Speed tuning in elementary motion detectors of the correlation type

J.M. Zanker[1], **M.V. Srinivasan**[1], **M. Egelhaaf**[2]

[1] Centre for Visual Sciences, RSBS, Australian National University, GPO Box 475, Canberra, ACT 2601, Australia
[2] Lehrstuhl für Neurobiologie, Universität Bielefeld, Postfach 100131, D-33501 Bielefeld, Germany

Abstract. A prominent model of visual motion detection is the so-called correlation or Reichardt detector. Whereas this model can account for many properties of motion vision, from humans to insects (review, Borst and Egelhaaf 1989), it has been commonly assumed that this scheme of motion detection is not well suited to the measurement of image velocity. This is because the commonly used version of the model, which incorporates two unidirectional motion detectors with opposite preferred directions, produces a response which varies not only with the velocity of the image, but also with its spatial structure and contrast. On the other hand, information on image velocity can be crucial in various contexts, and a number of recent behavioural experiments suggest that insects do extract velocity for navigational purposes (review, Srinivasan et al. 1996). Here we show that other versions of the correlation model, which consists of a single unidirectional motion detector or incorporates two oppositely directed detectors with unequal sensitivities, produce responses which vary with image speed and display tuning curves that are substantially independent of the spatial structure of the image. This surprising feature suggests simple strategies of reducing ambiguities in the estimation of speed by using components of neural hardware that are already known to exist in the visual system.

1 Introduction

The simplest way in which a visual system can determine how fast and in what direction an object travels would be to determine how long the object needs to cover the distance between two given points. Indeed, 'feature-tracking' mechanisms have been proposed along these lines to explain motion vision (Braddick 1980; Ullman 1983) and are sometimes used in machine vision

(Murray and Buxton 1990). The disadvantage of such schemes, however, is that they need to identify the object, or features within it, before carrying out the tracking. In another class of models, this problem is avoided by using information on local spatiotemporal changes of intensity to measure velocity. In one subclass of these 'intensity-based' models, the so-called gradient models, image speed is determined by computing the ratio of the local temporal and spatial gradients of intensity (Fennema and Thompson 1975; Limb and Murphy 1975). Gradient models have the property that they measure the speed of an image independently of its spatial structure. Modifications of this scheme have been proposed for measuring image velocity in two dimensions and for overcoming problems associated with low and sparsely distributed contrasts (Hildreth and Koch 1987; Johnston et al. 1992; Srinivasan 1990).

Another subclass of the intensity-based models is the so-called Reichardt or correlation detector, which extracts a motion signal from the spatiotemporal correlations that are present in the moving image (Reichardt 1957). In this model, which has been very successful in describing motion sensitivity in animal vision from insects to primates (review, Borst and Egelhaaf 1989), the signal from one input unit A is delayed or temporally low-pass filtered and multiplied with that from a neighbouring input B (see Fig. 1). As a consequence of this structure, the model produces a strong output only when the image moves in the direction (A → B). The standard design of the correlation model, the so-called balanced version, subtracts the output of this network from that of an anti-symmetric one which multiplies the signal from A with the filtered signal from B. The resulting scheme produces a positive response when the image moves in the direction (A → B), and a negative response when the image moves in the direction (B → A). Setting these two anti-symmetrical 'half-detectors' in opposition has the virtue of cancelling out, by subtraction, the direction-unspecific components of the response and makes the output of the overall network highly selective for the direction of motion (Borst and Egelhaaf 1990). Direction-unspecific response compo-
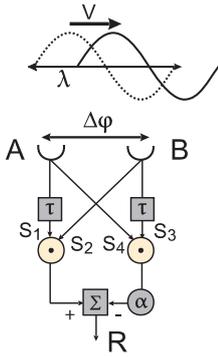
Fig. 1. Basic structure of the 'correlation' model of motion detection, generalised here to the case of variable balance between the half-detectors. *Circles with dots* denote non-linear interactions, namely the multiplication of the two input signals. *Square boxes* denote the temporal filters, namely pure time delays or temporal low-pass filters as indicated by τ. The symbol Σ represents a subtraction stage (summation with opposite signs), and α indicates a variable gain factor (coefficient of balance) which determines the degree to which the outputs of the two half-detectors are balanced. A gain of $\alpha = 1$ represents the well-known fully balanced correlation model in which both half-detectors contribute equally to the output; $\alpha = 0$ represents a fully unbalanced model in which only one half-detector contributes to the response; and intermediate values of α represent partially balanced detectors

nents may arise, for instance, from brightness changes in the visual field which are unrelated to motion. The so-called energy model is a variant of this basic scheme – by combining a set of spatial and temporal filters as they may be implemented in the primate visual cortex, including a so-called opponency stage, it generates very similar properties (Adelson and Bergen 1985; Verri and Poggio 1986).

One of the prominent features of the correlation model of motion detection is that the response depends not only upon the speed of the image, but also upon its spatial structure and contrast (Reichardt 1961; Varjú 1959; Varjú and Reichardt 1967; Buchner 1984). In particular, when stimulated by moving sinusoidal gratings, the response of the correlation model attains a maximum not at a constant velocity, but at a velocity at which the grating induces a specific temporal frequency. The value of the optimum temporal frequency depends upon the time delay or the time constant of the low-pass filter. Thus, for a model with fixed parameters, gratings with larger spatial wavelengths produce peak responses at larger velocities. This property accounts very well for the characteristics of the optomotor response in insects, for example, and the optokinetic nystagmus in vertebrates (review, Borst and Egelhaaf 1989). However, the correlation model does not readily explain other behaviours that involve the extraction of image velocity independently of the spatial structure, such as human speed discrimination (Thompson 1984), or some more recently investigated aspects of navigation and depth discrimination in insects (Srinivasan et al. 1996).

