity of their directional selectivity (Raymond, 1993). Similarly, Wilson and Kim (1994) describe a model of motion coherence and transparency that uses competitive inhibition between directionally tuned channels to compute the vector-sum of local motion signals. 'Direction repulsion' arises as a consequence of this mechanism when the components of bi-directional motion are widely separated (for plaids, by more than 108°). Be-

---

* Corresponding author. Fax: +44-171-608-6846.
E-mail address: s.dakin@ucl.ac.uk (S.C. Dakin)

0042-6989/00/$ - see front matter © 2000 Elsevier Science Ltd. All rights reserved.
PII: S0042-6989(99)00226-6

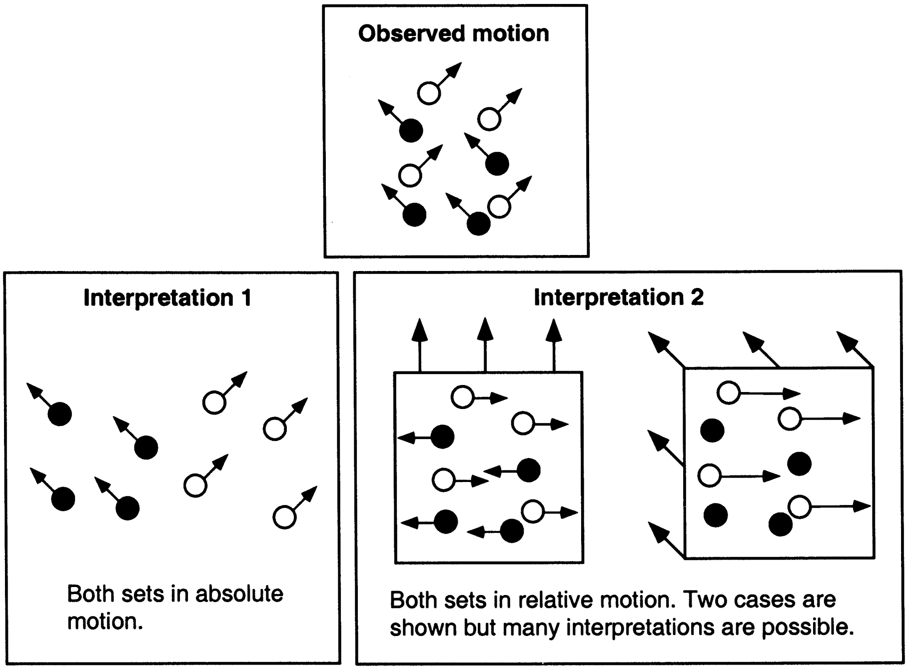

Fig. 1. A given set of motion signals is attributable either to absolute movement of objects, or to the result of their motion combined with a moving background.

cause this model is directed primarily towards explaining plaid motion, at smaller directional differences it predicts perception of coherent motion in the vector-sum direction. In contrast, using narrow-band 'dot' patterns, we will show that large illusory shifts in direction occur reliably when the directions of the two motion signals differ by as little as 45°.

Notice that, for the models described, these shifts of perceived direction are not computed deliberately but are 'errors' arising from the interaction of direction selective channels. We propose an alternative explanation; that shifts result from the purposeful computation of the *relative motion* of the two sets to some inferred background motion[1]. Fig. 1 illustrates the ambiguity that the visual system must resolve in order to calculate the direction and speed of an object. Motion signals may result from absolute motion, such as a bird flying across the sky, or from the motion of an object with respect to a moving background, such as a dog swimming across a fast moving river. Because our own movement interacts with motion signals arising from objects, relative motion calculation is probably frequently required (Johansson, 1973, 1975).

Because all of the interpretations shown in Fig. 1 are consistent with the observed motion the problem of deriving true underlying motion (in the absence of other visual cues to object identity) is fundamentally under-constrained. The results reported below suggest that the visual system does not interpret complex, intermixed motion signals as the result of absolute motion (Interpretation 1), but as the result of relative motion (Interpretation 2). Relative motion computation has previously been conceived of as an attempt by the visual system to minimise the complexity of its' interpretation of motion stimuli (Johansson, 1975). Might not 'direction repulsion' be the result of a simplified interpretation of these dot displays based on relative motion computation?

