Vision Research 48 (2008) 2321–2328

Contents lists available at ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

Discriminating depth in corrugated stereo surfaces: Facilitation by a pedestal is explained by removal of uncertainty

Mark A. Georgeson [a,*], Timothy A. Yates [a], Andrew J. Schofield [b]

[a] School of Life and Health Sciences, Aston University, Birmingham B4 7ET, UK
[b] School of Psychology, University of Birmingham, Birmingham B15 2TT, UK

ARTICLE INFO

Article history:
Received 28 February 2008
Received in revised form 10 July 2008

Keywords:
Stereopsis
Disparity modulation
Detection
Facilitation
Uncertainty
Signal detection theory

ABSTRACT

With luminance gratings, psychophysical thresholds for detecting a small increase in the contrast of a weak 'pedestal' grating are 2–3 times lower than for detection of a grating when the pedestal is absent. This is the 'dipper effect' – a reliable improvement whose interpretation remains controversial. Analogies between luminance and depth (disparity) processing have attracted interest in the existence of a 'disparity dipper'. Are thresholds for disparity modulation (corrugated surfaces), facilitated by the presence of a weak disparity-modulated pedestal? We used a 14-bit greyscale to render small disparities accurately, and measured 2AFC discrimination thresholds for disparity modulation (0.3 or 0.6 c/deg) of a random texture at various pedestal levels. In the first experiment, a clear dipper was found. Thresholds were about 2× lower with weak pedestals than without. But here the phase of modulation (0 or 180 deg) was varied from trial to trial. In a noisy signal-detection framework, this creates uncertainty that is reduced by the pedestal, which thus improves performance. When the uncertainty was eliminated by keeping phase constant within sessions, the dipper effect was weak or absent. Monte Carlo simulations showed that the influence of uncertainty could account well for the results of both experiments. A corollary is that the visual depth response to small disparities is probably linear, with no threshold-like nonlinearity.

1. Introduction

For patterns defined by luminance contrast, it is well known that observers can detect a test grating better when it is superimposed on a similar, barely visible, pedestal grating than when it is presented on its own without the pedestal (Legge & Foley, 1980; Nachmias & Sansbury, 1974). This improvement for low-contrast pedestals is often called 'facilitation' or the 'dipper effect', but its source remains controversial. It may be (a) based on nonlinear transduction of contrast in the visual system (Legge & Foley, 1980), or (b) based on linear transduction but with uncertainty about which visual channels to monitor (Pelli, 1985). On both models, the net effect is a nonlinear relation between contrast and detectability (d') that leads to facilitation in the pedestal experiment.

If the pedestal is shifted by 90° in orientation, or phase, from the test grating, or is presented shortly before or after the test pattern, masking is still observed but the low-contrast facilitation disappears (Foley, 1994; Foley & Chen, 1999; Georgeson & Georgeson, 1987). This suggests that, whatever the cause, facilitation requires summation of responses to the pedestal and test patterns within the same visual channel. Schofield and Georgeson (1999) used this logic to argue that luminance modulation (LM) and contrast modulation (CM) are detected in separate channels, because LM facilitated LM, and CM facilitated CM, but CM did not facilitate LM, nor vice-versa. During that work, we also noticed that images formed as combinations of LM and CM tended to look like corrugated surfaces in depth (Schofield, Hesse, Rock, & Georgeson, 2006). This prompted us to consider the combination of depth cues in signal-processing terms, and to ask whether stereo disparity modulation (DM), which also gives rise to the perception of corrugated surfaces, would combine with and mutually facilitate LM or CM.

As a step toward that goal, we here investigate whether DM facilitates itself: is there a 'disparity dipper'? More precisely, do textures with small amounts of disparity modulation facilitate the detection of increments in that disparity modulation? This question was tackled by Lunn & Morgan (1997) using disparity-modulated random textures (here called DM gratings), and for their two observers at two spatial frequencies (0.25, 0.5 c/deg) they found no disparity dipper. Disparity modulation thresholds were unaffected by the presence of a weak DM pedestal, although pedestals with greater modulation did produce a masking effect (threshold elevation). There are two reasons for re-visiting the question: (i) Lunn & Morgan's pedestal levels were all above threshold, and so might have been too high to reveal facilitation

* Corresponding author.
E-mail address: m.a.georgeson@aston.ac.uk (M.A. Georgeson).