Are the newly studied behaviours in insects mediated by an entirely different kind of motion-detecting system, or can they be explained by modifications of the well-

known correlation model? Here we show, analytically and by numerical simulation, that by variation of the balance coefficient in the correlation model, one can produce responses to moving gratings that are more closely related to velocity than to temporal frequency. While this finding does not prove that image velocity is indeed encoded in certain biological systems using such mechanisms, it demonstrates that the measurement can be made, in principle, by networks that require only relatively minor modifications of the neural hardware that is already believed to exist in biological visual systems.

2 Methods

For the simulations, we used a reduced version of a two-dimensional motion detector model (2DMD model; Zanker 1996a, 1997), which comprises a linear array of elementary motion detectors (EMDs) of the correlation type. Each of the 256 EMDs in the array has the same structure, looking in a 'retinotopic' manner at different regions of a 2D stimulus. The EMD operated either directly on the input intensity values, or the DC components were removed by spatial band-pass filtering in the two input lines of the EMD by means of isotropic 2D Difference of Gaussians (DOG) which mimic the spatial transfer characteristics of many visual systems. A first-order low-pass with time constant τ_L was used as a temporal filter. The temporally filtered signal from one location was multiplied with the direct input from a second location, separated horizontally by the sampling interval $\Delta\varphi$. The output of two such anti-symmetric subunits was subtracted with a balance coefficient α varying between 0 (half-detector) and 1 (fully opponent model) (see Fig. 1; for details of the 2DMD, see Zanker 1996b). In all simulations shown here, we fixed the spatial and temporal model parameters at $\Delta\varphi = 4$ pixels and $\tau_L = 2$ simulation steps (frames). Each simulation was run for 16 frames, which was sufficient for the response to attain a steady state.

The stimuli consisted of sequences of 16 two-dimensional images, comprising a vertically oriented sinewave grating that was displaced horizontally between consecutive frames. The contrast of the grating was set to maximum (intensity variation between 1 and 256 arbitrary units). The spatial period λ (measured in image pixels, inverse of spatial frequency) and the grating displacement between consecutive frames, i.e. the speed V (measured in pixels per frame) were treated as variables. In the present context, the output of the EMD array was averaged along its spatial dimension and over the second half of the 16 response frames, to provide a measure of the spatiotemporal average steady-state responses of the EMDs. The response, computed this way, was compared for various stimulus and model parameters. Basically, for any given set of model parameters, we computed responses for 33×33 combinations of λ and V, which were varied independently in 32 logarithmic steps between $\lambda = 4$ and 256 pixels and $V = 1/4$ and 16 pixels per frame. This yielded the 2D response contour

maps shown in Fig. 2. Since the spatial profile of the sinusoidal grating was defined continuously and not in a discrete manner, it was possible to generate sub-pixel displacements (cf. Morgan and Watt 1982).

3 Results

3.1 Theory

Consider a sinusoidal grating with mean intensity I_0, an intensity modulation m and spatial period λ (deg) moving rightward at an angular velocity V (deg/s). This would induce a temporal frequency of V/λ Hz in the inputs, corresponding to an angular frequency $\omega = 2\pi V/\lambda$. The signal in receptor A (cf. Fig. 1) is therefore $I_0 + m \cdot \sin(\omega t)$. Let $\Delta\varphi$ denote the angular separation between the neighbouring input units of the correlation model. The output of the model is generated by subtracting a fraction α of the response of one half-detector from the response of the other, symmetrical half-detector, as shown in Fig.1. Thus, $\alpha = 1.0$ represents a balanced detector, $\alpha = 0.0$ represents a half-detector, and $0 < \alpha < 1.0$ represents a partially balanced detector.

Denoting by $A(\omega)$ and $\Phi(\omega)$ the amplitude and phase response of the temporal filter, respectively, the temporally filtered input signal of the left detector input channel S_1 and the unfiltered signal of the right input signal S_2 (see Fig.1) can be expressed as

$$S_1(t) = A(0) \cdot I_0 + A(\omega) \cdot m \cdot \sin[\omega \cdot t - \Phi(\omega)]$$

$$\text{and} \quad S_2(t) = I_0 + m \cdot \sin\left[\omega \cdot t - \omega \cdot \frac{\Delta\varphi}{V}\right]$$

where A(0) denotes the zero-frequency (DC) gain of the temporal filter. Similar expressions can be written for the filtered signal S_3 of the right detector input and the unfiltered signal S_4 of the left input.

