Vision Res. Vol. 30, No. 7, pp. 1049–1067, 1990
Printed in Great Britain.

0042-6989/90 $3.00 + 0.00

EFFECT OF CONTRAST ON THE PERCEIVED DIRECTION OF A MOVING PLAID

L. S. STONE, A. B. WATSON and J. B. MULLIGAN
Vision Group, Human Interface Research Branch, NASA, Ames Research Center, MS 239-3, Moffett Field, CA 94035, U.S.A.

(*Received* 11 *September* 1989; *in revised form* 17 *November* 1989)

Abstract—We performed a series of experiments examining the effect of contrast on the perception of moving plaids. This was done to test the hypothesis put forth by Adelson and Movshon (1982) that the human visual system determines the direction of a moving plaid in a two-staged process: decomposition into component motion followed by application of the intersection of constraints rule. Although there is recent evidence that the first tenet of their hypothesis is correct, i.e. that plaid motion is initially decomposed into the motion of the individual grating components (Movshon, Adelson, Gizzi & Newsome, 1986; Welch, 1989), the nature of the second-stage combination rule has not as yet been established. We found that when the gratings within the plaid are of different contrast, the perceived direction is not predicted by the intersection of constraints rule. There is a strong (up to 20 deg) bias in the direction of the higher-contrast grating. A revised model, which incorporates a contrast-dependent weighting of perceived grating speed as observed for 1-D patterns (Thompson, 1982), can quantitatively predict most of our results. We discuss our results in the context of various models of human visual motion processing and of physiological responses of neurons in the primate visual system.

Motion Contrast Plaids Direction Aperture problem

INTRODUCTION

Deducing the direction of motion of a pattern as a whole from the motion of oriented components within that pattern is a challenge for all models of human visual motion processing. Adelson and Movshon (1982) studied this problem using moving *plaids*, the sum of two drifting gratings of different orientations. They proposed that the human visual system determines the direction of a moving pattern using a two-step procedure: first, the velocities of oriented components within the pattern are estimated, then at a later stage they are recombined to calculate the motion of the pattern as a whole. Their hypothesis was formulated to explain their psychophysical finding that, in order for two components to cohere (to move together as a plaid), the gratings must be similar in spatial frequency. They concluded from this finding that the human visual system analyzes plaid motion by first decomposing it into the motion of the grating components (Fig. 1A). They suggested that this decomposition is the natural consequence of having orientation and spatial-frequency tuned sensors at the front end of the system (for a review, see DeValois & DeValois, 1980). They also proposed that, at the second stage (Fig. 1B), the component velocities are recombined using the intersection of perpendicular constraints (Fennema & Thompson, 1979) to yield a measure of the motion of the plaid as a whole. When plaid motion is plotted in velocity space (Fig. 1B), the motion of each grating component within the plaid is ambiguous, consistent with a family of velocities lying along a constraint line (thick lines). Adelson and Movshon proposed that plaid velocity is recovered as the unique vector defined by the intersection of both constraint lines. The plaid-velocity vector is thus consistent with the rigid motion of both of the individual gratings and is, therefore, a measure of the motion of the coherent plaid. The lack of coherence for gratings of widely differing spatial frequencies was explained by assuming that the second stage only combines information from sensors with similar spatial-frequency tuning.

They found support for their hypothesis in the discovery of two types of motion-sensitive neurons in the monkey visual cortex: one sensitive to component motion and one at a higher level within the cortex, sensitive to the motion of the plaid as a whole (Movshon et al., 1986). Furthermore, a recent study has found that speed discrimination for moving-plaid stimuli is consistent with the two-staged approach

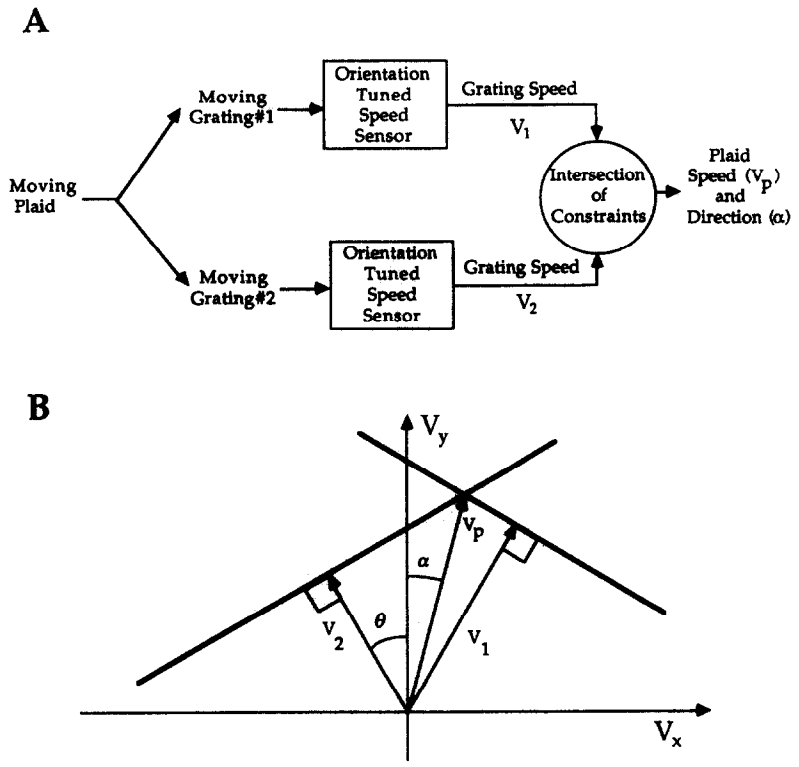

Fig. 1. The Adelson–Movshon model. This illustrates the basic two-stage framework where plaid motion is decomposed into the motion of the grating components, then reconstructed at a second stage. (B) The intersection of constraints rule is demonstrated by showing the motion of a plaid in velocity space. The direction of plaid motion (α), a function of the speeds of the grating components (V_1 and V_2) and of the plaid angle (θ) is given by equation (4).

that Adelson and Movshon proposed (Welch, 1989). However, the second-stage recombination rule proposed by Adelson and Movshon has recently been challenged (Ferrera & Wilson, 1988, 1989).

In this study, we extend the investigation of how the brain determines the direction of motion. In particular, we examine the effect of contrast on the perceived direction of a moving plaid. It has been shown that the perceived speed of a single grating is a function of contrast (Thompson, 1982). At temporal frequencies below 8 Hz, a low-contrast grating appears to move more slowly than a standard higher-contrast grating moving at the same physical speed. If this contrast-dependent distortion in the perceived speed of the components is passed on to the second stage of the model in Fig. 1, then a significant contrast-dependent distortion in the perceived direction of motion of the plaid as a whole should result. We confirmed this prediction: making the contrast of the individual components within the plaid unequal results in errors in judgments of direction. The perceived direction of motion can differ by up to 20 deg from that predicted by the model in Fig. 1.

We, therefore, propose a revised model that incorporates Thompson's finding as a contrast-dependent distortion of component speed. Since the proposed contrast dependence in the revised model is a function of the component contrast in threshold units, we first measured detection threshold of moving gratings in the presence of a moving grating mask of different orientation, the geometric arrangement being the same as with the plaid stimuli. Simulations of the revised model show that, in most cases, if the contrast-distorted estimate of grating speed rather than the true grating speed is the input to the second stage of processing, then the observed errors in perceived direction can be quantitatively predicted. Preliminary results have been presented elsewhere (Stone, Mulligan & Watson, 1988a, b).

GENERAL METHODS

The stimulus used in this study was a vignetted plaid, the sum of two sinusoidal grat-

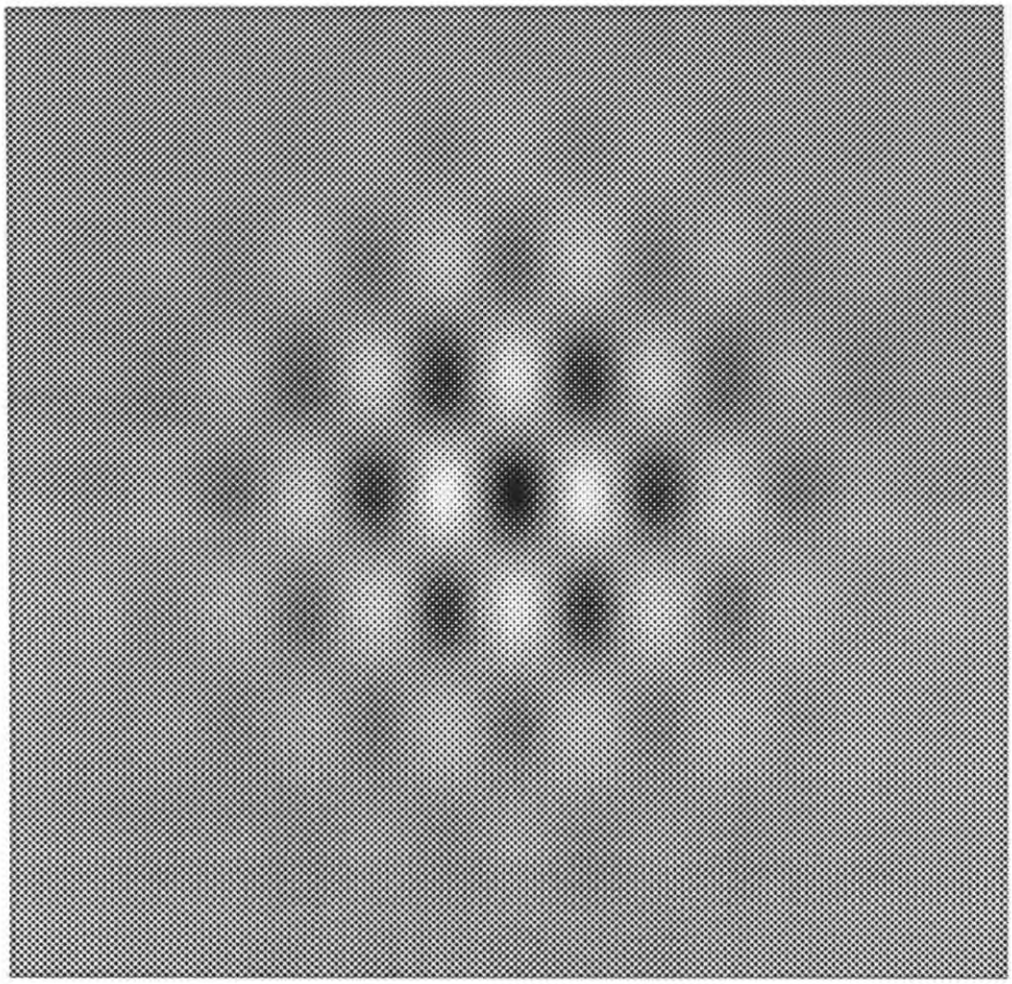

Fig. 2. The standard plaid stimulus. The two gratings were 1.5 c/deg, oriented 60 deg symmetrically from vertical, and viewed through a Gaussian window. The half-toning for this figure was not that used for the actual stimulus. For a detailed discussion of the half-toning used to display our stimuli see Mulligan and Stone (1989).