Following Johansson (1975) we propose that there are two stages in relative motion computation: estimation of the background motion signal, and adjustment of observed motion to compensate for the effects of the background. We consider two constraints on computation of the background. First it should weight the contribution of motion signals by their strength. Second, it should reflect in some way a 'common component' of observed motion (Johansson, 1975). Consider Interpretation 2 shown in Fig. 1. The rightmost part shows the case where the motion of one set (the solid dots) is taken to be entirely due to background motion (i.e. this set is interpreted as being static on a mobile background). This means that the direction of the other set (the open dots) contributes nothing to the computation of background motion. The leftmost case on Interpretation 2, in contrast, equally weights the contribution of both sets to the computation of the background motion signal. Given that the visual system

---

[1] Note that we are not arguing against a central role for inhibitory processing per se, but against an approach where inhibition is present *only* to sharpen channel tuning and where repulsion arises as an undesirable biproduct of that inhibition.

has no a priori evidence to the contrary, it would be sensible to weight the contribution of each set in proportion to it's density. Therefore, it would seem appropriate to utilise a population statistic, such as the vector-sum, to infer background motion. The vector-sum both weights the contribution of a motion signal by its speed and leads to an estimate which in some sense reflects a 'common component' of all observed motion.

There are other reasons for using the vector-sum. Individual motion signals are sensitive to noise, and might be unrepresentative of true local motion. The use of individual motion signals would lead to local 'pairwise' computation of relative motion, which would be highly sensitive to variations in local motion signal density. As a consequence of these problems any estimate of relative motion might frequently be locally inconsistent. A population statistic, on the other hand, is resistant to noise and leads to a more economical and locally consistent representation of motion. The vector-sum is a computationally simple, unique statistic. It is also known that statistics such as the vector-sum, and more generally the *mean*, are employed by the visual system in a variety of domains: to represent motion in plaids (Wilson & Kim, 1994), and in random dot patterns when the directional signals are of equal strength (Zohary, Scase & Braddick, 1996), to represent position in the spatio-luminance domain (Westheimer & McKee, 1977; Watt & Morgan, 1983; Badcock & Westheimer, 1985), as an estimate of local orientation in texture (Dakin & Watt, 1997), etc. Finally, the vector-sum can be simply computed using a channel-based representation of motion and, by interpolating between broadly-tuned channels, serves to greatly increase the accuracy of the system (compared to, say, a measurement based on peak channel activity).

To summarise, using the vector-sum motion as the inferred background signal has a number of computational advantages and embodies the constraints that motion signals should contribute in proportion to their speed, and that the background signal should reflect a 'common component' of observed motion.

The second stage of the relative motion computation is the estimation of the underlying motion that, when combined with the inferred background signal would lead to the observed motion. This amounts to simple addition of the vector-sum, rotated through 180°, to all motion signals.

This model generates clear predictions related to the 'direction repulsion' phenomenon. For stimuli composed of two sets of moving dots, not only the direction of the sets, but also their relative density and speed alter the vector-sum, and therefore the illusory shift, in a predictable manner. To test this, we measured the perceived direction of a target set of dots, intermixed with a distractor set moving in a different direction. We then observed the effect of systematically varying the relative density and speed of the two sets, on this estimate.

2. General methods

2.1. Apparatus

A Macintosh 7500 microcomputer generated stimuli and recorded subjects' responses. Stimuli were displayed on a Nanao Flexscan 6500 monochrome monitor (75 Hz frame refresh rate). The luminance of the display was linearised using measurements from a UDT photometer in conjunction with routines from the VideoToolbox package (Pelli, 1997). The screen was viewed binocularly at a distance of 95 cm and had a mean background luminance of 23.6 cd/m^2.

2.2. Stimuli

Stimuli consisted of ten frame movie sequences, of 500 ms duration, showing two moving sets of spatially band-pass elements. Elements were isotropic Laplacian-of-Gaussians:

$$\nabla^2 G(x, y, \sigma) = \frac{1}{\sigma^2}\left(1 - \frac{x^2 + y^2}{\sigma^2}\right) e^{-(x^2+y^2)/2\sigma^2} \qquad (1)$$

with $\sigma = 0.1°$. Micro-patterns had a peak spatial frequency of approximately 3.8 c/deg.

The directions of the 'target' (the set subjects attended to) and 'distractor' sets were symmetrically oriented around vertical, and differed by 45, 90 or 135°. Stimuli were contained in 384 pixel square images (subtending 7.5 deg^2) viewed through a circular aperture of radius 192 pixels.