The response of the detector is given by $\overline{S_1 \cdot S_2} - \alpha \cdot \overline{S_3 \cdot S_4}$, where the bar denotes a time average. It can be shown (cf. Egelhaaf et al. 1989) that in the general case this response is

$$R = A(0) \cdot (1 - \alpha) \cdot I_0^2$$
$$+ A(\omega) \cdot \frac{m^2}{2} \cdot \left[\cos\left(\omega \frac{\Delta\varphi}{V} + \Phi(\omega)\right)\right.$$
$$\left. - \alpha \cdot \cos\left(-\omega \frac{\Delta\varphi}{V} + \Phi(\omega)\right)\right] \quad (1)$$

If the model is balanced ($\alpha = 1$) and uses a pure time delay of magnitude ΔT, the amplitude and phase response are given by $A(0) = A(\omega) = 1$ and $\Phi(\omega) = -\Delta T \cdot \omega$, respectively. Inserting this into equation (1) we obtain

$$R = m^2 \cdot \sin\left(\frac{2\pi\Delta\varphi}{\lambda}\right) \cdot \sin\left(\frac{2\pi V \Delta T}{\lambda}\right) \quad (2)$$

It is clear from (2) that the response does not depend upon the velocity of the grating per se, but upon the temporal frequency V/λ that the moving grating induces. The maximum response occurs at a temporal frequency $1/4\Delta T$, regardless of the period of the grating.

If the model contains a first-order low-pass filter with time constant τ (instead of a pure time delay), leading to an amplitude and phase response of $A(\omega) = \frac{1}{\sqrt{1+\tau^2\omega^2}}, A(0) = 1$, and $\Phi(\omega) = -\arctan(\tau\omega)$, respectively, the response is given by

$$R = \frac{m^2}{\sqrt{1 + \frac{4\pi^2 V^2 \tau^2}{\lambda^2}}} \cdot \sin\left(\frac{2\pi\Delta\varphi}{\lambda}\right) \cdot \sin\left[\arctan\left(\frac{2\pi V \tau}{\lambda}\right)\right] \quad (3)$$

Here again, it is clear that the response does not depend upon the velocity of the grating per se, but upon the temporal frequency V/λ that the moving grating induces. It can be shown that the maximum response occurs at a temporal frequency $1/2\pi\tau$, regardless of the period of the grating.

Consider next a half-detector ($\alpha = 0$). With a pure time delay ($A(0) = A(\omega) = 1$ and $\Phi(\omega) = -\Delta T \cdot \omega$), (1) simplifies to

$$R = I_0^2 + \frac{m^2}{2} \cdot \cos\left[\frac{2\pi V}{\lambda}\left(\frac{\Delta\varphi}{V} - \Delta T\right)\right] \quad (4)$$

We see from (4) that this response is no longer dependent on temporal frequency per se. Rather, the response attains a maximum at a constant velocity $V_{\text{opt}} = \Delta\varphi/\Delta T$, regardless of the period of the grating. In other words, this detector is tuned to velocity rather than temporal frequency.

If the half-detector uses a first-order low-pass filter with time constant τ instead of a pure time delay ($A(\omega) = \frac{1}{\sqrt{1+\tau^2\omega^2}}, A(0) = 1$, and $\Phi(\omega) = -\arctan(\tau\omega)$), from (1) we derive

$$R = I_0^2 + \frac{m^2}{2 \cdot \sqrt{1 + \frac{4\pi^2 V^2 \tau^2}{\lambda^2}}} \cdot \cos\left[2\pi\frac{\Delta\varphi}{\lambda} - \arctan\left(\frac{2\pi V \tau}{\lambda}\right)\right] \quad (5)$$

We see that this response is not strictly a function of velocity or temporal frequency. If we consider the variation of the response as a function of velocity, it can be shown that the maximum response occurs at a velocity

$$V_{\text{opt}} = \frac{\lambda}{2\pi\tau} \cdot \tan\left(\frac{\pi\Delta\varphi}{\lambda}\right) \quad (6)$$

The optimum velocity depends on the period λ of the grating. However, the dependence is weak because an increase of λ, for example, would cause the first term to increase and the second to decrease. In particular, if $\pi\Delta\varphi/\lambda \ll 1$, i.e. if the period of the grating is much larger than the separation of the inputs, then (5) can be approximated as

$$V_{\text{opt}} \cong \frac{\Delta\varphi}{2\tau} \quad (7)$$

Under these conditions, the optimum velocity is independent of the grating period.

Consider, finally, a partially balanced detector ($0 < \alpha < 1$). The response of such a detector can be obtained by combining the responses of the left half-detector, R_{HL}, and the right half-detector, R_{RL}. We have already calculated R_{HL} for a detector with a pure time delay (4) and a low-pass filter (5). In each case, the response of the right half-detector, R_{RL}, is obtained from R_{HL} simply by reversing the polarity of the velocity V of the moving grating. The reason is that, since the right half-detector is structurally a mirror image of the left half-detector, the response evoked in the right half-detector by a grating of a given contrast, period and speed, moving from A to B, would be identical to the response evoked in the left half-detector by an identical grating moving from B to A. The overall response of the partially balanced detector is then given by

$$R = R_{HL} - \alpha R_{HR} \qquad (8)$$

Inserting (4) into (8), we see that the response of a partially balanced detector with pure time delays is given by

$$R = (1 - \alpha) \cdot I_0^2 + \frac{m^2}{2} \left\{ \cos\left[\frac{2\pi V}{\lambda}\left(\frac{\Delta\varphi}{V} - \Delta T\right)\right] \right.$$
$$\left. - \alpha \cdot \cos\left[\frac{2\pi V}{\lambda}\left(\frac{\Delta\varphi}{V} + \Delta T\right)\right] \right\} \quad (9)$$

Inserting (5) into (8), the response of a partially balanced detector with first-order low-pass filters is

$$R = (1 - \alpha) \cdot I_0^2$$
$$+ \frac{m^2}{2 \cdot \sqrt{1 + \frac{4\pi^2 V^2 \tau^2}{\lambda^2}}} \cdot \left\{ \cos\left[2\pi\frac{\Delta\varphi}{\lambda} - \arctan\left(\frac{2\pi V \tau}{\lambda}\right)\right] \right.$$
$$\left. - \alpha \cdot \cos\left[2\pi\frac{\Delta\varphi}{\lambda} + \arctan\left(\frac{2\pi V \tau}{\lambda}\right)\right] \right\} \quad (10)$$

Naturally, (9) and (10) reduce to the expressions (2) and (3), respectively, for the case of fully balanced detectors ($\alpha = 1$), and to expressions (4) and (5), respectively, for the case of half-detectors ($\alpha = 0$).