ings of different orientations viewed through a two-dimensional Gaussian window. Figure 2 shows an example of such a stimulus. We generated moving plaids on a Mitsubishi 19-in high-resolution monochrome monitor (model M-6950) using an Adage RDS 3000 image display system. The luminance output of the monitor was calibrated and corrected for its gamma nonlinearity using a lookup table procedure described elsewhere (Watson, Nielson, Poirson, Fitzhugh, Bilson, Nguyen & Ahumada, 1986). A detailed account of the animation procedure that was used to generate moving plaids can be found in Mulligan and Stone (1989). Briefly, the plaid stimulus was a 512 × 512, 8-bit/pixel image created using both locally developed programs and the HIPS image-processing software package (Landy, Cohen & Sperling, 1984). First, four 2-D sinusoidal gratings were generated (sine- and cosine-phase components of gratings with two different orientations symmetric with respect to the vertical axis). These four images were then multiplied by a two-dimensional Gaussian (x and y standard deviations of 90.5 pixels). This procedure eliminated the sharp edges at the boundaries of the stimulus. The images were than halftoned using a modified error-diffusion method (Floyd & Steinberg, 1975; Mulligan, 1986). The resulting four bit-mapped images were then loaded into the four lower-order bit-planes. A 3 × 3 pixel white fixation cross was drawn into a fifth bit-plane in the center of the image. The remaining three bit-planes were blank. The image could be loaded into the framebuffer within a few seconds. Then, by varying the color lookup table on a frame-by-frame basis (at 60 Hz), we modulated the contrast of the sine- and cosine-phase components of each grating in temporal quadrature so that they appeared as a single drifting grating. In this way, we had complete control over the speed and contrast of both gratings within the plaid without having to load new images into the framebuffer. Furthermore, the initial spatial phases of the grating within the plaid were randomized so that position cues could not be used to assess motion.

There were small but measurable departures from linearity of spatial summation in our display monitor which conflict with one of the basic assumptions underlying halftoning techniques. However, using a technique described elsewhere (Mulligan & Stone, 1989), for a stimulus of 40% total contrast, we estimated the contrast of the largest artifact to be less than 0.2% and, in particular, those artifacts harmonically related to the stimulus were even smaller.

The standard plaid stimulus consisted of two 1.5 c/deg gratings whose normal vectors were oriented symmetrically ±60 deg from the vertical axis (Fig. 2). We defined the *plaid angle* as the angle between the normal vector defining each grating and the bisecting axis or half the angle between the normal vectors (θ Fig. 1B). It was, therefore, 60 deg for the standard plaid. This arbitrary definition was chosen because it simplifies the equations presented below. The grating contrasts were 10% each. For a pair of sinusoidal gratings, the *total contrast* is simply the sum of the grating contrasts or 20%. The speed of the coherent plaid was held constant at 2 deg/sec. In some experiments, the spatial frequency was 0.75 or 3.0 c/deg, the total contrast was 5, 10 or 40%, and the plaid speed was increased to 6.0 deg/sec.

Subjects viewed the screen binocularly through natural pupils from a distance of 273 cm. This distance made the image subtend 5.4 × 5.4 deg (20 pixels/cm) and made the high-frequency halftoning noise invisible except at 40% total contrast. The mean luminance of the image was 100 cd/m^2. The stimulus presentation lasted 300 msec. The contrast rose with a Gaussian time course (standard deviation of 0.71 frames) reaching full contrast after 50 msec (3 frames), stayed at full contrast for 200 msec, then fell with the same Gaussian time course over the final 50 msec. We used four male observers (three of whom were naive with respect to the purpose of the experiment) aged between 16 and 30.

Experiment 1: detection threshold for moving plaid components

Before performing the main experiment of this study, we measured the detection threshold of our subjects for each of the components within the plaid. This was necessary to ensure that both gratings within the plaid were above threshold in expt 2 and to convert absolute contrast into threshold units which were needed for the simulations.

Methods. We determined the threshold for detecting the presence of a moving sinusoidal grating (signal) in the presence of a second moving grating (mask) of higher contrast and different orientation using an unconventional procedure: we held total contrast (mask plus signal) constant at (5, 10, 20 or 40%) to allow

easy comparison with the data from expt 2. The signal and mask were both 1.5 c/deg sinusoidal gratings oriented either +60 or −60 deg from the vertical axis and moving at 1 deg/sec. The choice of which of the two gratings was signal and which was mask was made randomly before each trial. Threshold was determined using a two-interval forced-choice protocol. The signal contrast level was chosen from a finite set which varied from 2.5 to 0.025% in fifteenth of a log unit steps. A trial consisted of two stimulus intervals (300 msec each separated by a 500 msec blank interval) presented in random order: one in which both signal and mask were present at a fixed total contrast and another in which only the mask was present. Thus, although the mask varied from trial to trial, it was identical in both intervals within a single trial. The signal contrast on a given trial was determined by one of two independent, randomly interleaved staircases. Within each staircase, contrast was reduced after two correct responses and increased after a single incorrect response.

Subjects were instructed to watch the screen and to fixate the small cross which appeared for 500 msec immediately before the onset of each stimulus, and was extinguished while the moving stimulus was displayed. They were then asked to indicate whether the signal was in the first or second interval. The resulting proportion correct (P) vs signal contrast (x) was fit with the best-fitting Weibull function (Watson, 1979; Weibull, 1951):

$$P = \min\left[0.99, 1 - 0.5_e^{-\left(\frac{x}{T}\right)^{3.5}}\right] \qquad (1)$$

where T is detection threshold. Thus, threshold was defined as the signal contrast at which the observer distinguished correctly 82% of the time between a weak "signal" grating moving within a moving plaid (i.e. in the presence of a "mask" grating) and the moving mask grating alone.

Results. Figure 3 plots $\log_{10}$ threshold contrast (T) as a function of $\log_{10}$ mask contrast (M). The data for all four subjects have a flat region below some critical mask contrast followed by a linear rising slope. This result is similar to the detection threshold results of Legge and Foley (1980) for pairs of stationary gratings of different spatial frequency. To quantify the results, we did a simple piecewise linear fit. For the flat portion, we made the assumption that the mean threshold at 5% total contrast

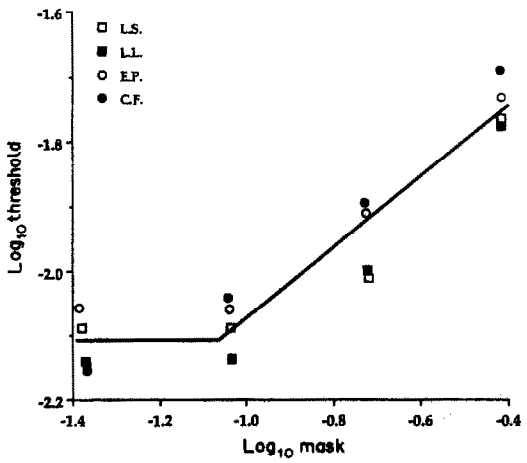

Fig. 3. Detection threshold vs mask contrast. This log–log plot contains the thresholds of all four subjects at four different total contrasts (5, 10, 20 and 40%). The solid line is the fitted curve used for the simulations and is given by equation (2).

(leftmost data points in Fig. 3), was the mean unmasked threshold (c). To measure the linear rising phase, we fit the three clusters of points generated for total contrasts of 10% and higher using linear regression to determine the slope (a) and intercept (b). Although it is arbitrary to include the points generated at 10% total contrast, any resulting error is probably small. The mean curve, shown in Fig. 3 as a thick line, is given by the following equation:

$$T = \max[10^{(a \log M - b)}, c] \qquad (2)$$

with $a = 0.548$, $b = 1.526$ and $c = 0.0078$. Equation (2) is merely a power law with an exponent of 0.548 for mask contrasts above 8.6%. The exponent found here is similar to that for stationary grating masks of different spatial frequency (range: 0.50–0.79 in Legge & Foley, 1980) and of different orientation (range: 0.40–0.72 in Phillips & Wilson, 1984). We will use equation (2) to estimate threshold for the simulation in Fig. 11.