In two separate conditions the relative density and the relative speeds of the two sets were systematically varied. In the relative density condition, the target: distractor densities tested were 500:0, 455:45, 400:100, 363:167, 250:250, 167:363, 100:400 and 45:455 elements. Target and distractor moved at 1.25°/s. In the second condition we tested relative speeds of 0.9:2.5, 1.25:2.5, 1.8:2.5, 2.5:2.5, 3.5:2.5, 5.0:2.5, and 7.1:2.5°/s. Stimuli were composed of two sets each containing 250 elements.

2.3. Procedure

Subjects (the authors, who are both corrected-to-normal myopes) were presented with a movie and instructed to 'indicate the direction of the more rightwardly moving set' by adjusting the orientation of a line at the bottom of the display. They adjusted line orientation using four keys on the computer keyboard (two giving 10° steps, and two 2° steps, in the clockwise and anti-clockwise directions). They were free to review the movie sequence, and change their estimate until satisfied.

A block of trials consisted of sixteen adjustments, and blocks from various conditions were randomly interleaved.

3. Results

Fig. 2a and b shows the effect of the relative density of target and distractor sets on perceived direction of the target, when sets differ in direction by 45, 90 and 135° for the two observers. As the dot density of the target set decreases (i.e. moving rightwards on the abscissa), relative to the distractor set, there is a systematic increase in the perceived directional shift of the target. Making the target set sparse can increase the size of the effect, compared with the 1:1 case, by up to 50%. Despite the fact that the magnitude of an illusion is being measured, note the degree of inter-subject consistency. We confirmed that the bias was dependent on *relative* density by obtaining similar results, for a 90° directional difference, using stimuli containing 2000 dots.