Summarising the above results, we see that balanced detectors ($\alpha = 1$) produce a response that depends critically on the temporal frequency of periodic gratings, while half-detectors ($\alpha = 0$) produce a response that critically depends on velocity. We expect that partially balanced detectors ($\alpha = 0.5$) will exhibit a behaviour that lies between these two extremes. That is, their responses should exhibit speed tuning curves that depend weakly on the grating period. In the next section we investigate the properties of fully balanced detectors, partially balanced detectors and half-detectors by means of computer simulations.

3.2 Simulation

To get a general impression of the behaviour of motion detectors as revealed by computer simulations, the complete wavelength-speed profiles of the average detector responses are shown in Fig. 2 as contour plots. The fully opponent model (Fig. 2a) generates the characteristic response profile of a correlation detector. For grating displacements that are larger than half the period between successive frames (top right corner of plot), the stimulus appears to move in the opposite

Fig. 2a–c. Response profiles of three motion detector models (sketched in the *insets*) for the variation of spatial frequency and speed of moving gratings. Response contours are shown in 10% steps of the maximum response, *broken lines* indicate regions of negative response. **a** In the fully balanced correlation detector, two anti-symmetrical subunits are subtracted from each other. **b** The 'half-detector' is a correlation detector stripped down to only one (asymmetrical) subunit. **c** In this simple implementation of a gradient detector, after spatially bandpass filtering of the input signal, the temporal gradient (approximated by differences between successive simulation steps) is divided by the spatial gradient (approximated by the difference between neighbouring points, to which a small constant ε is added)

balanced
correlation
detector

a

correlation
half-
detector

b

gradient
type
mechanism

c
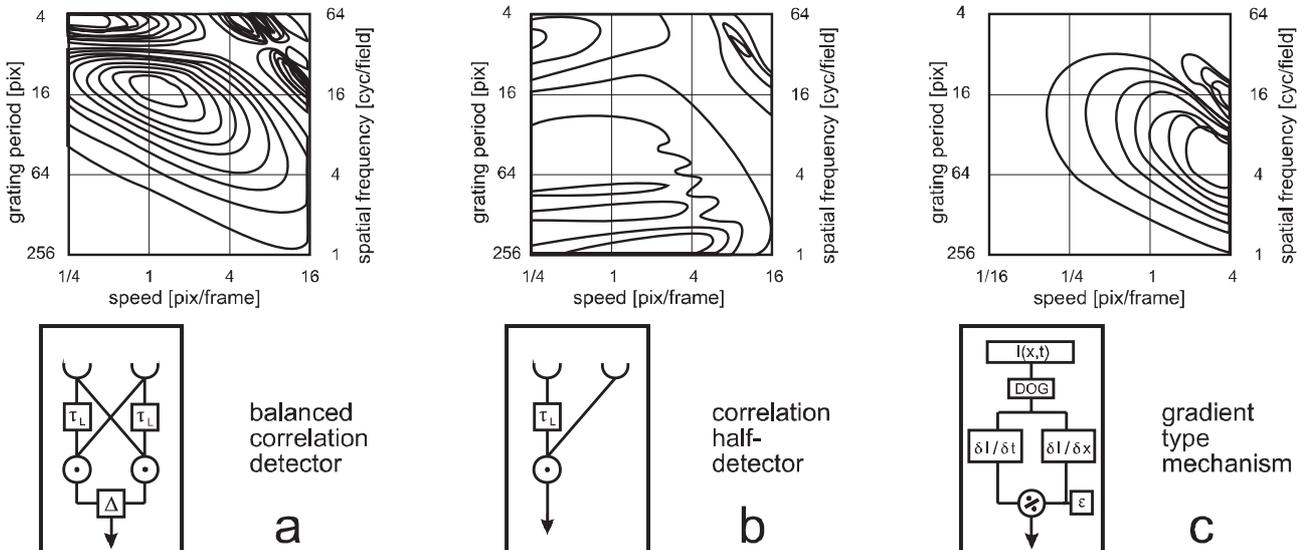
direction, as a consequence of which the response is inverted (dotted contours indicate negative response regions). This phenomenon, known as 'temporal aliasing', is an artefact of the temporally discrete nature of the simulation and would not occur in a biological network using continuous time-dependent signals. A second region of negative response, often referred to as 'geometric interference' or 'spatial aliasing', can be observed for fine gratings (top left area of plot). When the period λ is smaller than twice the sampling distance of the detector, the 'Nyquist limit' $2\ \Delta\varphi$, the detector response is inverted (e.g. Götz 1964). By introducing a bandpass filter into the input lines as mentioned in the Methods, this effect can be minimised, with significant responses being more or less limited to a range of gratings with a spatial period between 8 and 32 pixels, for instance (data not shown). In the dominating region of positive response, the contours are slanted and skewed, leading to an oblique overall orientation of the profile. As a consequence, the maximum response occurs at higher speeds when λ is increased. This characteristic pattern of results, which is less pronounced though qualitatively the same for a detector with a spatial bandpass in the input (data not shown), will be shown in more detail in Figs. 3 and 4 below.