0042-6989/$ - see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2008.07.009

and (ii) the limitations of an 8-bit graphics system (256 grey levels) might have prevented them from rendering very small disparities adequately. Indeed, Lunn & Morgan (1997, p. 364) noted that "A more detailed cyclopean experiment aimed specifically at locating a dip would be required to resolve the issue". That is our aim here.

Disparity thresholds are strikingly low – often requiring just a few seconds of arc spatial displacement between the eyes. On digital displays this requires a spatial precision that is a small fraction of one pixel. Such sub-pixel resolution can be gained (in effect) by exploiting greyscale resolution in place of pixel resolution, as described later. We used a 14-bit display system (16,384 grey levels) that does enable very fine sub-pixel resolution, and we report here two experiments that used a two-alternative forced-choice (2AFC) staircase procedure to measure DM increment thresholds, for a wide range of DM pedestal levels (Fig. 1).

2. Methods

Images of a random, Gaussian white noise texture (512×512 pixels) were created in *Matlab* 5.2 on a Macintosh G4 computer, and displayed on a Clinton fast-phosphor monitor via a CRS Bits++ interface in true 14-bit greyscale mode. *PsychToolbox* (Brainard, 1997) and in-house software was used to calibrate the display system and run the experiments. Stereo viewing was achieved by using frame-interleaving goggles (FE1; CRS Ltd) to present separate images to the two eyes. The high frame rate (150 Hz; 75 Hz per eye) ensured that the alternating display appeared as a steady 3-D image with no visible flicker. Luminance measurements and calculations showed that physical crosstalk between the two eyes' views was very low. We calculated that for a sine-wave grating less than 1% of its contrast would effectively 'leak' through to the other eye.

To create the appearance of horizontal corrugations, with very small sinusoidal modulations of disparity, each row of the texture had to be shifted to the left or right by a small (often sub-pixel) distance that was a sinusoidal function of the row's vertical position (y) in the image. These small shifts were achieved by blurring each row separately with a Gaussian kernel whose space constant (σ) was 1 min arc, and whose peak was displaced by $\pm\delta/2$ for the left and right eyes respectively, where δ is the desired disparity of the row. Thus the convolution (blur) kernel r was defined by:

$$r(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(\frac{-(x \pm \delta/2)^2}{2\sigma^2}\right)$$

where x is sampled in 1-pixel steps. To ensure that the blurring was isotropic (circularly symmetric), a similar blurring was then imposed on each column of pixels, but with no spatial offset. We evaluated the adequacy of sub-pixel resolution, and the rendering of disparity, by plotting the difference between the Fourier phase spectra of corresponding rows for the left and right eyes, using *Matlab*'s fft function. For a given disparity, the phase difference should increase linearly with spatial frequency, with a predictable slope. We concluded from many such analyses, simulating graphics systems with different greyscale resolutions and allowing for the influence of gamma correction, that the precision of disparity rendering with our 14-bit greyscale was good down to disparities of about 0.1–0.2 s. This is important because it ensured that very low DM thresholds (as low as 2–3 s) could be measured accurately, and that disparity-defined surfaces would always be smooth rather than step-like even at very low pedestal amplitudes (as low as 1 s). The rendering of disparity was about two orders of magnitude less precise when we simulated a standard 8-bit graphics card.

At the viewing distance of 114 cm, the image texture subtended 10×10 deg square, set in a uniform background (17 deg wide - 13 deg high) of the same mean luminance (26 cd/m^2 when viewed through the frame-interleaving goggles). The RMS contrast of the texture (standard deviation of luminance divided by mean luminance) was 0.2. The spatial frequency f of disparity modulation was either 0.3 or 0.6 c/deg in different experiments. These spatial frequencies are close to the peak of disparity modulation sensitivity (Bradshaw & Rogers, 1999; Schumer & Julesz, 1984; Tyler, 1974), at least for central vision (Prince & Rogers, 1998). In this paper we quantify the disparity modulation by its sinusoidal amplitude a in sec arc, i.e. half the difference between the peak and trough disparities. Thus the stereo surface was defined by its disparity profile:

$$\delta(y) = a.\cos(2\pi fy - \varphi)$$

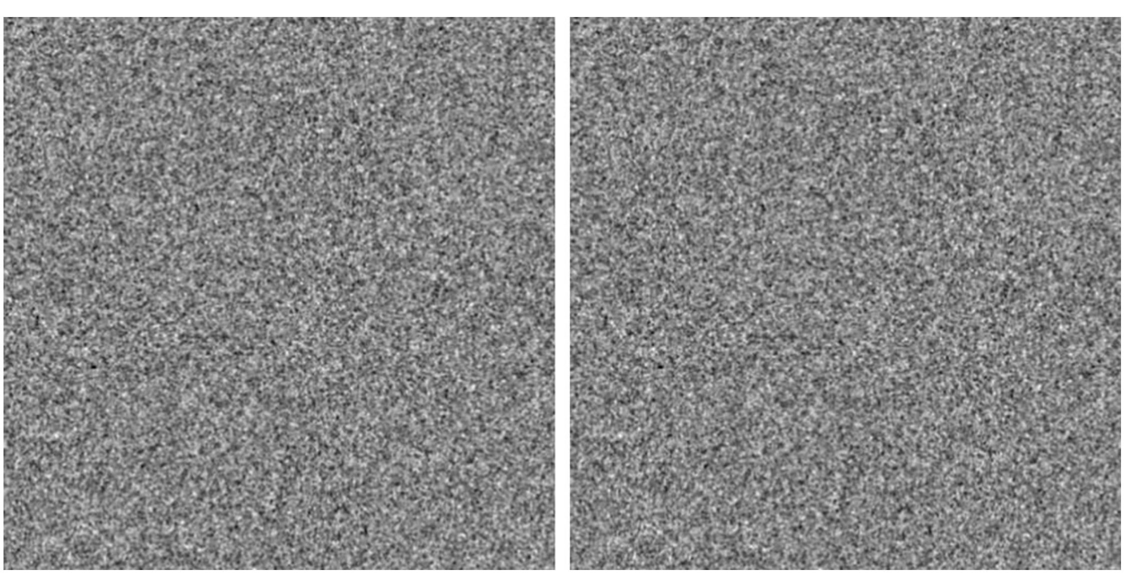

Fig. 1. Illustrates the type of slightly blurred Gaussian noise texture used in the experiments. Fusion of the left and right images should reveal a horizontal corrugated surface, produced by sinusoidal modulation of disparity. Here there are 1.5 cycles of modulation; in the experiments there were 3 or 6 cycles. The 2AFC task was always to report which of two presentations had the greater depth.

where δ, a, f and y are defined above, and φ is the phase of the modulation relative to the screen centre ($y = 0$).

Following several previous studies (e.g. Prince & Rogers, 1998; Tyler & Kontsevich, 2001) we used a 500 ms presentation time that made the observer's task comfortable. At this duration we cannot exclude the possibility of vergence movements, but studies with briefer durations (e.g. 100 ms, Schumer & Julesz, 1984) have found DM sensitivity curves quite similar to those at longer durations. Within one trial, the corrugated surface (DM stereogram) was presented twice, with a 500 ms inter-stimulus interval. The two intervals, marked by audible tones, contained different random samples of texture but more importantly they differed in the DM amplitude presented: one interval contained the pedestal, the other contained the pedestal plus an increment in modulation that was the test signal. The phase φ was the same in the two intervals. The 2AFC task was to discriminate the level of modulation – which of the two presentations had the greater depth? A high- or low-tone gave feedback about correctness. A small, dark, central fixation point – acting as a cue to convergence and fixation – was shown on a uniform (mean luminance) field at all times except during the texture presentations. The test disparity amplitude was varied in 2 dB steps from trial to trial using a standard staircase procedure (3-correct-down, 1-wrong-up). The proportions of correct responses at each test level were fitted with a Weibull function using a maximum likelihood procedure to estimate, for each pedestal level, the threshold level of added modulation that would give 81.6% correct performance.