For two subjects, we measured the effect of temporal and spatial frequency on detection threshold (Fig. 4). Temporal frequency (speed) had little effect on threshold (Fig. 4A). However, threshold was very sensitive to changes in spatial frequency (Fig. 4B): it increased at lower spatial frequency (0.75 c/deg) and decreased at higher spatial frequency (3.0 c/deg). This is consistent with previous studies of human spatiotemporal contrast sensitivity (Robson, 1966; Koenderink & van Doorn, 1979). The thresholds for these two observers at these temporal and spatial frequencies, given by equation (2)

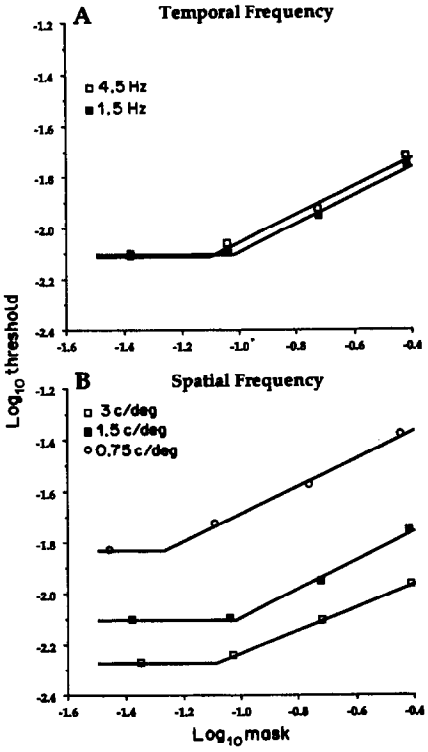

Fig. 4. Detection threshold vs mask contrast. These plots contain the thresholds of two subjects at four different total contrasts at two different temporal frequencies (A) and three different spatial frequencies (B). The solid curves are given by equation (2) using the parameters shown in Table 1.

with a, b and c as shown in Table 1, were used for the simulations in Figs 12 and 13.

Experiment 2: effect of contrast on the perceived direction of plaid motion

These experiments were conducted to measured systematically the effect of the relative contrast of the two gratings within the moving plaid on the perceived direction of motion. They were designed to test the model shown in Fig. 1 which predicts that changes in contrast will have no systematic effect on the perceived direction of plaid motion.

Methods. Subjects were presented with a single stimulus interval and were asked whether the stimulus moved to the left or right of subjective vertical. The true direction of plaid motion was varied by making the appropriate change in the ratio of the speeds of the two gratings (speed ratio) while the speed of the coherent plaid was held constant at 2 deg/s. The direction was changed within two interleaved up–down staircases to determine the direction for which subjects chose left or right with equal probability: we call this point *perceived vertical* and, for simplicity, we express it in degrees with respect to true vertical.

Table 1. Threshold parameters depend on temporal and spatial frequency

Spatial frequency	Temporal frequency	a	b	c
0.75	1.5	0.536	1.148	0.0148
1.5	1.5	0.554	1.532	0.0080
3.0	1.5	0.448	1.781	0.0054
1.5	4.5	0.553	1.497	0.0078

While total contrast was held constant at 5, 10, 20 or 40%, the ratio of the contrasts of the two gratings (contrast ratio) was varied in steps of $\sqrt{2}$. For example, the possible contrast pairs with 40% total contrast are: 20%, 20%; 23.4%, 16.6%; 26.7%, 13.3%; 29.6%, 10.4%, etc. and the symmetric counterparts. The series of stimuli included all possible contrast ratios which were powers of $\sqrt{2}$ and for which both gratings were above detection threshold. Since these series contained so many conditions, they were split into two interleaved subseries: one with contrast ratios which were even powers of $\sqrt{2}$ and another with odd powers of $\sqrt{2}$. The two subseries were presented in separate sessions.

For example, if the contrast of one of the gratings is so low that its perceived speed is half that of its true speed and if perceived rather than true speed feeds into the second stage of the model in Fig. 1, then the intersection of constraints rule will yield a severe directional bias towards the direction of motion of the grating of higher contrast (Fig. 5A). To quantify this bias, the true direction of plaid motion is altered by varying the speed ratio of the components. When the speed ratio reaches 2:1, the plaid will appear to move straight up (Fig. 5B). Perceived vertical measured this way (bias in Fig. 5B) is equal and opposite to the bias seen when the plaid is actually moving straight up (bias in Fig. 5A) assuming the contrast-dependent distortion is a multiplicative speed distortion which is independent of temporal frequency.

Our staircase method yielded typical psychometric curves (Fig. 6). We fit the data for each condition with a cumulative Gaussian using a weighted least-squares procedure (Mulligan & MacLeod, 1988) based on probit analysis (Finney, 1971). The standard deviation of the best-fitting cumulative Gaussian was defined as the *precision* in the observer's direction judgments. The location of the inflection point represents the bias that we refer to as

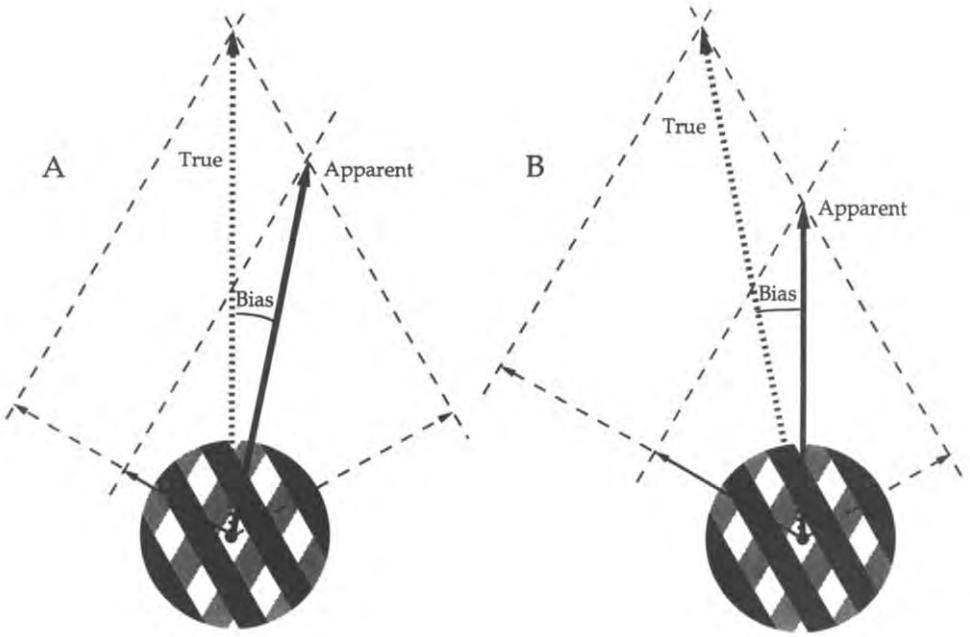

Fig. 5. Measuring the directional bias of a plaid moving straight up but composed of unequal contrast gratings. (A) If Thompson (1982) is correct, the lower-contrast grating (grey) will appear to move more slowly and the intersection of constraints rule applied to the perceived grating speeds will predict a bias toward the direction of motion of the higher-contrast grating (black). (B) If the speed ratio is changed until the plaid is perceived to move straight up then the true direction of plaid motion will have an equal and opposite bias to that in (A).

perceived vertical (the direction of motion that is perceived as pure vertical).

Results. For the standard plaid with a contrast ratio of 1, Table 2 shows the mean precision of four subjects averaged over three runs. Observers were apparently able to determine the direction of plaid motion to around ±4 deg. Although there seems to be some idiosyncratic variability, on average, there is no bias in the mean perceived vertical indicating that there was little or no systematic bias.

Figure 6 shows raw data for naive observer C.F. The psychometric curves shift along the x-axis for different contrast ratios: perceived vertical goes from 14.9 deg rightward for a contrast ratio of 0.125 to 1.1 deg rightward at equal contrast, and finally to 10.2 deg leftward for a contrast ratio of 8. However, the precision for the three conditions remained nearly unchanged at 3.4, 3.4 and 2.5 deg, respectively. This illustrates, at the raw-data level, that there are systematic changes in perceived vertical that occur as a function of contrast ratio and which cannot be explained by changes in precision.

A complete analysis of the performance of all four subjects shows that varying the contrast ratio away from 1 produced a large distortion in the perceived direction of motion. Figure 7 plots

Table 2. Precision for standard plaid with contrast ratio of 1

Subject	Perceived vertical (deg)	Precision (deg)
L.S.	−0.7 ± 1.3	4.0 ± 1.0
L.L.	0.5 ± 0.4	3.4 ± 0.4
E.P.	−3.4 ± 3.7	5.5 ± 1.6
C.F.	4.0 ± 2.6	2.7 ± 0.7
Mean	0.1	4.0

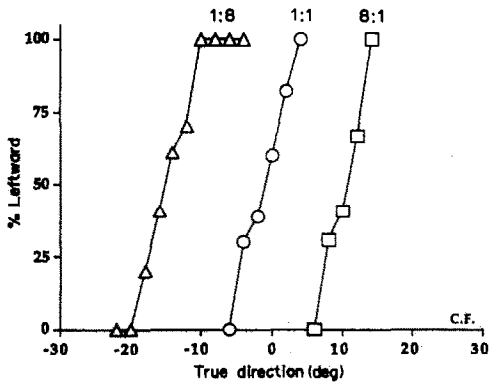

Fig. 6. Raw psychometric curves for plaid direction discrimination. This figure plots the percentage of stimulus presentations that were perceived as leftward vs the true direction of motion of the stimulus, for a single subject, for three runs at the three different contrast ratios indicated above the curves. Positive angles indicate leftward motion.