The half-detector (Fig. 2b) generates a response profile that shares some similarities with the fully opponent model, but is dominated by an overall positive response that is largely independent of the stimulus pa-

rameters. This is a consequence of the fact that in this model the effects of DC components in the input signals are no longer cancelled by subtracting the two anti-symmetrical units [cf. (5)]. Accordingly, stimulus aliasing inherent to the apparent motion stimulus and geometric interference do not generate negative responses, but lead to smaller positive responses. When spatial bandpass filters are added to the front end of the model to remove the DC components from the input, the usual sign inversions are observed (data not shown). The regions of response reduction/inversion due to geometric interference are larger than those observed for the fully balanced model. This further supports the view that a perfectly balanced model, incorporating two anti-symmetrical half-detectors, possesses increased directional selectivity (cf. Borst and Egelhaaf 1990).

The major observation in this contour plot, however, is that the lobes of positive response are not skewed, but rather stacked vertically with respect to the speed axis. As a result, the response maxima are largely independent of the structure of the stimulus pattern, as will be analysed in more detail in Fig. 3 and 4.

For purposes of illustration and immediate comparison (Fig. 2c), the same simulations were carried out for a simple implementation of the gradient scheme. For this model, after spatial bandpass filtering of the input with a DOG filter, the difference of the signals of two consecutive frames is divided by the difference between signals from neighbouring points, after adding a small constant ($\varepsilon = 0.01$) to the denominator. In this contour plot it is clear that, apart from perturbations related to stimulus aliasing, the response increases approximately linearly with speed, largely independently of the structure of the stimulus pattern. This shows that this implementation, at least within the limits tested here, meets the expectations of the gradient scheme (Fennema and Thompson 1975; Horn and Schunck 1981; Limb and Murphy 1975) and that one could indeed retrieve a largely unambiguous speed signal with such a model. In particular, there is no decrease of the response after reaching a maximum speed, a property that is qualitatively different from the correlation-type models. We also note that, despite the fact that a spatial bandpass filter is operating on the input images, the spatial frequency tuning of the gradient scheme is rather broad,

Fig. 3a–c. Speed characteristics of three variants of the correlation model. Each panel shows the average motion detector response to rightward motion of gratings as function of speed V (*abscissa*) for three different grating periods λ (indicated by different symbols). All data are normalised to the maximum response that is elicited by each grating of a given period. Note that the data are plotted on different scales for the ordinate. The continuous functions are derived from (10), whereas the data points are the results of computer simulations of a motion detector array. The results from these two procedures differ for large pattern displacements where the apparent motion stimuli of the digital simulations begin to alias. **a** For a fully opponent correlation detector (balance coefficient $\alpha = 1.0$), the speed optimum is shifted to higher values for larger grating periods. **b** For an 'imperfectly' balanced detector ($\alpha = 0.5$), the speeds eliciting the largest response come closer together for the various gratings. **c** For the correlation half-detector ($\alpha = 0$), the optimum speed and the speed tuning curves are virtually identical for a range of grating periods
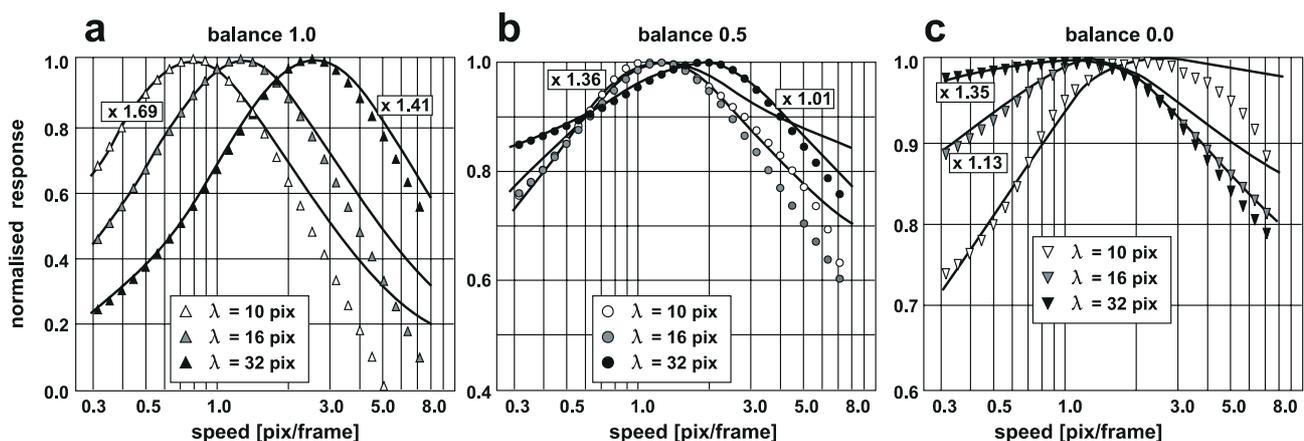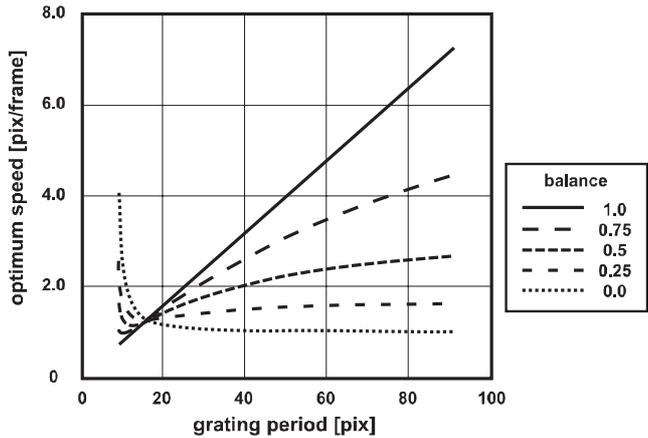
Fig. 4. Variation of optimum speed V_{opt} with grating period λ, for five variants of the correlation model with different balance coefficients α (indicated by different line types), as derived from (10). With decreasing balance coefficient, the overall orientation of the lines changes from diagonal (*continuous line*) to horizontal (*dotted line*), indicating a change from temporal frequency tuning to speed tuning. This general pattern is overlaid by systematic deviations in V_{opt} that are predicted by the theory for small values of α when pattern wavelengths approach the sampling interval $\Delta\varphi$