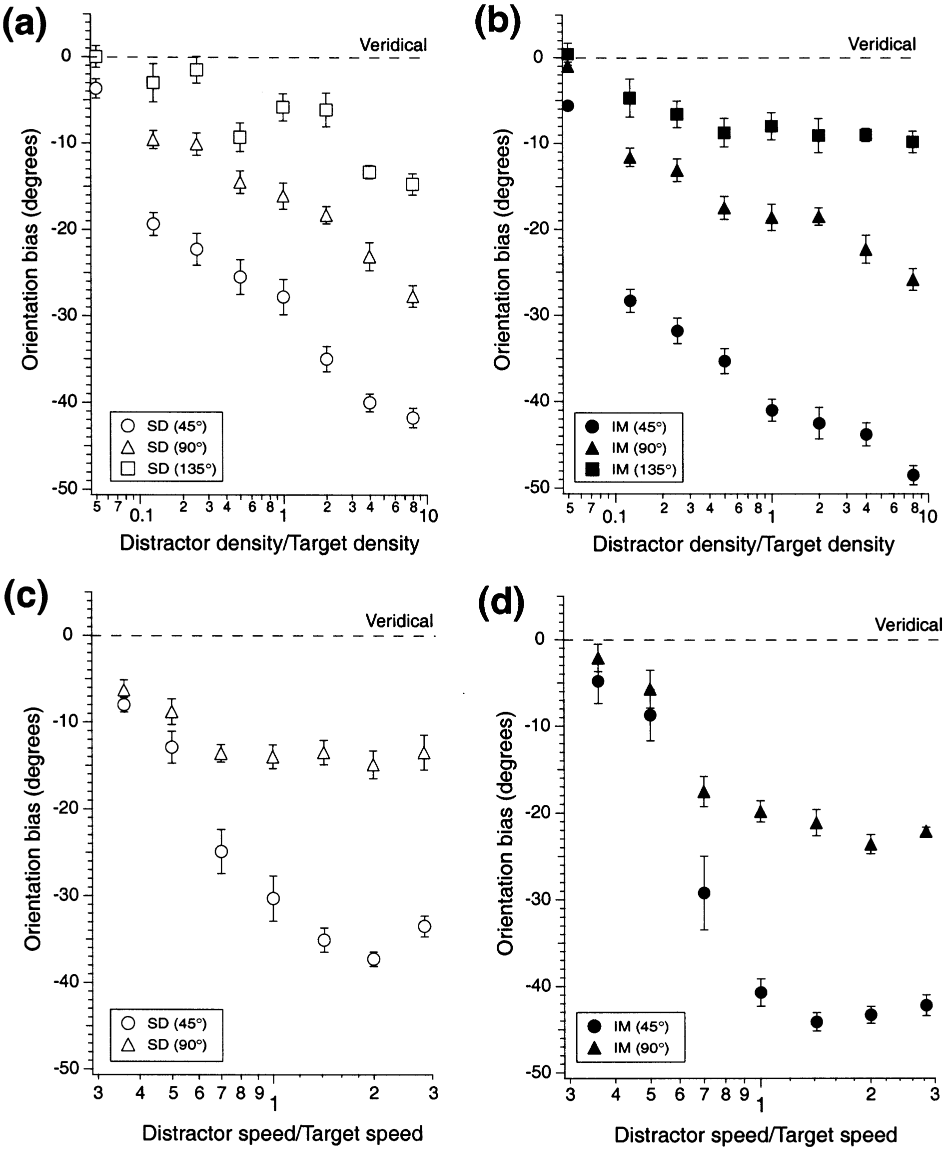

Fig. 2. Bias on the perceived direction of a target dot set, spatially intermixed with a distractor set, as a function of their (a and b) relative densities, and (c and d) speeds. Subjects systematically misjudged the direction of the target as being closer to horizontal (0°). The magnitude of this bias varied in inverse proportion to both the angular difference between sets and the relative density of the target. Slow or sparse distractors shift the target less from its veridical direction. However, the plateau in the data shown in (c) and (d) indicates that speeding up the distractor does not increase the magnitude of the illusion beyond the shift shown in the equal-speed condition.

Vertical shifts in data sets shown in Fig. 2a and b indicate that as the angular difference between the distractor and the target sets is decreased the directional bias is increased. Notice that, contrary to earlier reports examining dot patterns (Marshak & Sekuler, 1979; Mather & Moulden, 1980), but consistent with results for plaids (Kim & Wilson, 1996), we find reliable directional shifts for dot sets with directions differing by more than 90°. It is possible that this inconsistency is because subjects were allowed to review stimuli in our study. Because these perceived shifts are quite small subjects may have simply missed them in previous experiments.

Notice that, for the 1:1 density case for the 45, 90 and 135° angular difference, the average perceived shifts are 30, 18 and 5°, respectively. Since all of these stimuli have a vector-sum direction that is vertical, according to the model described in Section 1 we would predict perception of horizontal motion, i.e. perceived shifts of 67.5, 45 and 22.5°, respectively. Although the ordering of the magnitude of the effect is correct, clearly this model overestimates subjects' perceived shift. We return to this point in Section 4.

In a similar manner to decreasing the density of the target set, decreasing its relative speed (rightwards on the abscissa) results in an increase in its directional shift (Fig. 2c and d). However notice that, as the speed of the distractor exceeds the target's speed, the magnitude of the shift plateaus at the degree of shift obtained in the equal-speed condition. This pattern of results is in agreement with preliminary reports by Marshak and Sekuler (1979). They report that sets of matched speed give the greatest shift and that speed differences attenuate the effect (although they do not identify which set was moving faster, we must presume it was the target set). These results are contrary to those described for multiaperture bars by Kim and Wilson (1996), who report that shifts are a linear function of the distractor set's speed. However these authors only considered ratios of distractor-to-target speeds up to 1.67, leaving open the possibility that the perceived direction of motion would not be influenced further at the higher ratios (up to 2.8) that we tested.

Why don't increasingly fast distractor sets induce increasingly great repulsion? It is known that reduction in the distractor contrast reduces the degree of direction repulsion for multiaperture bars (Kim & Wilson, 1996). It is possible that the speeds we tested were such that the effective contrast of the distractor set was greatly attenuated, consequently reducing its contribution to the computation of the mean. It is informative to note that the speed of the reference set (2.5°/s) corresponds to a temporal frequency of 9.5 Hz, which is close to the peak of the human temporal contrast sensitivity function. At spatial frequencies around 3.8 c/deg, this curve shows a rapid decline with increasing temporal frequency but a shallow roll-off at lower temporal frequencies. For example, it is known that contrast sensitivity is reduced by a factor of ten at a temporal frequency of around 30 Hz (de Lange, 1958), which approximately corresponds to the velocity of the fastest set used (3.8 c/deg moving at 7.1°/s). Thus, our data do not allow us to discount the possibility that reduced visibility of distractor sets leads to the plateau observed in Fig. 2c and d.

4. Modelling

4.1. Estimation of the vector-sum

As discussed above, a computation simply using the vector-sum as an inferred background motion signal will consistently overestimate the magnitude of the illusory shift. Recall from Fig. 1 that any two motion signals cannot unambiguously specify a particular form of relative motion. There are no 'correct' and 'incorrect' interpretations of our stimuli in terms of relative motion and it appears that the influence of the background motion signal is limited.