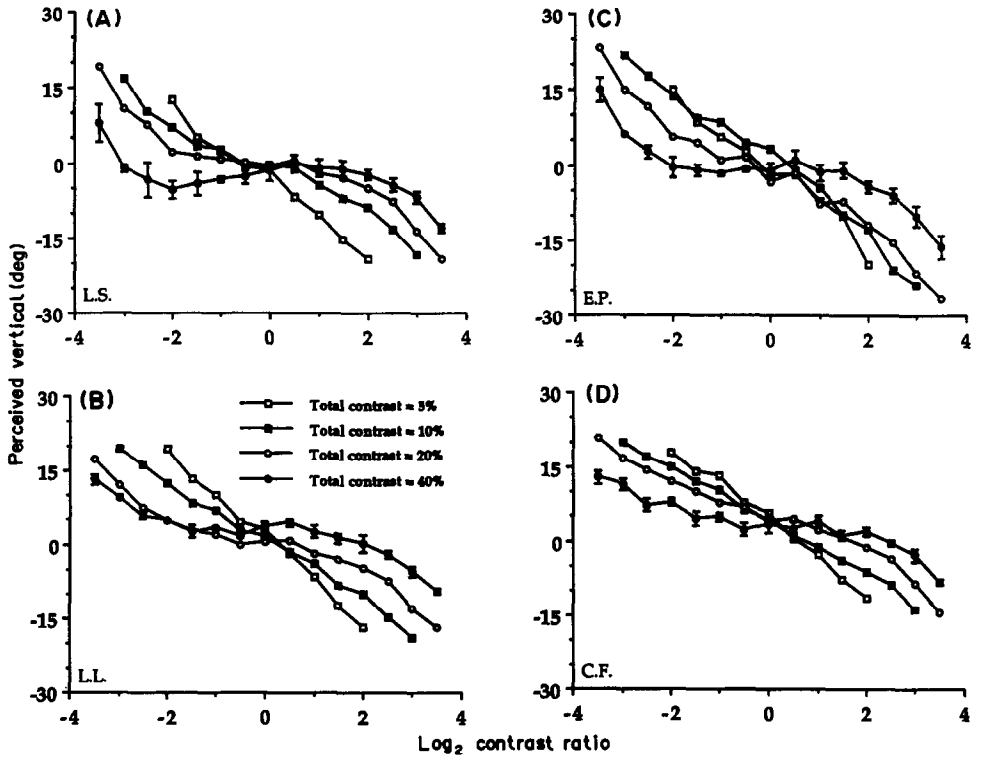

Fig. 7. Perceived vertical vs contrast ratio. (A–D) This figure plots perceived vertical for all four subjects at four different contrasts (5, 10, 20 and 40%). Error bars indicate standard deviations. Positive values indicate biases to the left of actual vertical.

perceived vertical in degrees away from true vertical as a function of $\log_2$ contrast ratio at total contrasts of 5, 10, 20 and 40% for all four subjects. (Typical standard deviations are plotted for 40% contrast.) When the gratings were of unequal contrast, the perceived direction of motion was shifted toward that of the higher-contrast grating. The effect increased systematically with increased contrast ratio although it was different for different total contrasts. All four subjects showed the same qualitative behavior.

The precision of the direction judgements was insensitive to changes in contrast except at extreme contrast ratios. Figure 8 plots precision as a function of $\log_2$ contrast ratio at four different total contrasts for all four subjects. Although subjects varied in their overall sensitivity to the direction of motion, there were no systematic effects of contrast ratio on precision.

A REVISED MODEL

The Adelson–Movshon model as shown in Fig. 1 fails to explain the results of expt 2 because it tacitly implies that the speeds of the components are accurately determined regardless of contrast. In this section, we amend the model to incorporate the finding of Thompson (1982) that the perceived speed of moving gratings is a function of contrast. The revised model is then tested with a variety of moving plaid stimuli.

Theory

Figure 9 shows the revised model. The modification is that the second stage is passed a contrast-distorted version of grating speed (for convenience hereafter called perceived speed) rather than actual grating speed. We construct perceived speed by multiplying actual speed by f, the contrast-dependent weighting function. For each grating, f is a function of the grating's contrast in threshold units (C_T), determined by dividing absolute contrast by threshold calculated using equation (2) with the other grating acting as the mask. Figure 10 shows the contrast-dependent weighting function that we used for the simulations that follow. We chose a Weibull function (Weibull, 1951) which goes to zero at threshold and which rises rapidly to just about 1 for contrasts

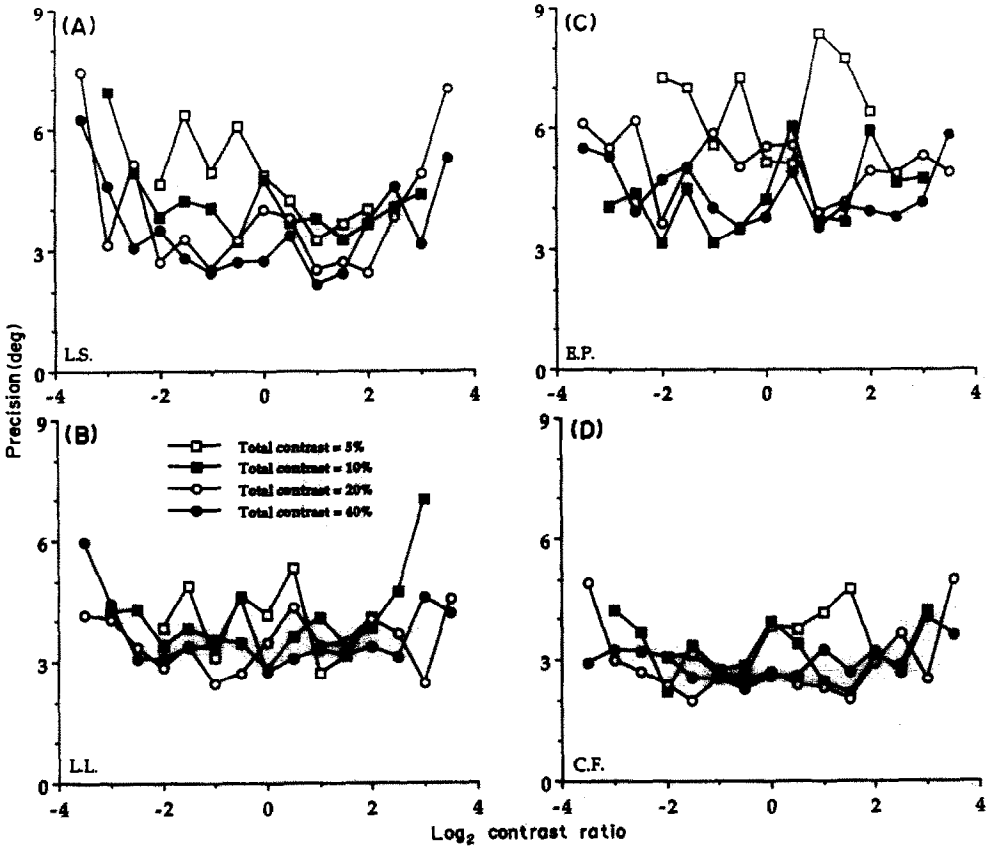

Fig. 8. Precision vs contrast ratio. (A–D) This figure plots precision for all four subjects at four different total contrasts (5, 10, 20 and 40%). Note that the scale is greatly amplified as compared to Fig. 7.

exceeding 10 threshold units. The explicit formula was:

$$f = 1 - e^{-\left(\frac{C_T - 1}{k_1}\right)^{k_2}} \quad \text{for } C_T \geq 1$$

$$= 0 \quad \text{for } C_T < 1 \qquad (3)$$

with $k_1 = 1.99$ and $k_2 = 0.76$ (by a least-squares fit of the data in Fig. 11).

Once f is determined for each grating human performance can be simulated. The simple intersection of constraints rule illustrated in Fig. 1B predicts that the perceived direction of motion (α) is given by the following equation (Stone, 1988):

$$\alpha = \arctan\left[\operatorname{cotan}\theta\left(\frac{V_1 - V_2}{V_1 + V_2}\right)\right] \qquad (4)$$

where V_1 and V_2 are the speeds of the two grating components and θ is the plaid angle. Note that, when $V_1 = V_2 = V$, α is zero and the plaid is perceived to move straight up. If, however, perceived rather than true speed is the input to the intersection of constraints stage and the perceived speed of the ith grating is $f_i V$, then the perceived direction of motion when the plaid is actually moving straight up is:

$$\alpha = \arctan\left[\operatorname{cotan}\theta\left(\frac{f_1 - f_2}{f_1 + f_2}\right)\right] \qquad (5)$$

Equation (5) allows us to simulate human perception of the direction of motion and to compare the result with the data in Fig. 7.

Results. Figure 11A plots the average bias of the four observers as a function of $\log_2$ contrast ratio for four different total contrasts by condensing the data presented in Fig. 7. Because the performance of all four subjects was qualitatively the same and because there is an inherent symmetry in the series of contrast ratios (i.e. the curves in Fig. 7 are nearly antisymmetric), we collapsed the data over symmetric pairs of data points and averaged over all subjects. We defined the contrast-dependent *bias* for a given contrast ratio (and its inverse) as the difference in perceived vertical for symmetric contrast-ratio pairs divided by two. Figure 11B plots the simulated output of the model (using equations 2, 3 and 5) under the same conditions. Using the simplex method of nonlinear curve fitting (Press, Flannery, Teukolsky & Vetterling, 1988),