unlike the situation in correlation-type models with bandpass filters in their input lines.

Figure 3 examines more closely the dependence of the average detector response on grating speed, with the grating period λ treated as simulation parameter. To ease the comparison between the tuning curves for different λ, the results are normalised with respect to the maximum response for a given grating (the scaling factors are shown as labels next to each curve in the figure). Furthermore, the negative response contributions due to aliasing are excluded from the simulation data points shown in Fig. 3, which at large speeds differ for that reason from the graphs resulting from (10) that are shown as continuous lines. In the perfectly balanced correlation detector, the response curve shifts rightward along the speed axis as the grating period λ is increased, as expected (Fig. 3a). The optimum speed is clearly proportional to λ, leading to a constant optimum temporal frequency $f_{opt} = v_{opt}/\lambda$. This means that the response is a maximum when a fixed number of stripes pass the visual field of the detector per unit time. An imperfectly balanced model ($\alpha = 0.5$) behaves differently (Fig. 3b). Apart from generally elevated response levels (note that the ordinate scales vary in Fig. 3), the speed characteristics for the three gratings approach each other for this unbalanced correlation detector. As a consequence, the maximum response occurs at rather similar speeds for gratings of different periods. This trend becomes more prominent when the balance coefficient α is set to zero. In such a half-detector (Fig. 3c), the grating speed that elicits the strongest response is virtually identical for the two larger values of λ. The grating with the smallest period, close to the interference limit $2\Delta\varphi$, elicits the maximum response at a slightly higher speed. Under these extreme stimulus conditions, the behaviour of the standard, fully balanced correlation detector is even inverted.

To examine further the influence of stimulus temporal frequency and speed, we plot in Fig. 4 the optimum speed as a function of the grating period, as derived by using (10). Five detector variants that differ in the balance coefficient (indicated by different symbols in the figure) are compared, the extreme cases being the fully balanced detector (balance coefficient $\alpha = 1.0$, continuous line) and the half-detector ($\alpha = 0.0$, dotted line). For the fully balanced detector, all significant data points lie on a diagonal, indicating that the optimum velocity V_{opt} is strictly proportional to grating period λ. In other words, the response peaks at a constant temporal frequency. For the half-detector, the optimum velocity remains approximately constant for grating periods ranging between $\lambda = 20$ pixels and $\lambda = 80$ pixels, but increase when the interference limit $\lambda = 8$ pixels is approached, where all responses are virtually zero. This generally 'inverse' behaviour for very small wavelengths confirms the observation made with the simulations shown in Fig. 3c, and can be expected from the restriction for approximating (6) by (7). It is clear, however, that the half-detector does not exhibit the constant f_{opt} despite the variation of λ – that is characteristic of the standard, fully balanced correlation detector (this behaviour is referred to as temporal frequency tuning, without implying that the detector response depends unambiguously on the temporal frequency of a moving pattern). As predicted in the 'theory' section, for large grating periods this detector exhibits constant V_{opt} for the variation of λ (this behaviour is referred to as speed tuning, without implying that the detector response is determined unambiguously by pattern speed). The three other data sets in Fig. 4 illustrate that correlation models with intermediate values of the balance coefficient α display intermediate characteristics. Thus, in this way, a continuum between temporal frequency tuning and speed tuning can be achieved by varying the degree to which two opponent half-detectors are balanced.