2.1. Experiment 1

We tested a total of 10 pedestal amplitudes, from 1 to 512 s arc, as well as a no-pedestal condition. Since an accurate estimate of baseline threshold was important, the no-pedestal case (detection task) was tested twice as often as the other conditions: 240 trials instead of 120 trials per subject per threshold estimate at each phase. In each session, three different pedestal levels at each of the two phases were tested in six randomly interleaved staircases. Thus the phase φ of the corrugation – an apparent ridge or trough at the centre of the screen – varied between trials.

2.2. Experiment 2

Procedure was similar to experiment 1 in most respects, but phase was constant within-session. This apparently minor detail turned out to be a key factor. Based on experiment 1, pedestal levels were chosen to be in the expected region for facilitation: 0, 2, 4, 8, 16 s arc only. Spatial frequency was 0.6 c/deg.

3. Results

Three observers were tested, including one author and two other experienced observers. Fig. 2 shows DM increment thresholds as a function of pedestal amplitude, for the three individuals and as the group average, with $f = 0.3$ c/deg. A 'disparity dipper' is clearly evident in these results: near-threshold pedestals made depth differences easier to see. Thresholds were about a factor of 2 lower than the average detection baseline. For all three observers (Fig. 2), deeper pedestals ($a > 20$ s arc) made changes in depth progressively harder to see, in a manner analogous to contrast masking.

For one observer, TAY, experiment 1 was repeated with $f = 0.6$ c/deg. The baseline sensitivity, facilitation and masking effects (Fig. 3) were all similar to those obtained at 0.3 c/deg (Fig. 2), suggesting that there was no important difference in DM processing between the two spatial frequencies. In Fig. 2, when judged against the group mean baseline, all three observers appeared to show facilitation in the region $1 < a < 10$ s arc. But Fig. 4 (top row) illustrates that, relative to their individual baseline thresholds, observers DHB and TAY showed similar amounts of facilitation, while SAW showed only minor facilitation.

In experiment 1, the test phase φ was randomized (0 or 180 deg) from trial to trial. We show later how, in the framework of signal detection theory, this kind of uncertainty about the stimulus can lead to an apparent facilitation in the presence of a pedestal. Such a view clearly implies that the observed facilitation (Figs. 2 and 3) may be partly or wholly a consequence of the experimental procedure itself. A critical test, therefore, is to examine the pedestal effect with phase uncertainty removed. In experiment 2, with $f = 0.6$ c/deg, the same three observers were tested with pedestals in the range $a = 2$–16 s arc, where facilitation was observed in experiment 1, but now the two phases ($\varphi = 0$ or 180 deg) were tested in separate, interleaved sessions.

Fig. 4 (bottom row) shows that, when phase did not vary from trial to trial, little or no facilitation was seen for any observer. This lack of facilitation (experiment 2) is similar to the findings of Lunn & Morgan (1997), who also used the same phase of corrugation on every trial. [We might also expect that removal of uncertainty would improve (lower) the baseline thresholds in experiment 2, but because the spatial frequencies were different in the main parts of experiments 1 and 2 (0.3 vs. 0.6 c/deg) we cannot test that prediction against our data].