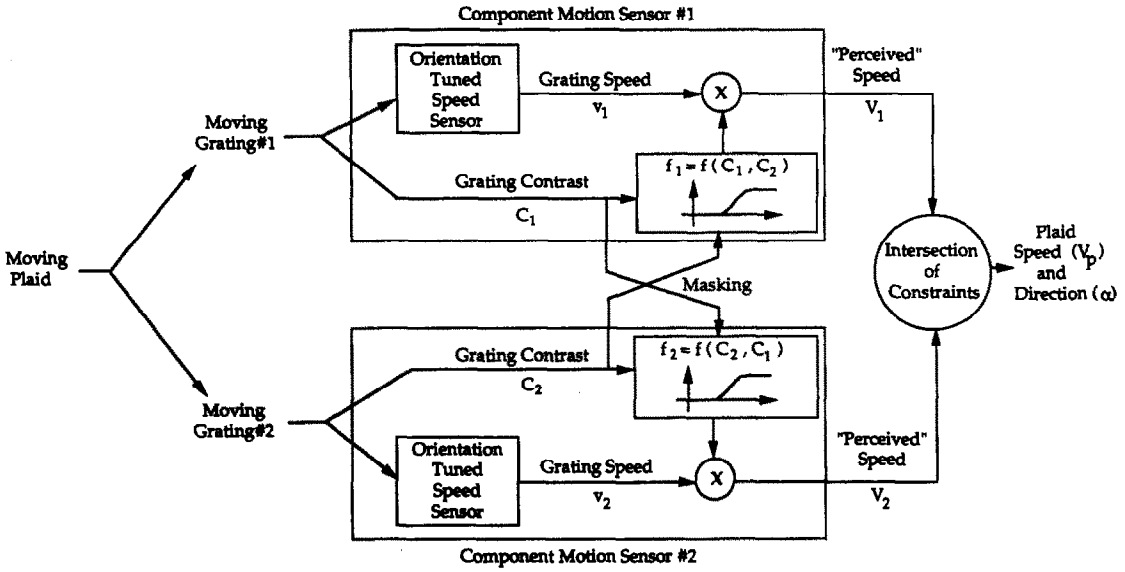

Fig. 9. A revised model. A contrast-dependent nonlinearity is added to each channel in the Adelson–Movshon model. The nonlinearity is a function of the contrast in threshold units of the input grating for the particular channel. Since threshold will be altered in the presence of the other grating (masking), the nonlinearity actually becomes a function of the contrast of both gratings.

we selected the only two floating parameters of the model (k_1 and k_2 of equation 3) such that simulations of equation (5) optimally (least-squared error) fit the data shown in Fig. 11A. Both the actual and simulated data show systematic shift in the perceived direction of motion up to about 20 deg toward the direction of the higher-contrast grating.

The model in Fig. 9 also qualitatively predicts the effect of changing the spatial and temporal frequency of the stimulus. Figure 12A plots the average bias of two observers as a function of $\log_2$ contrast ratio at three different spatial frequencies. Figure 12B plots the output of our model under the same conditions. The model qualitatively predicts that the bias will be larger at 0.75 c/deg and smaller at 3.0 c/deg. This prediction results from the changes in threshold as a function of spatial frequency (Fig. 4B): changes in the threshold parameters (Table 1) produce changes in C_T (through equation 2) which via equations (3) and (5) yield changes in the simulated bias. There is, however, significant quantitative discrepancy between the data and the simulations. Specifically, at 3.0 c/deg, there is small decrease in threshold so C_T is slightly larger and, therefore, f is slightly closer to 1. This leads to a small decrease in the simulated bias. However, there is a large decrease in the actual bias seen by our two observers. Similarly, at 0.75 c/deg, although there is a large increase in threshold and, therefore, a large increase in the simulated bias, there is only a small increase in the actual bias seen by our two observers.

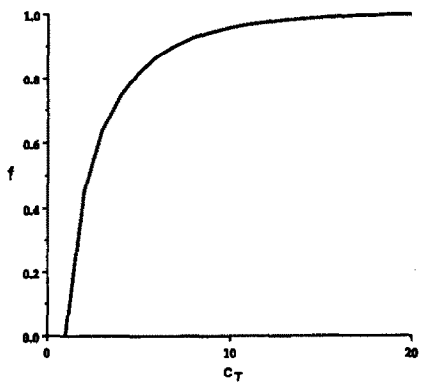

Fig. 10. The contrast-dependent weighting function. This rapidly saturating function given by equation (3) was derived by choosing k_1 and k_2 such that the mean squared error between the simulated and actual data was minimized.

Increasing the temporal frequency had little effect on the perceived direction of plaid motion. Figure 13A shows the average data for the same two observers at mean temporal frequencies of 1.5 and 4.5 Hz (plaid speeds of 2 and 6 deg/sec). Figure 13B shows the simulation of the model under the same conditions (using Table 1 and equations 2, 3 and 5). Temporal frequency changes had little effect on threshold (Fig. 4A) and therefore little effect on the simulated bias.

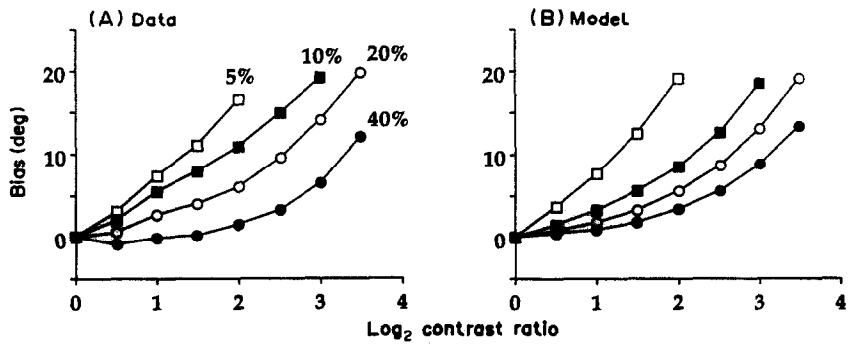

Fig. 11. Simulated vs actual bias. (A) This panel is a plot of the same data as in Fig. 7 averaged over subjects and over symmetric contrast-ratio pairs. (B) This panel shows simulations of the model in Fig. 9 under the same conditions as (A).

Our revised model does not appear robust to changes in the plaid angle. When the plaid angle is decreased to 30 deg, some subjects (3 of 7) show a bias toward the direction of motion of the lower-contrast grating. Figure 14 shows separately the bias of two different observers at three different plaid angles. Decreasing the plaid angle had different effects for the two subjects that were tested extensively. The subject in Fig. 14A showed a reduced bias for a plaid angle of 45 deg and, at 30 deg, a reversal of the bias toward the motion of the lower-contrast grating. The model in Fig. 9 does not predict this reversal. However, the subject in Fig. 14B did not show this reversal.

DISCUSSION

Perception of motion for unequal-contrast plaids

The results presented here and recent results by others (Kooi, DeValois & Wyman, 1988) clearly show that the simple intersection of constraints rule model proposed by Adelson and Movshon (1982) cannot account for the perceived direction of motion of plaids when the components are of unequal contrast. However, a simple modification of their model which takes into consideration the fact that the perceived speed of a moving grating is dependent on its contrast (Thompson, 1982) can, in most circumstances, account for the perceived direction of a moving plaid.

The modified model is robust in that it predicts qualitatively the changes (or lack thereof) in perceived direction as a function of temporal and spatial frequency. The quantitative discrepancy between the predicted and actual effect of changing spatial frequency (Fig. 13) can be explained if one postulates that f, the contrast-dependent weighting function, is itself a function of spatial frequency. In all of our simulations, f was defined by equation (3) as determined by fitting the data for a 1.5 c/deg stimulus. If one allows the two floating parameters to vary with spatial frequency, one can quantitatively account for the data in Fig. 13A. Unfortunately, we do not have a satisfactory explanation for the reversed biases seen by some observers at low plaid angles (Fig. 14).

Our results cannot be explained by incoherent plaid motion. Our standard stimulus subjectively appeared to move coherently for all

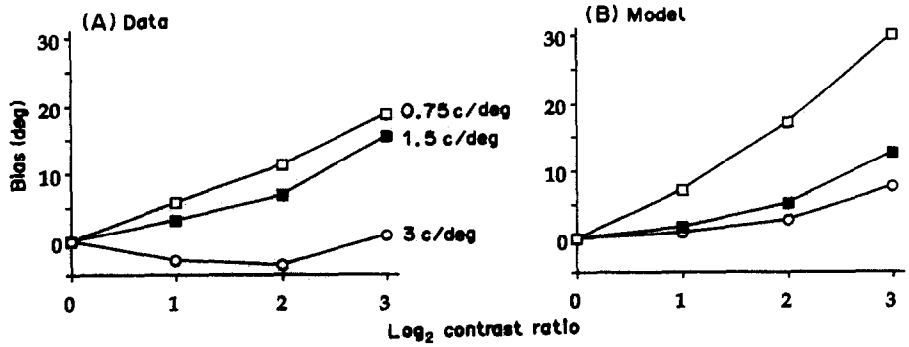

Fig. 12. Effect of spatial frequency. (A) This panel is a plot of the bias of two subjects at three spatial frequencies averaged over symmetric contrast-ratio pairs. (B) This panel shows simulations of the model in Fig. 9 under the same conditions as (A).

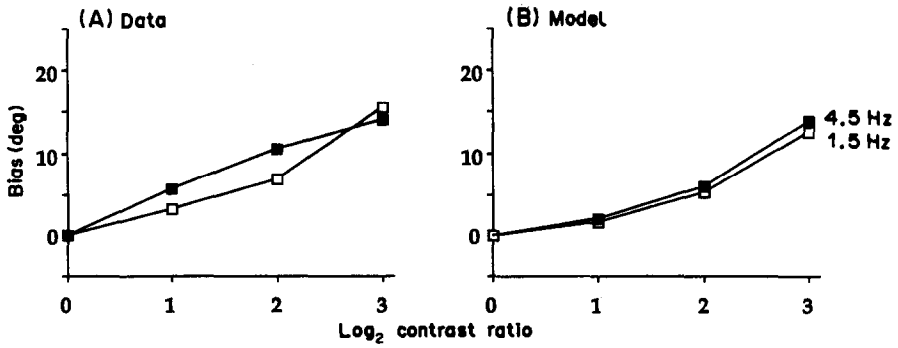

Fig. 13. Effect of temporal frequency. (A) This panel is a plot of the bias of two subjects at two temporal frequencies averaged over symmetric contrast-ratio pairs. (B) This panel shows simulations of the model in Fig. 9 under the same conditions as (A).