4 Discussion

It has commonly been assumed that motion detectors based on the correlation scheme always respond best to a characteristic temporal frequency rather than a characteristic speed. Indeed, this property has been used as a characteristic 'fingerprint' to probe for the existence of such schemes in animal vision. Here we have shown that this kind of temporal frequency tuning is exhibited only by a perfectly balanced correlation model. If we consider imperfections in the subtraction stage that lead to an imbalance between the two anti-symmetric subunits, and eventually to a half-detector, the characteristics of a correlation detector are gradually transformed from temporal frequency tuning to a kind of speed tuning. The half-detector, which represents the case of extreme imbalance, responds best to a preferred speed, largely independent of the spatial period of the stimulus grating, but at the expense of directional selectivity (Egelhaaf et al. 1989). This has two important consequences. Firstly, a signal closely related to speed could

easily be made available to the visual system even if movement detection is mediated by a mechanism that relies on correlation. It is easy to imagine that the speed and temporal frequency of a moving grating are coded simultaneously, and can be accessed simply by using the signals appearing before and after the subtraction stage, respectively, of a fully opponent correlation model. The important point is that in such a scheme of double 'readout', with the restrictions discussed in the next paragraph, there is no need to invoke additional mechanisms of higher complexity to account for the measurement of speed. Secondly, this means that on the basis of experimental findings indicating that some speed signal is used to control behaviour, for instance, it will be difficult to reject the hypothesis that the underlying mechanism is of the correlation type.

Although a half-detector is speed-tuned, it would not be capable of measuring the speed of a grating unambiguously. One reason is that the bell-shaped response-versus-speed curve implies that a response of a given magnitude could have been evoked by one of two different speeds. This ambiguity can be eliminated by using an additional half-detector with a different speed tuning (i.e. a different optimum speed). The relative strengths of the responses of the two detectors would then provide unambiguous information on speed. The contrast of the grating is another factor that confounds the ability of a single detector to measure speed unambiguously. This is because, for a grating of a given period and speed, the response increases as the square of the contrast [cf. (1)], and changes in detector output could therefore be due to changes in speed or to changes in contrast. This ambiguity can again be eliminated by comparing the responses of two or more detectors with different speed tunings. The ratio of their responses, for example, would provide an indication of speed that is largely independent of contrast. Furthermore, it has to be remembered that all these considerations refer to the steady-state response; the outcome may be different under dynamic conditions.

Apart from gradient-based detectors, other schemes have been proposed for estimating the velocity of image motion independently of spatial structure and contrast. Usually, this involves a combination of several channels tuned to different spatial and temporal frequencies (e.g. Heeger 1987; Zanker and Braddick, in preparation), which also resolve other ambiguities, such as those caused by variations in contrast. In this context, two independent studies have used half-detectors of the correlation type to estimate speed (Glünder 1990; Snippe and Koenderink 1994). Glünder (1990) presents a model for estimating 2D velocity. This uses an array of half-detectors with sampling bases ($\Delta\varphi$) of various sizes and various directions, and estimates the speed and direction of image motion in terms of the centroid of the distribution of detector response in this parameter space. Snippe and Koenderink (1994) discuss the velocity tuning properties of half-detectors and show how the tuning can be sharpened by using multiple delays or time constants, together with multiple sampling bases. Contrast invariance is achieved by using detectors tuned to dif-

ferent velocities and adopting a winner-takes-all principle. However, both of these studies retrieve velocity from an ensemble of correlation detectors and do not consider the properties of the single element in any detail. Our study, in contrast, focuses on comparing the properties of a half-detector with those of a full detector or a partially balanced detector.

Electrophysiological studies of the large-field, directionally selective neurones in the visual system of the fly have often revealed the presence of second harmonics (signals at twice the temporal frequency of the moving grating) in their response (Egelhaaf et al. 1989). This has been interpreted as evidence that the underlying motion detectors are not perfectly balanced, and the amplitude of the second harmonic has been used to infer the degree of imbalance. In the light of our findings, it would be interesting to investigate experimentally the temporal frequency and velocity tuning of such units, and to examine whether the dependence of optimum speed on spatial structure is related to the prevalence of the second harmonic signals in the response of those neurones.

Recent behavioural studies on freely flying insects suggest the presence of motion-detecting mechanisms that retrieve the speed of the image on the eye, rather independently of the spatial structure of the image. For example, honeybees flying through a tunnel maintain equidistance from the side walls by balancing the velocities of the images on the two eyes, even when the walls carry gratings of different spatial periods (Kirchner and Srinivasan 1989; Srinivasan et al. 1991). Bees can use image motion to distinguish objects at different distances, independently of their absolute size (Lehrer et al. 1988; Kirchner and Srinivasan 1989); they can also distinguish objects of different absolute sizes, irrespective of their ranges (Horridge et al. 1992). These findings imply that bees can disentangle size cues from velocity cues, and measure the two parameters independently. Bees can also be trained to find a food reward at a constant distance from a landmark, independently of the size of the landmark (Lehrer and Collett 1994). This again suggests that the motion of the image of the landmark is being measured independently of the size of the image on the retina. Bees flying through long tunnels measure distance travelled by integrating, over time, the image velocity that they experience en route (Srinivasan et al. 1996). This computation is again robust to changes of the texture lining the tunnels. All of these behaviours require image speed to be measured independently of spatial structure. It is possible that these behaviours are mediated by motion-detecting mechanisms that are entirely different in nature from the well-known correlation scheme, as suggested by some recent studies (e.g. Douglass and Strausfeld 1996; Srinivasan et al. 1991, 1993; Srinivasan and Zhang 1997). Another possibility, however, is that sensitivity to velocity is achieved by using half-detectors or imperfectly balanced correlation mechanisms, as described here.