We reflect this in our model by using a weighted vector-sum estimate. The modified two-stage scheme for deriving relative motion is illustrated in Fig. 3a. In the first stage the vector-sum direction and speed of all motion signals is computed. The direction component, and a fixed proportion of the speed component, serve as a background motion signal to which observed motion is assumed to be relative. The vector-sum direction of motion is defined as:

$$\Theta = \tan^{-1}\left(\frac{\sum_{i=1}^{N} V_i \sin(\theta_i)}{\sum_{i=1}^{N} V_i \cos(\theta_i)}\right) \tag{2}$$

where θ_i is the direction, and V_i the velocity of the ith component. For computation of the speed, a weighting parameter, k, is used:

$$S = k\sqrt{\sum_{i=1}^{N} V_i^2 \sin^2(\theta_i) + \sum_{i=1}^{N} V_i^2 \cos^2(\theta_i)} \tag{3}$$

Justification for this type of computation comes from experiments examining human estimation of direction and speed in plaid patterns (Ferrera & Wilson, 1991; Wilson & Kim, 1994). They report that the perceived direction of motion is predicted by a vector-sum calculation of the motion signals present but that the perceived speed is inversely proportional to the angular difference between plaid components. In fact the weighting parameters used in the simulations described are (within the range tested) in broad agreement with the speed judgement errors reported in Ferrera and Wilson (1991).

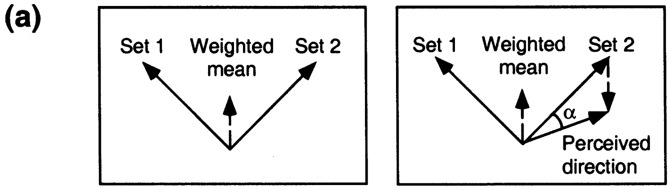

Stage 1. Estimate background motion, using mean direction and weighted mean velocity.

Stage 2. Adjust motion signals

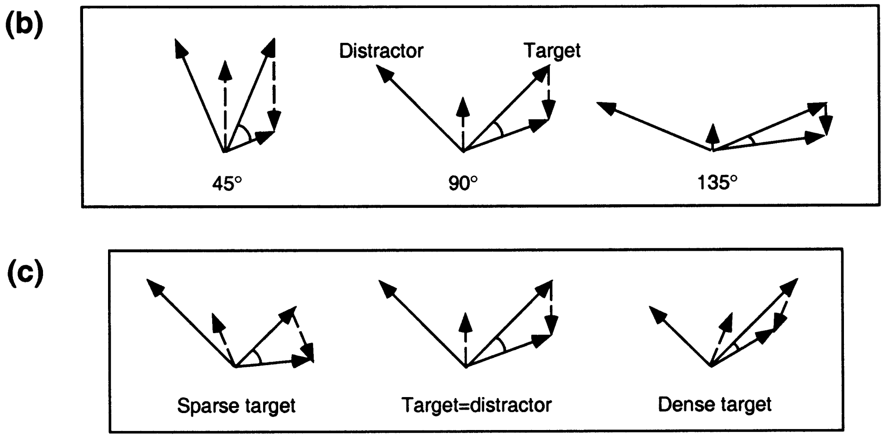

Fig. 3. Modified model for computing relative motion, incorporating speed weighting. (b) Increasing angular separation of the two component dot sets decreases the speed of the vector-sum and consequently the size of perceived shift in the target direction. (c) Similarly, increasing target density shifts the direction of the vector-sum towards the target, producing smaller perceived shifts.

4.2. Inference of object motion

The second stage is the inference of an object's motion that based on the speed and direction of the background, would lead to the observed motion signal. As stated above, this consists simply of adding the inferred background signal, rotated through 180°, to all observed motion. The motion that results from this operation is attributed to an object. Fig. 3b illustrates how this type of scheme can account for 'direction repulsion' effects. As the directions of two motion signals diverge, the velocity of the vector-sum decreases. Consequently, the contribution of the background to the second stage of the scheme decreases, and this leads to an decrease in α (the difference between perceived and veridical direction). Fig. 3c illustrates how density changes can lead to a similar effect. As the density of the target increases, the vector-sum (background) motion vector shifts towards the target direction and consequently has a diminishing effect on the perceived direction of the target.