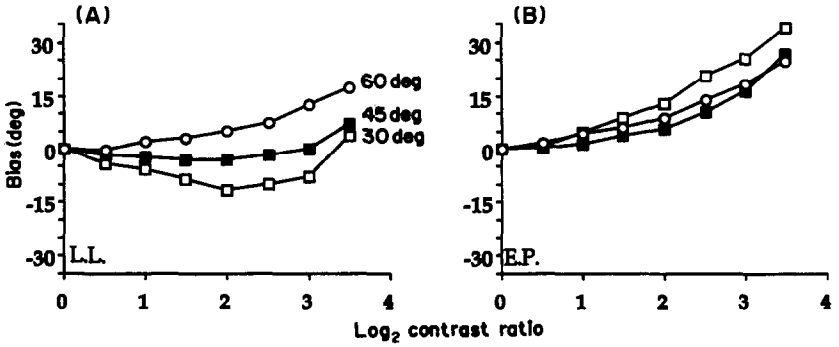

Fig. 14. Effect of plaid angle. (A–B) These two figures plot separately the bias of two subjects at three different plaid angles.

subjects even at contrast ratios as high as $8\sqrt{2}$. Attempts at measuring coherence thresholds were unsuccessful because subjects never reported sliding. At first glance this may appear inconsistent with previous observation on coherence (Movshon et al., 1986). However, it should be emphasized that Movshon and colleagues used a very different plaid stimulus for the results presented in their Fig. 5A. Our stimulus was composed of a plaid composed of two 1.5 c/deg gratings, moving at 1 deg/sec, viewed through a Gaussian window for 300 msec. The plaid in the Movshon study consisted of two gratings of unequal spatial frequencies (1.5 and 2.0c/deg), moving at 3 deg/sec, viewed through a sharp circular aperture for 1500 msec. In particular, Movshon and colleagues point out that under ideal conditions, plaids whose components are of different contrasts will cohere even when the low-contrast grating approaches detection threshold (Movshon et al., 1986) and our stimulus conditions are close to ideal in this respect (Movshon, personal communication). Finally, an objective measure of the coherence of our standard stimulus is the fact that the precision of the direction judgment is nearly independent of contrast ratio (Fig. 8). If some trials were perceived as coherent and moving in the veridical direction while other trials were perceived as incoherent and moving with a bias, we would not have observed the smooth shifting of the psychometric curves as a function of contrast ratio (Fig. 6).*

*If one assumes that for contrast ratio of 1, the percept is completely coherent and that for a contrast ratio of 8, the percept is completely incoherent then, from the slopes of the curves in Fig. 6, the variance in direction judgments for these conditions must be fortuitously identical (σ^2). This coincidence seems unlikely and argues against the possibility that our stimuli appeared incoherent. However, assuming equal variance for direction judgments under the pure coherent and pure incoherent conditions, for intermediate contrast ratios, if the coherent percept is seen with probability P, then it can be shown that the precision (standard deviation) in the mixed percept (Δ) will be related to the directional bias of the mixed percept (b) by the following equation:

$$\Delta = \sqrt{\sigma^2 + Pb^2/(1-P)} \quad (P < 1).$$

If, for example, one supposes that our plaid of 10% total contrast with a contrast ratio of 4 coheres only half the time ($P = 0.5$), the observed directional bias of around 11 deg predicts a precision of around 12 deg which is around three times that observed.

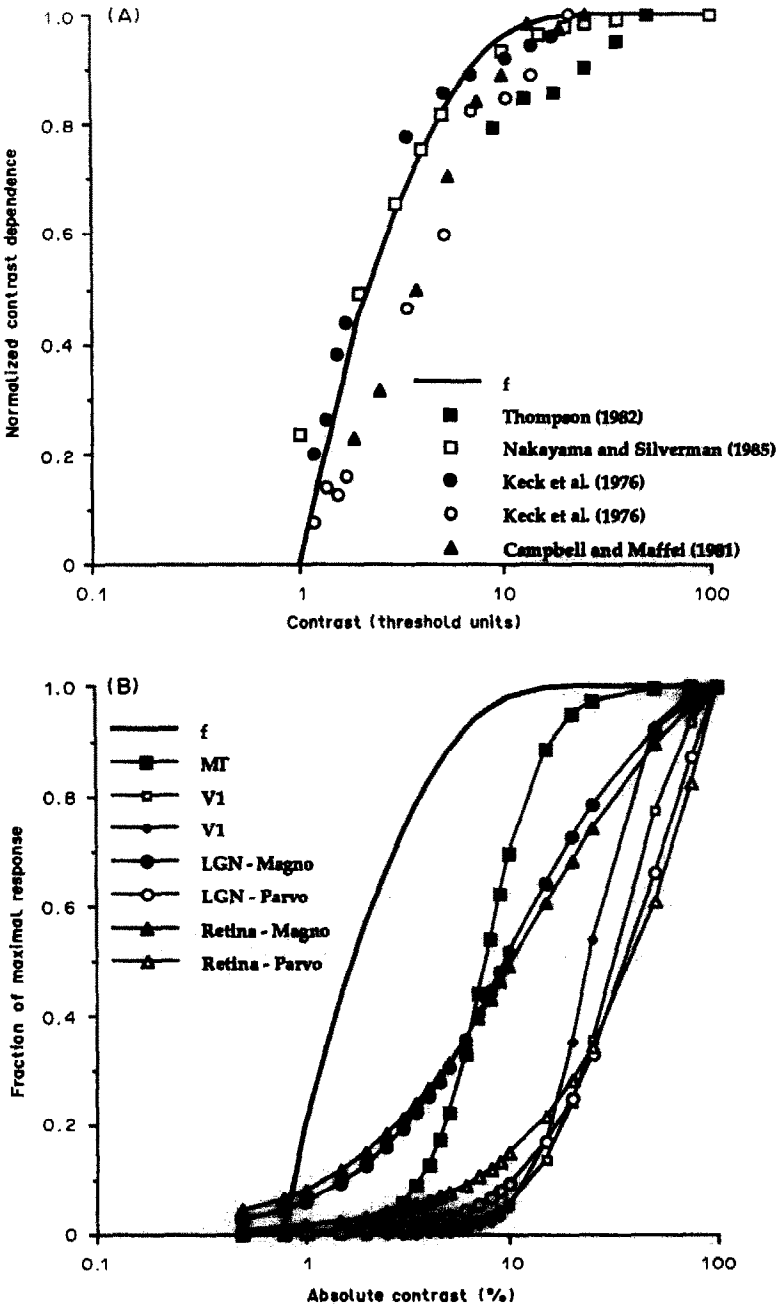

Fig. 15. Comparison of our contrast-dependent nonlinearity with normalized contrast response functions found in the literature. (A) This panel replots the same function f found in Fig. 10 on a log-scale plot together with data from four different psychophysical studies. The studies looked at the effect of contrast on perceived grating speed (solid squares), on detection of grating displacement (open squares), on motion after-effect duration (solid circles) and initial speed (open circles), and on perceived rotational frequency (solid triangles). (B) This panel replots f as a function of absolute contrast with the data from three different physiological studies: MT, V1 (squares), and LGN data are from Sclar et al. (1990), V1 data (diamonds) are from Abrecht and Hamilton (1982), and retinal ganglion-cell data are from Kaplan and Shapley (1986). All of the physiological contrast response functions were plotted using the mean or median best-fitting hypobolic function normalized to their response at 100% contrast. For the ganglion cells, the exponent of the hyperbolic function was fixed at 1.

An important caveat when interpreting our results and those of others (Adelson & Movshon, 1982; Movshon et al., 1986; Ferrera & Wilson, 1988, 1989; Kooi et al., 1988; Welch, 1989), is that moving plaids are strong stimuli for optokinetic eye movement and that eye movements may contribute to the perception of plaid motion. In this study, the brief

stimulus duration (300 msec) makes it unlikely that eye-movement contamination dominates the percept although it does not eliminate that possibility. Initial experiments showed that leaving the fixation light on during the stimulus presentation did not effect direction discrimination. In the main experiments, we chose to extinguish the fixation light because it would not have fully suppressed eye movements but it would have provided a cue for relative motion.

Finally, our success in salvaging the Adelson–Movshon hypothesis should not be construed as proof that the hypothesis is correct. Recently, Welch (1989) provided strong support for the first tenet of the hypothesis: that the motion of the plaid is first decomposed into the motion of the individual components. However, Ferrera and Wilson (1988, 1989) have found evidence that the intersection of constraints rule is not always used at the second stage of processing. Explaining our data with the revised Adelson–Movshon model should not be viewed as an endorsement of the intersection of perpendicular contraints rule. It is likely that our data could be explained using a different second-stage rule. However, a contrast-dependent nonlinearity would still be required.

Contrast-dependent effects in motion processing

The contrast-dependent weighting function (Fig. 10), determined by a least-squares two-parameter fit to our bias data in Fig. 11, saturates at very low contrast (reaches 0.5 at 2.3 times threshold or below 2% contrast). Low-contrast saturation is associated with many psychophysical phenomena involing moving stimuli. Figure 14A replots f vs $\log_{10}$ contrast (in threshold units) together with psychophysical measurements made in four different studies. As stated above, Thompson (1982) measured directly the perceived speed of gratings as a function of contrast and found that low-contrast gratings appear to move more slowly than a high-contrast reference moving at the same speed. The solid squares plot the mean perception of two subjects as the ratio of perceived speed of a test grating to that of a standard grating of 25% contrast, assuming a detection threshold contrast of 0.5%. A higher detection threshold would shift the curve to the left and would, therefore, improve overlap with f. In Thompson's study, the contrast effect appears to saturate slower but he used a stimulus that differed slightly in spatial (his 2 c/deg vs our 1.5 c/deg) and temporal (his 2 Hz vs our 1.5 Hz) frequency, in mean luminance (his 32 vs our 100 cd/m^2), and greatly in duration (his 2.5 sec vs our 0.3 sec). In addition, from his data, it is difficult to assess the precision and possible bias associated with his matching technique. We assumed that a test grating of 25% matched the standard perfectly. All of these factors may explain the small quantitative differences between the f measured in this study and Thompson's results.