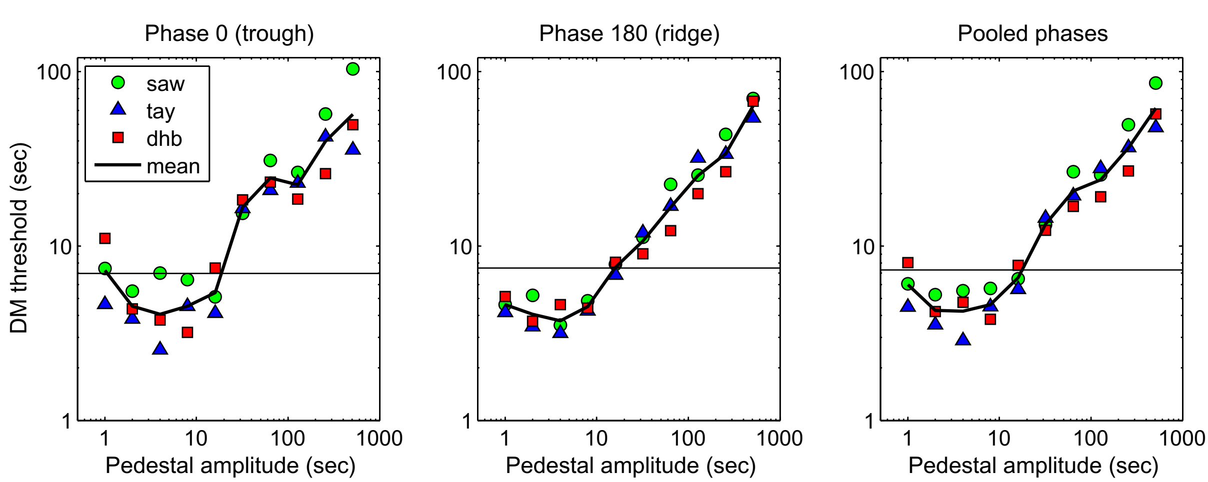

Fig. 2. Experiment 1: disparity modulation (DM) increment thresholds as a function of pedestal amplitude. Symbols indicate three different observers; solid line joins the geometric means. Thresholds were estimated separately for the two test phases (left and centre panels), or after pooling the raw data for both phases (right panel). Horizontal line is the mean baseline threshold (no pedestal). Spatial frequency was 0.3 c/deg.

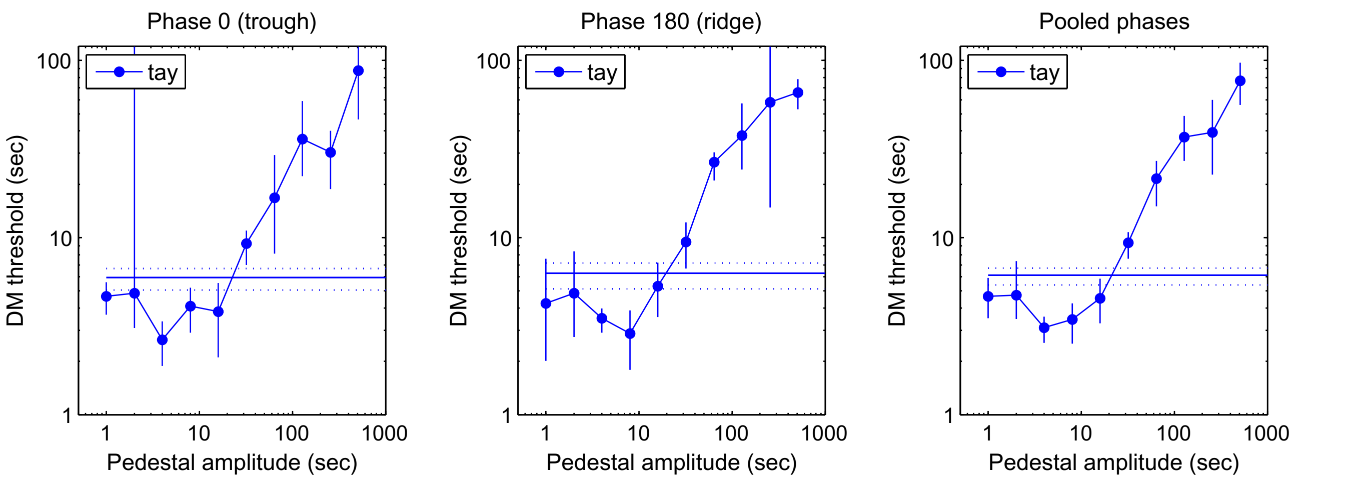

Fig. 3. Disparity modulation (DM) increment thresholds for one observer (TAY) at 0.6 c/deg, under the conditions of experiment 1. Error bars and horizontal dashed lines are 95% confidence limits obtained by a parametric bootstrapping method (Wichmann & Hill, 2001).