References

Adelson EH, Bergen JR (1985) Spatiotemporal energy models for the perception of motion. J Opt Soc Am 2:284–299

Borst A, Egelhaaf M (1989) Principles of visual motion detection. Trends Neurosci 12:297–306

Borst A, Egelhaaf M (1990) Direction selectivity of blowfly motion-sensitive neurons is computed in a two-stage process. Proc Natl Acad Sci USA 87:9363–9367

Braddick OJ (1980) Low-level and high-level processes in apparent motion. Philos Trans R Soc Lond Biol 290:137–151

Buchner E (1984) Behavioural analysis of spatial vision in insects. In: Ali MA (ed) Photoreception and vision in invertebrates. Plenum, New York, pp 561–621

Douglass JK, Strausfeld NJ (1996) Visual motion-detection circuits in flies: parallel direction and non-direction-sensitive pathways between the medulla and the lobula plate. J Neurosci 16:4551–4562

Egelhaaf M, Borst A, Reichardt W (1989) Computational structure of a biological motion-detection system as revealed by local detector analysis in the fly's nervous system. J Opt Soc Am 6:1070–1087

Fennema CL, Thompson WB (1975) Velocity determination in scenes containing several moving objects. Comput Graph Image Proc 9:301–315

Glünder H (1990) Correlative velocity estimation: visual motion analysis, independent of object form, in arrays of velocity-tuned bilocal detectors. J Opt Soc Am 7:255–263

Götz KG (1964) Optomotorische Untersuchung des visuellen Systems einiger Augenmutanten der Fruchtfliege Drosophila. Kybernetiks 2:77–92

Heeger DJ (1987) Model for the extraction of image flow. J Opt Soc Am 4:1455–1471

Hildreth E-C, Koch C (1987) The analysis of visual motion: from computational theory to neuronal mechanisms. Annu Rev Neurosci 10:477–533

Horn BKP, Schunck BG (1981) Determining optical flow. Art Int 17:185–203

Horridge GA, Zhang SW, Lehrer M (1992) Bees can combine range and visual angle to estimate absolute size. Philos Trans R Soc Lond [Biol] 337:49–57

Johnston A, McOwan PW, Buxton H (1992) A computational model of the analysis of some first-order and second-order motion patterns by simple and complex cells. Proc R Soc Lond [Biol] 250:297–306

Kirchner WH, Srinivasan MV (1989) Freely flying honeybees use image motion to estimate object distance. Naturwissenschaften 76:281–282

Lehrer M, Collett TS (1994) Approaching and departing bees learn different cues to the distance of a landmark. J Comp Physiol [A] 175:171–177

Lehrer M, Srinivasan MV, Zhang SW, Horridge GA (1988) Motion cues provide the bee's visual world with a third dimension. Nature 332:356–357

Limb JO, Murphy JA (1975) Estimating the velocity of moving images in television signals. Comput Graph Image Proc 4:311–327

Morgan MJ, Watt RJ (1982) Mechanisms of interpolation in human spatial vision. Nature 299:553–555

Murray DW, Buxton BF (1990) Experiments in the machine interpretation of visual motion. MIT Press, Cambridge, Mass.

Reichardt W (1957) Autokorrelations-Auswertung als Funktionsprinzip des Zentralnervensystems. Z Naturforsch 12b:448–457

Reichardt W (1961) Autocorrelation, a principle for the evaluation of sensory information by the central nervous system. In: Rosenblith WA (ed) Sensory communication. MIT Press, Cambridge, Mass., pp 303–317

Snippe HP, Koenderink JJ (1994) Extraction of optical velocity by use of multi input Reichardt detectors. J Opt Soc Am 11:1222–1236

Srinivasan MV (1990) Generalized gradient schemes for the measurement of two-dimensional image motion. Biol Cybern 63:421–431

Srinivasan MV, Zhang SW (1997) Visual control of honeybee flight. In: Lehrer M (ed) Orientation and communication in arthropods. Birkhauser, Basel, pp 95–113

Srinivasan MV, Lehrer M, Kirchner WH, Zhang SW (1991) Range perception through apparent image speed in freely flying honeybees. Visual Neurosci 6:519–535

Srinivasan MV, Zhang SW, Chandrashekara K (1993) Evidence for two distinct movement-detecting mechanisms in insect vision. Naturwissenschaften 80:38–41

Srinivasan MV, Zhang SW, Lehrer M, Collett TS (1996) Honeybee navigation en route to the goal: visual flight control and odometry. J Exp Biol 199:237–244

Thompson P (1984) The coding of velocity of movement in the human visual system. Vision Res 24:41–45

Ullman S (1983) The measurement of visual motion – computational considerations and some neurophysiological implications. Trends Neurosci 6:177–179

Varjú D (1959) Optomotorische Reaktionen auf die Bewegung periodischer Helligkeitsmuster (Anwendung der Systemtheorie auf Experimente am Rüsselkäfer Chlorophanus viridis). Z Naturforsch 14b:724–735

Varjú D, Reichardt W (1967) Übertragungseigenschaften im Auswertesystem für das Bewegungssehen II. Z Naturforsch 22b:1343–1351

Verri A, Poggio T (1986) Motion field and optical flow: qualitative properties. A I Memo 917:1–32

Zanker JM (1996a) Looking at the output of two-dimensional motion detector arrays. Invest Ophth Vis Sci 37:S 743

Zanker JM (1996b) On the elementary mechanism underlying secondary motion processing. Proc R Soc Lond [Biol] 351:1725–1736

Zanker JM (1997) Is facilitation responsible for the "motion induction" effect? Vision Res. 37:1953–1959