Nakayama and Silverman (1985) measured the effect of contrast on the minimum displacement of a sinusoidal grating that can be discriminated (left or rightward motion of a vertical grating). The minimum displacement (in degrees of phase) necessary for discrimination decreases to 5–10 deg as contrast increases to about 3% then remains nearly constant. The open squares in Fig. 14A plot what the authors called the normalized "effective contrast" of the stimulus (mean of the best-fitting hyperbolic functions for two subjects assuming again that threshold was 0.5%). Their effect saturates at nearly the same rate (reaches 0.5 at about 2.1 times threshold or around 1% contrast) as the contrast effect in this study.

Keck, Palella and Pantle (1976) studied the effect of contrast on motion after-effect (MAE) and found that both the duration (solid circles) and perceived initial speed (open circles) of MAE (normalized with respect to that of a 12.5% grating) saturate at low contrast. Their MAE duration data nearly superimpose on the f derived in this study. Their MAE speed data saturate more slowly and at nearly the same rate as Thompson's data.

Campbell and Maffei (1981) looked at the effect of contrast on perceived rotational speed and found that a rotating low-contrast grating patch was perceived to rotate more slowly than an otherwise identical high-contrast stimulus rotating at the same rate. The solid triangles in Fig. 14A plot the ratio of perceived rotational frequency of a test grating to that of a 20% contrast grating. The effect again saturates at low contrast although not as rapidly as the effect in this study.

The fact that all of these disparate psychophysical studies seem to saturate similarly at low contrast probably reflects a fundamental property of a shared input stage for human judgments of motion. It is important for contrast responses within the motion processing system to saturate early so as to disambiguate

signals related to contrast from those related to motion. The interesting finding of this and the other studies is not that there are contrast-dependent misperceptions of motion but actually that these misperceptions only occur at the extreme low end of the contrast scale.

Examination of the contrast response properties of neurons within the monkey visual cortex suggest that this shared input is at a higher stage than striate cortex (Albrecht & Hamilton, 1982; Sclar, Maunsell & Lennie, 1990). Figure 14B replots f as a function of log absolute contrast together with the normalized mean contrast response function of ganglion cells (triangles), lateral geniculate neurons (circles), and neurons within the striate (open squares and diamonds) and extrastriate visual cortex (solid squares) of macaque monkeys. Albrecht and Hamilton found that striate neurons (V1) show a wide range of contrast response functions. Some neurons begin responding at 1% contrast and saturated by 10%. Others do not begin responding until 10% contrast. V1 neurons, on average, reached 50% of their maximal response at 23.9% contrast with those tuned for 1.5 c/deg doing so at around 20%. Similarly, Sclar and colleagues (1990) found that, on average, V1 neurons reached 50% of maximal response at 31.6%. They, also, found that neurons within the middle temporal area (MT), an area of extrastriate visual cortex specifically associated with motion processing (Maunsell & Van Essen, 1983; Rodman & Albright, 1987; DeYoe & Van Essen, 1988), saturate at much lower contrast. On average, MT neurons reached 50% of their maximal response at only 7.6% contrast with some individual neurons reaching 50% saturation at as low as 1.6% contrast. Sclar and colleagues (1990) and Kaplan and Shapley (1986) measured the contrast response functions of neurons at earlier stages in the visual pathway, in the LGN and retina, respectively. MT neurons have higher contrast sensitivities apparently because they receive a selective input from the magnocellular pathway (solid symbols) beginning at the retina (Kaplan & Shapley, 1986; DeYoe & Van Essen, 1988) and because they receive pooled information from lower-level neurons with smaller receptive fields (Sclar et al., 1990). On average, LGN neurons and ganglion cells within the magnocellular pathway reached 50% of their maximal response at 9.6 and 10.4% contrast, respectively. A separate parvocellular pathway has lower contrast sensitivity with neurons reaching 50% of their maximal response at 36.5% contrast in the LGN and at 38.9% in the retina.*

The psychophysical phenomena described in Fig. 14A still saturate faster than the responses in MT (or responses at earlier stages in the magnocellular pathway) suggesting that either the common input for psychophysical judgments is from a higher cortical level or that the psychophysical judgments use information pooled from several MT neurons. It is also possible that the humans make psychophysical judgments based on input from a selective group of MT neurons since a subset appear to saturate as fast as the psychophysics (Sclar et al., 1990).

In neither V1 nor MT does it appear that speed is encoded in the firing rate of individual neurons (Maunsell & Van Essen, 1983; Rodman & Albright, 1987; Kennedy & Bullier, 1986; Movshon, 1975). Therefore, it is a reasonable conjecture that the speed of a moving grating is encoded as some integral of the collective output of an ensemble of MT neurons or in some "higher" cortical area that receives pooled input from MT. However, regardless of the specific scheme used to encode speed, at low contrasts, because neuronal activity within MT is affected by both speed and contrast, a reduction in contrast will cause a decrement in the collective output of the ensemble by reducing the number of cells that respond and the amplitude of the responses of those that do. This might be misinterpreted as a reduction in stimulus speed which just as easily could have been the cause. This would lead to the psychophysical findings present here and in Thompson's study (1982). At higher contrast, changes in the activity of MT neurons only reflect changes in the motion of the stimulus so the perception of speed is veridical.

*When the average response (R) is given by a hyperbolic function of contrast (C), i.e.

$$R = \frac{R_{max} C^n}{C^n + C_{50}^n};$$

it should be emphasized that C_{50} is not the contrast value which produces half the maximum response. Because C varies only between 0 and 100%, R_{max} is often the extrapolated maximum and the true contrast value at half-maximum response is:

$$\frac{100\, C_{50}}{\sqrt[n]{100^n + 2C_{50}^n}};$$

which reduces to C_{50} for $C_{50} \ll 100\%$.

Implications for models of human motion processing

The fact that contrast can systematically and dramatically distort the perceived direction of plaid motion puts a constraint on future models of human motion processing. As stated above it cannot be explained by the schematic model in Fig. 1 (Adelson & Movshon, 1982) but can by the simple modification presented in Fig. 9. However, both the Adelson–Movshon and the revised model are mere cartoons that provide an organizational structure for motion processing but are not true models. We now examine the performance of a few, more complete models of visual motion processing, (some of which conform to the structure proposed in Fig. 9 and some of which do not) to determine if they mimic our psychophysical findings.

One class of models that would not exhibit the same behavior as our subjects consists of cross-correlation models (e.g. Leese, Novak & Taylor, 1970). When a plaid with components of unequal contrast moves straight up, a pure cross-correlation technique will show no bias because cross-correlation determines the maximum overlap between two successive frames and overlap is perfect (neglecting noise) for true upward motion. Therefore, our results clearly indicate that the human visual system is not using a full-field cross-correlation technique. Bulthoff, Little and Poggio (1989) have recently proposed a neural-network implementation of a variant of the cross-correlation method. Rather than performing a simple 2-D cross-correlation over the whole image, it does a local cross-correlation over a patch. If this patch is the whole image, the model reduces to a simple cross-correlation model. In response to a plaid whose components are of unequal contrast, it would therefore show no bias. If the patch is small as compared to the spatial frequency of the gratings, the model will exhibit a spatially nonuniform response: at different points within the image it will detect either the motion of an individual component or of a node. It is hard to say what the exact response to a plaid whose components were of unequal contrast would be, as it depends critically on the patch size, but it seems likely that it would detect either the true direction of motion or nonrigid motion but not the systematic biases that we observed empirically.

A second class of models that one would expect to be invariant to changes in contrast comprises those that track the motion of the nodes by tracking the motion of edges (zero-crossings of the second spatial derivative) within the image (e.g. Marr & Ullman, 1981; Hildreth, 1984). The nodes in our stimuli always moved exactly upward regardless of the contrast ratio and therefore should only provide information about the true direction of motion. However, it is not the determination of the edge velocity *per se* but but how the edge velocities are combined to determine the global motion of the plaid that is important. For example, one model (Perrone, 1990) that identifies moving edges within images and hence, ostenstibly, tracks the moving nodes, uses a cosine-weighted voting scheme to determine the global motion. The Perrone model shows a directional bias toward the direction of motion of the higher-contrast grating in response to plaids composed of unequal contrast gratings (Perrone & Stone, 1988) although the biases are twice as large as those found here. The simulated bias occurs because the nodes change shape and become spatially asymmetric for contrast ratios different than 1. This asymmetry causes a shift in the distribution of edge-velocity vectors. The voting scheme then causes a shift in the output of the model. Therefore, models that look at the motion of edges within the image can predict directional biases despite the fact that the features that they are tracking are moving in the true direction of motion of the pattern as a whole. However, feature-based model that identify and track the nodes *per se* would no show such biases.

A third class of motion models consists of those that work directly with the spatial and temporal gradients of the image intensity (e.g. Limb & Murphy, 1975; Horn & Schunk, 1981). Recently, a neural-network implementation of this approach has been shown to respond to plaids composed of unequal contrast gratings with directional biases toward the direction of motion of the higher-contrast grating (Koch, Wang, Mathur, Hsu & Suarez, 1989). The bias is caused by the contrast dependence of their first-stage neurons (U cells) at low contrast although the asymmetry of the spatial gradient, when the components have unequal contrasts, may also contribute. However, the proposal that the U cells are located in V1 is not plausible because the output of U cells is proportional to speed and no such units have been found in either V1 or anywhere else in the primate visual cortex.