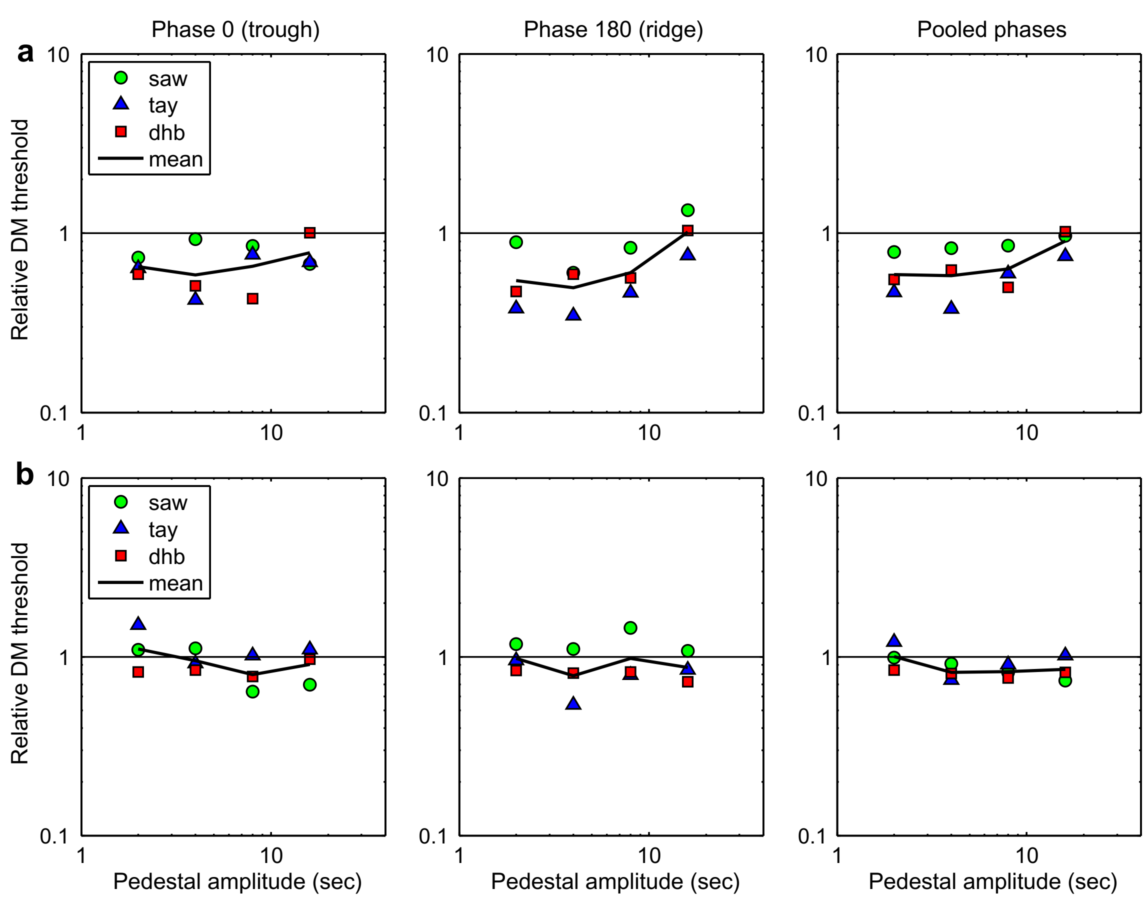

Fig. 4. Disparity modulation (DM) increment thresholds, expressed relative to baseline thresholds for each individual observer, as a function of pedestal amplitude. Symbols indicate three different observers; solid line joins the geometric means. Thresholds were estimated separately for the two test phases (left and centre panels), or after pooling the raw data for both phases (right panel). Top row: results of Experiment 1 re-plotted from Fig. 2, for pedestal amplitudes in the facilitation region only (a = 2, 4, 8 and 16 s arc). Bottom row: results of Experiment 2 revealed little or no facilitation when the two phases were tested in separate sessions.

4. Discussion

4.1. Facilitation

A weakly corrugated surface (the 'pedestal') reliably improved disparity modulation thresholds in a 2AFC task, but only when the phase of modulation (ridge or trough at the centre of gaze) varied from trial to trial. Consistent with Lunn and Morgan (1997), experiment 2 showed that facilitation was weak or absent when the same phase of corrugation was used from trial to trial. Taken together, the two studies suggest that limitations imposed by graphics hardware or choice of pedestal levels (see Section 1) were not a major factor. Instead our results suggest that, for disparity modulation, an apparent facilitation arises from the effect of stimulus uncertainty – the phase uncertainty in experiment 1 – that was absent in experiment 2. We now develop that idea in the context of signal detection theory (SDT), and apply a simple model to the experimental results.