Note that the model as presented not only makes specific predictions about the perceived angular shift, but also about the velocity at which the shifted set moves. This is reflected in the length of the vector representing the shifted set. For example, in Fig. 3b, note that as the angle separating the two sets increases, the velocity of the shifted set approaches veridical velocity (compare the 135° case to the 45° case). It is the authors' impression that perceived velocity of the shifted set broadly agrees with this prediction (i.e. perceived speed of the target varies in proportion to the angular difference of sets' directions). We are presently investigating perceived speed in these displays in greater detail.

4.3. Modelling results

The solid line in Fig. 4a is the predicted perceived direction obtained using the vector-sum direction and weighted speed as a background. The speed weightings used were fitted and are $k = 0.70$, 0.55 and 0.30 for angular differences of 45, 90 and 135°, respectively. Subjects' performance is in good agreement with the predictions from this model. Vertical shifts in data sets shown in Fig. 4a indicate that as the angular difference between the distractor and the target sets is decreased the directional bias is increased. This dependence is also quantitatively predicted by the relative motion computation.

Fig. 4b shows the predicted perceived direction (solid lines) as a function of the relative speed of two equally dense sets. Subjects' perceived shifts are quantitatively predictable by the relative motion model (solid line in Fig. 4b) for cases when the target is slower than, or of equal-speed to the distractor (data to the left of 1.0 on the abscissa). Speed weighting parameters were 0.72 and 0.45 (SCD), and 0.80 and 0.55 (IM), for 45 and 90° angular differences respectively. Decreasing the relative target speed shifts the background towards the distractor and increases the difference between the relative and absolute motion of the target set. Psychophysically, however, increasing the relative speed of the distractor set, beyond the equal-speed case, does not produce the predicted increase in the magnitude of the effect. Instead the size of perceived shift plateaus at the 1:1 speed ratio. Once the distractor set is travelling at the same speed as the target its effect is constant. This constraint can be included in the relative motion computation as a ceiling on the influence of the distractor set speed (modified predictions are indicated by the dashed lines, Fig. 4b). However, whether this represents a real constraint on the visual systems computation of relative motion, or whether this plateau is a consequence of the reduced visibility of distractor sets at high temporal frequencies, cannot be determined form our data.

4.4. Control experiments

To further establish the use of the weighted vector-sum as background motion signal we conducted a series of control experiments, using stimuli composed of three dot sets. In agreement with the relative motion model we find that when one dot set is static, it appears to move in a direction opposite to the vector-sum motion. The static dot set does not selectively stimulate any direction selective channels, so it is difficult to see how repulsive mechanisms operating between them could produce this percept. When the three dot sets move in directions separated by 120°, vector-sum motion has zero magnitude and, accordingly, we report veridical perception of the direction of all sets. We also observe, for a variety of combinations of set directions, that when one of the three sets moves in the vector-sum direction, its direction is not shifted. A final condition (suggested by an anonymous referee) involved two dot sets of equal-speed and density separated by 45°, and a third set moving in the direction opposite to the vector-sum. We measured the perceived shift in direction of one of the equal-speed pair as a function of the speed of the third set. A repulsive inhibitory mechanism predicts no change in the size of the illusory shift with the speed of the third set while we predict a systematic reduction of the illusion as the speed of the third set reduces the vector-sum (of the entire configuration) to zero. Fig. 5 plots data from this condition for one subject showing that the illusion is indeed reduced and approaches zero, as the vector-sum approaches zero (at 2.3°/s on the abscissa). All of these results support the notion of recalibration of observed motion by the weighted vector-sum. Given the relatively prolonged exposure duration of our stimuli (500 ms) we performed a final experiment to control for the contribution of tracking eye movements. We tripled the speed of our displays (reducing presentation times to approximately 160 ms) and found a similar magnitude of effect (for the equal density condition at 45 and 90° directional differences).