A fourth class of models includes those that look at motion in the frequency domain (e.g. Watson & Ahumada, 1983, 1985), those that calculate motion energy (e.g. Adelson & Bergen, 1985), and the related elaborated Reichardt detectors (van Santen & Sperling, 1985). Motion-energy models are expected to be seriously effected by contrast manipulations since motion energy is basically proportional to the square of the contrast. To address this weakness, Heeger (1987) proposed a modified motion-energy model which normalizes the response by dividing the output of each sensor by the total energy for a given orientation. Because of this normalization, Heeger's model determines the true direction of motion for moving plaids with contrast ratios as high as 1:32, which is inconsistent with our results. Furthermore, the model without contrast normalization will yield larger biases over a wider range of contrasts than the biases observed here. New approaches to reduce the inherent contrast dependence of motion-energy models need to be developed because a version with partial normalization might reproduce our psychophysical findings.

Watson and Ahumada (1985) designed their model of human motion processing to be robust to contrast variations. They determine the direction of motion by examining the temporal frequency of the response of linear spatio-temporal filters, a measure which is independent of contrast. Because of this, the Watson–Ahumada model finds the true direction of motion for plaids with contrast ratios as high as 1:10 which is inconsistent with our results. Their model, however, can also be modified, by incorporating a contrast-dependent nonlinearity.

Our results are therefore inconsistent with the specific versions of a number of current models of human motion processing. In many cases, simple modifications can be made to account for our data. This discussion is not meant to be an exhaustive survey of motion models but is merely intended to show how our psychophysical results can be used in many cases to refine and in some cases to rule out certain models. It is also meant to show how the quantitative comparison of empirical and simulation data is needed for the meaningful analysis of the biological plausibility of existing models of human visual motion processing.

Acknowledgements—This work was supported by NASA RTOP No. 506-47-11 and a National Research Council associateship to L.S. The authors would like to thank John Perrone, J. Matt Valeton and Al Ahumuda for their comments on an earlier draft, Bill Paulsen for help with Figs 1 and 9, and John Maunsell for graciously providing electrophysiological data.

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.

Adelson, E. H. & Movshon, J. A. (1982). Phenomenal coherence of moving visual patterns. *Nature, London, 300*, 523–525.

Albrecht, D. G. & Hamilton, D. B. (1982) Striate cortex of monkey and cat: Contrast response function. *Journal of Neurophysiology, 48*, 217–237.

Bulthoff, H., Little, J. & Poggio, T. (1989). A parallel algorithm for real-time computation of optical flow. *Nature, London, 337*, 549–554.

Campbell, F. W. & Maffei, L. (1981). The influence of spatial frequency and contrast on the perception of moving patterns. *Vision Research, 21*, 713–721.

DeValois, R. L. & DeValois, K. K. (1980). Spatial vision. *Annual Review of Psychology, 31*, 309–341.

DeYoe, E. A. & Van Essen, D. C. (1988). Concurrent processing streams in monkey visual cortex. *Trends in Neuroscience, 11*, 219–226.

Fennema, C. L. & Thompson, W. B. (1979). Velocity determination in scenes containing several moving images. *Computer Graphics and Image Processing, 9*, 301–315.

Ferrera, V. P. & Wilson, H. R. (1988). Perceived direction of moving 2-D patterns. *Investigative Opthalmology and Visual Science* (Suppl.), *29*, 264.

Ferrera, V. P. & Wilson, H. R. (1989). Perceived speed of moving 2-D patterns. *Investigative Ophthalmology and Visual Science* (Suppl), *30*, 75.

Finney, D. J. (1971). *Probit analysis*. Cambridge: Cambridge University Press.

Floyd, R. W. & Steinberg, L. (1975). An adaptive algorithm for spatial gray scale. *SID 1975 International Symposium Digest of Technical Papers*, pp. 36–37.

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

Hildreth, E. C. (1984). *The measurement of visual motion*. Cambridge, Mass.: MIT Press.

Horn, B. K. P. & Schunk, B. G. (1981). Determining optical flow. *Artificial Intelligence, 17*, 185–203.

Kaplan, E. & Shapley, R. M. (1986). The primate retina contains two types of ganglion cells, with high and low contrast sensitivity. *Proceedings of the National Academy of Science, 83*, 2755–2757.

Keck, M. J., Palella, T. D. & Pantle, A. (1976). Motion aftereffect as a function of the contrast of sinusoidal gratings. *Vision Research, 16*, 187–191.

Koch, C., Wang, H. T., Mathur, B., Hsu, A. & Suarez, H. (1989). Computing optical flow in resistive networks and in the primate visual system. *Proceedings of the Workshop on Visual Motion* (pp. 62–72). Washington D.C.: IEEE Computer Society Press.

Koenderink, J. J. & van Doorn, A. J. (1979). Spatiotemporal contrast detection threshold surface is bimodal. *Optical Letters, 4*, 32–34.

Kooi, F. L., DeValois, R. L. & Wyman, T. K. (1988). Perceived direction of moving plaids. *Investigative Opthalmology and Visual Science* (Suppl.), *29*, 265.

Landy, M. S., Cohen, Y. & Sperling, G. (1984). HIPS: A unix-based image processing system. *Computer Vision, Graphics and Image Processing, 25*, 331–347.

Leese, J. A., Novak, C. S. & Taylor, V. R. (1970). The determination of cloud pattern motions from geosynchronous satellite image data. *Pattern Recognition, 2*, 279–292.

Legge, G. E. & Foley, J. M. (1980). Contrast masking in human vision. *Journal of the Optical Society of America, 70*, 1458–1471.

Limb, J. O. & Murphy, J. A. (1975). Estimating the velocity of moving images in television signals. *Computer Graphics and Image Processing, 4*, 311–327.

Marr, D. & Ullman, S. (1981). Directional selectivity and its use in early visual processing. *Proceedings of the Royal Society, London, B 211*, 151–180.

Maunsell, J. H. R. & Van Essen, D. C. (1983). Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. *Journal of Neurophysiology, 49*, 1127–1147.

Movshon, J. A. (1975). The velocity tuning of single units in cat striate cortex. *Journal of Physiology, London, 249*, 445–468.

Movshon, J. A., Adelson, E. H., Gizzi, M. S. & Newsome, W. T. (1986). The analysis of moving visual patterns. *Experimental Brain Research* (Suppl.), *11*, 117–151.

Mulligan, J. B. (1986). Minimizing quantization errors in digitally controlled CRT displays. *Color Research and Applications, 11*, S47–S51.

Mulligan, J. B. & MacLeod, D. I. A. (1988). Reciprocity between luminance and dot density in the perception of brightness. *Vision Research, 28*, 503–519.

Mulligan, J. B. & Stone, L. S. (1989). Halftoning method for the generation of motion stimuli. *Journal of the Optical Society of America, A 6*, 1217–1227. Also published as Mulligan, J. B. and Stone, L. S. (1988). Efficient use of bit-planes in the generation of motion stimuli, NASA TM–101022.

Nakayama, K. & Silverman, G. H. (1985). Detection and discrimination of sinusoidal grating displacements. *Journal of the Optical Society of America, A 2*, 267–274.

Orban, G. A., Kennedy, H. & Bullier, J. (1986). Velocity sensitivity and direction selectivity of neurons in areas V1 and V2 of the monkey: Influence of eccentricity. *Journal of Neurophysiology, 56*, 462–480.

Perrone, J. (1990). A simple technique for optical flow estimation. *Journal of the Optical Society of America, A 7*, 264–278.

Perrone, J. & Stone, L. S. (1988) Two-dimensional motion analysis: Information in the nodes of moving plaids. *Optical Society of America Technical Digest, 11*, 141.

Phillips, G. C. & Wilson, H. R. (1984) Orientation bandwidths of spatial mechanisms measured by masking. *Journal of the Optical Society of America, A 1*, 226–232.

Press, W. H., Flannery, B. P., Teukolsky, S. L. & Vetterling, W. T. (1988). *Numerical recipes in C: The art of scientific computing*. New York: Cambridge University Press.

Robson, J. G. (1966). Spatial and temporal contrast-sensitivity functions of the visual system. *Journal of the Optical Socity of America, 56*, 1141–1142.

Rodman, H. R. & Albright, T. D. (1987). Coding of visual stimulus velocity in area MT of the macaque. *Vision Research, 27*, 2035–2048.

van Santen, J. P. H. & Sperling, G. (1985). Elaborated Reichardt detectors. *Journal of the Optical Society of America, A 2*, 300–321.

Sclar, G., Maunsell, J. & Lennie, P. (1990). Coding of the image contrast in central visual pathways of the macaque monkey. *Vision Research, 30*, 1–10.

Stone, L. S. (1988). Precision in the perceived direction of a moving pattern. NASA TM-101080.

Stone, L. S., Mulligan, J. B. & Watson, A. B. (1988a). Contrast affects the perceived direction of a moving plaid. *Optical Society of America Technical Digest, 11*, 141.

Stone, L. S., Mulligan, J. B. & Watson, A. B. (1988b). Neural determination of the direction of motion: Contrast affects the perceived direction of a moving plaid. *Society for Neuroscience Abstracts, 14*, 1251.

Thompson, P. (1982). Perceived rate of movement depends on contrast. *Vision Research, 22*, 377–380.

Watson, A. B. (1979). Probability summation over time. *Vision Research, 19*, 515–522.

Watson, A. B. & Ahumada, A. J. (1983). A look at motion in the frequency domain. NASA TM-84352. Also published in *Motion: Representation and Perception* (1983). Baltimore, MD: Association of Computing Machinery.

Watson, A. B. & Ahumada, A. J. (1985). Model of human visual motion sensing. *Journal of the Optical Society of America, A 2*, 322–342.

Watson, A. B., Nielsen, K. R. K., Poirson, A., Fitzhugh, A., Bilson, A. J., Nguyen, K. & Ahumada, A. J. (1986). Use of a raster framebuffer in vision research. *Behavioural Research Methods, Instruments and Computers, 18*, 587–594.

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

Welch, L. (1989). Perception of moving plaids reveals two motion processing stages. *Nature, London, 337*, 734–736.