In a 2AFC experiment, the observer compares two observations (X_1, X_2), one from each interval, and must decide which was more likely to be the signal. Fig. 5 offers a representation of the 2AFC experiment, with and without the pedestal, that helps to explain how stimulus uncertainty can produce the facilitation effect. For this exposition, we suppose that each stimulus has a sign and a magnitude that is encoded by a noisy bipolar channel – defined as one that can encode the sign of the stimulus by the sign of the response, X (Klein, 1985). For our experiment the sign represents phase (0 or 180 deg). Whether the sensory channel encodes disparity, or disparity difference, or disparity curvature (cf. Lunn & Morgan, 1997) or some other quantity is not important here. If the sign of the signal is known (say, positive, S+, Fig. 5A), then the observer's best bet is always to choose the more positive observation. Given that one observation was drawn from S+ and the other from N, and assuming $X_2 > X_1$, then X_2 is more likely than X_1 to have been drawn from the S+ distribution. [More formally, it is the likelihood ratio, $f(X_i|S+)/f(X_i|N)$ that is higher for X_2 than X_1. Green and Swets (1966, p. 46), show that to maximize the proportion correct the observer's decision about the signal should always favour the higher likelihood ratio. This is true provided that the likelihood ratio is a monotonic function of X]. This decision rule remains optimal when a pedestal is present (Fig. 5B). Moreover, if the transducer (defined as the function that maps stimulus value onto mean response X) is *linear*, then (by definition) responses are additive, so the pedestal adds the same amount to the N and S+ means, and their difference (which is the discriminability, d') remains unchanged. Thus, with a linear transducer and no phase uncertainty, the pedestal should have no effect on discrimination.

The situation is very different when the sign of the signal is not known (Fig. 5C). Now the sign of the response is uninformative: the more positive response is more likely to be the signal only if the signal is positive, but that remains unknown. The best that can be said is that non-signal responses (drawn from N) will tend to be closer to zero, while signal responses (drawn from S+ or S−) tend to deviate more from zero in either direction. The rational observer should ignore the sign and choose the response with the greater absolute magnitude. [More precisely, it can be shown that the likelihood ratio is not a monotonic function of X in this situation, but it is a monotonic function of $abs(X)$. Therefore, choosing the greater absolute value is the optimal strategy]. This loss of sign represents a loss of information, and so performance must be worse with phase uncertainty. Errors will be made when, as in Fig 5C, the sign would have led to the correct response (X_2) but using magnitude alone leads to the wrong response (X_1). When the pedestal is present, these potential disagreements between sign and magnitude as the basis for decision will be less frequent (Fig. 5D). When the pedestal is sufficiently large, most pairs of responses (X_1, X_2) will have the same sign, and then unsigned magnitude is a good basis for the decision, even when sign is not known in advance. When sign is unknown, adding the pedestal should improve performance, but when sign is known, the pedestal should have no effect. It is perhaps surprising that knowledge of the sign is important, even though the sign of the stimulus does not have to be reported.

To confirm and quantify these ideas, we ran a Monte Carlo simulation. The sensory channel was linear (mean response proportional to mean signal), perturbed by additive, zero-mean Gaussian noise with equal variance at all signal levels. Without loss of generality, the signals were always positive, but of course with noise the responses might be positive or negative. For the phase-known case, the model observer exploited the sign of response on each 2AFC trial by choosing the more positive response as the signal. In the phase-uncertain case the model observer chose the