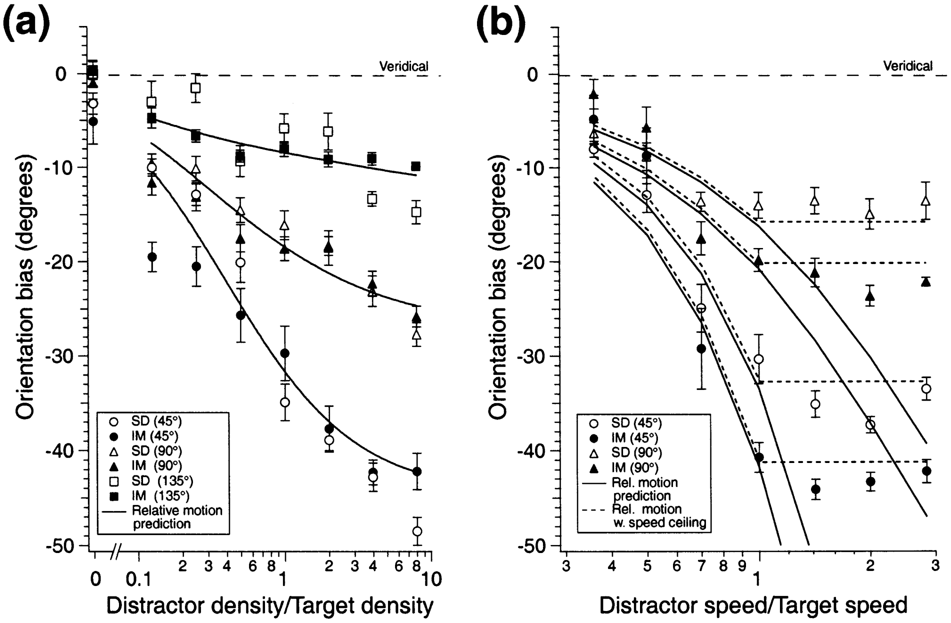

Fig. 4. Effect of (a) density, and (b) speed on predictions from the relative motion scheme compared to psychophysical data from Fig. 2. (a) Predictions for the density condition match subjects' performance well. (b) Contrary to the predictions of a pure relative motion computation (solid lines), there appears to be a ceiling on the influence that the speed of the distractor can have on the magnitude of the effect. Including this ceiling in the relative motion computation improves predictions (dashed lines).

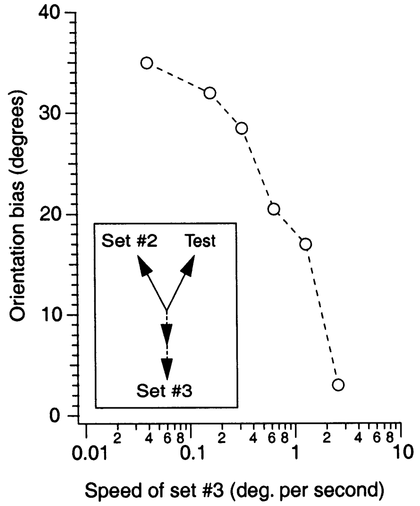

Fig. 5. Perceived directional bias on one of three equal density moving dot sets (test) as a function of the speed of set 3. Sets 1 and 2 moved at 1.25°/s and at 67.5° (test set) and 112.5° (set 2). Set 3 moved in the direction opposite to the vector-sum of the test set and set 2, i.e. 270° and at a number of speeds. Notice the systematic reduction in the size of the illusion as the speed of set 3 increases.

5. Discussion

To summarise, we propose that when we attend to a target set of moving dots, intermixed with a second dot set, shifts in the targets' perceived direction are attributable to relative motion computation. This computation uses the vector-sum direction and weighted speed of all motion signals as an assumed background. It then adjusts observed motion signals to take into account the contribution of this background to the observed motion. To test this model, we systematically varied the relative speed and density of target and distractor, and found the following results. First, and contrary to previous reports (Marshak & Sekuler, 1979; Mather & Moulden, 1980), when the directions of target and distractor differ by as much as 135°, we still observe reliable, although smaller, shifts in the perceived direction of the target. Second, increasing the density of the distractor set increases the magnitude of the illusion (in accord with the effect of increasing contrast of the distractor component of a multiaperture bar stimulus; Kim & Wilson, 1996). Third, we find a small effect of slow distractors, which increases with distractor speed and then plateaus when the distractor is travelling at the same speed as the target.

The first two results are in close accord with computation of the relative motion of the target set. The third result suggests that the speed of components is not treated in the way supposed by our model.

As is often the case the conditions under which the model fails prove more interesting than those under which it succeeds. Why should the influence of the second set plateau as its speed increases? The first possibility is discussed above; fast distractor sets may simply exert less influence on the mean because of observers' poor contrast sensitivity at high temporal frequencies. This could be tested by extending the model to employ the output of mechanisms tuned for temporal frequency. The second possibility is that this effect arises from the way in which the visual system makes the correction for background motion (the 'inference of object motion' stage in our model). Within a local vector representation of direction and velocity it is easy to see how an inferred background motion can be simply subtracted from all observed motion to compute relative motion. However, how could the visual system 'correct' a channel coded representation of the two motion signals represented in Fig. 3b? It seems likely that motion in oblique directions (in the example shown) is coded not only by the activity of channels tuned to the appropriate directions and speeds but by the distributed activity of channels tuned to all speeds and all directions. If that were the case, then the correction stage proposed could be implemented by selectively inhibiting these 'speed' and 'direction' components (in the case of Fig. 3b, the upward motion component) without the possibility of exciting any further channels. To put it another way, in order to implement the true effect of high speed distractor sets on the background estimate would require introduction of new motion components into the channel based code which is not possible within a scheme based exclusively on inhibition.

These findings have implications for the development of models of visual motion processing. We propose that relative motion computation is a fundamental and (normally) effortless component of motion processing. Its primary goal being to resolve the conflict that arises when contradictory motion signals occur in the same space, the visual system may apply a variety of heuristics to decide if two motion signals do arise from relative motion. For motion in natural scenes, cues such as depth from stereo might provide enough information to reliably determine when motion is relative and when it is not. When the relative motion computation is fundamentally under-constrained, as is the case of two intermixed sets of dots, the operation of this system leads to large 'errors' in the perceived direction of motion components. The accuracy of this system with unambiguous cues to relative motion remains to be determined.

Acknowledgements

This work was conducted while the authors were at the McGill Vision Research Unit. SCD was supported by an MRC grant (# MT 108-18 to R.F. Hess), and I.M. by FCAR fellowship. We thank Frans Verstraten, Robert Hess, Curtis Baker, Joshua Solomon and Michael Morgan for valuable discussion of this work.

References

Badcock, D., & Westheimer, G. (1985). Spatial localization and hyperacuity: the centre/surround localization contribution function has two substrates. *Vision Research*, *25*, 1259–1267.

Campbell, F., & Robson, J. (1968). Application of Fourier analysis to the visibility of gratings. *Journal of Physiology*, *197*, 551–566.

Dakin, S. C., & Watt, R. J. (1997). The computation of orientation statistics from visual texture. *Vision Research*, *37*, 3181–3192.

de Lange, H. (1958). Research into the dynamic nature of the human fovea-cortex systems with intermittent and modulated light. II Phase shift in brightness and delay in color perception. *Journal of the Optical Society of America*, *48*, 784–789.

Ferrera, V. P., & Wilson, H. R. (1991). Perceived speed of moving two-dimensional patterns. *Vision Research*, *31*, 877–893.

Johansson, G. (1973). Visual perception of biological motion and a model for its analysis. *Perception and Psychophysics*, *14*, 201–211.

Johansson, G. (1975). Visual motion perception. *Scientific American*, *232*, 76–89.

Kim, J., & Wilson, H. R. (1996). Direction repulsion between components in motion transparency. *Vision Research*, *36*, 1177–1187.

Levinson, E., & Sekuler, R. (1976). Adaptation alters perceived direction of motion. *Vision Research*, *16*, 779–781.

Marshak, W., & Sekuler, R. (1979). Mutual repulsion between moving visual targets. *Science*, *205*, 1399–1401.

Mather, G., & Moulden, B. (1980). A simultaneous shift in apparent direction: further evidence for a 'distribution-shift' model of direction coding. *Quarterly Journal of Experimental Psychology*, *32*, 325–333.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming number into movies. *Spatial Vision*, *10*, 437–442.

Raymond, J. E. (1993). Movement direction analysers: independence and bandwidth. *Vision Research*, *33*, 765–775.

Snowden, R. J. (1989). Motion in orthogonal directions are mutually suppressive. *Journal of the Optical Society of America*, *6*, 1096–1101.

Tolhurst, D. J., & Movshon, J. A. (1975). Spatial and temporal contrast sensitivity of striate cortical neurons. *Nature*, *257*, 674–675.

Watt, R. J., & Morgan, M. J. (1983). Mechanisms responsible for the assessment of visual location: theory and evidence. *Vision Research*, *23*, 97–109.

Westheimer, G., & McKee, S. P. (1977). Integration regions for visual hyperacuity. *Vision Research*, *17*, 89–93.

Wilson, H. R., & Kim, J. (1994). Perceived motion in the vector sum direction. *Visual Neuroscience*, *6*, 1205–1220.

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