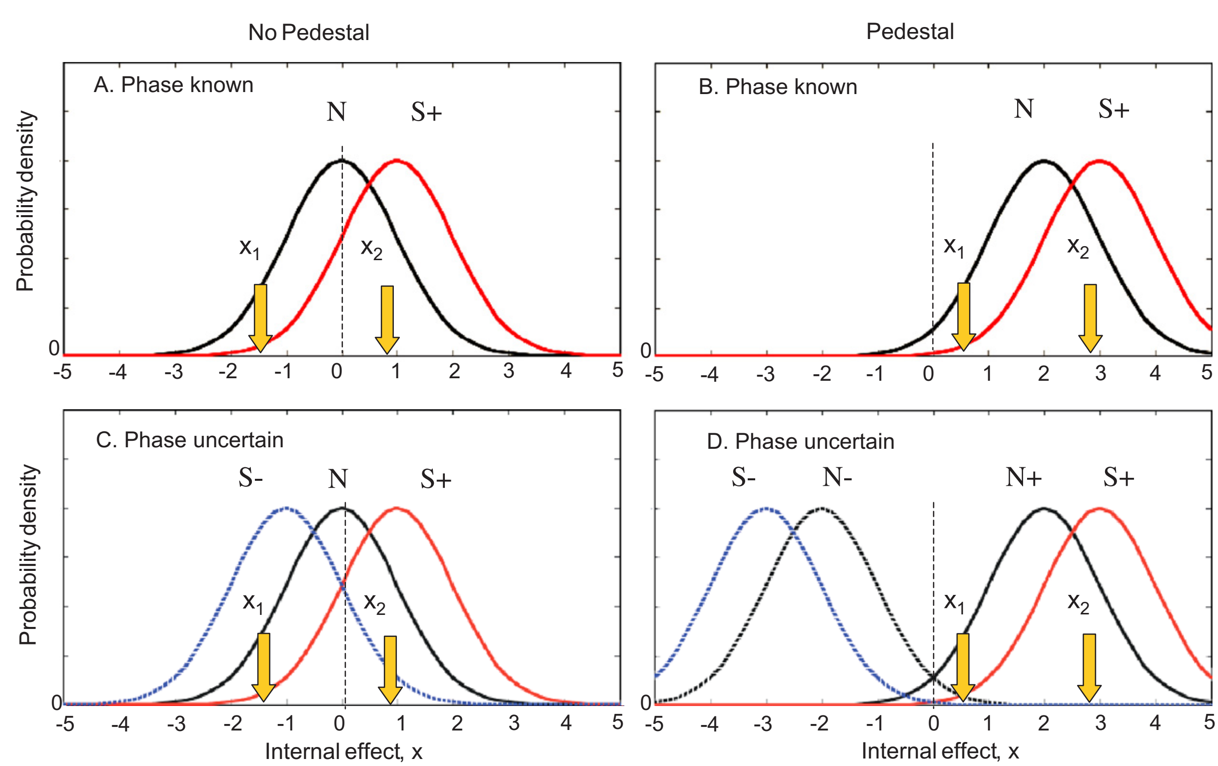

Fig. 5. Application of signal detection theory to the pedestal experiment. Each 2AFC trial yields two observations (X_1, X_2). The observer must decide which was more likely to have been drawn from the signal distribution, S. The non-signal or pedestal distribution is N. A rational observer will behave differently when phase is known (top row) and when it is uncertain (bottom row). C: When phase is uncertain, the response to the signal is drawn from S+ or S−, on different trials. D: Phase uncertain, like C, but with a pedestal; responses to the signal + pedestal and pedestal alone are drawn from (S+, N+) or from (S−, N−), on different trials. See text for details.

response with the greater absolute magnitude, as discussed above. There were 10,000 simulated trials at each of 11 test amplitudes. Weibull functions were fitted to the psychometric function (proportion of correct trials), and thresholds were plotted against pedestal level, just as for the experimental data.

Fig. 6 shows that this model does indeed predict facilitation by the pedestal when stimulus phase is uncertain, but no facilitation when phase is known. The magnitude of the predicted facilitation (Fig. 6, right) is similar to that observed in experiment 1. At the higher pedestal levels, predicted facilitation levels off at 3.3 dB below the baseline (and this was confirmed over a much wider range of pedestal levels; not shown). Thus the model predicts the observed facilitation fairly well, but does not predict the transition from facilitation to masking (Figs. 2 and 3). We should not expect this model to predict masking, because some nonlinearity is needed to account for the masking phenomenon, and at higher pedestal levels this simple model has none.

The experimental data (Fig. 4, bottom right) showed just a hint of facilitation even when phase was known. Such an effect can be explained by additional uncertainty in the observation process itself (Pelli, 1985). For a given stimulus, some sensory channels carry informative responses, while others do not, but the observer may be uncertain about which are the relevant channels. For example, when signal phase is 0 or 180 deg, a mechanism tuned to 90 deg spatial phase carries no signal and so is non-informative, but the observer may not know this. Monitoring irrelevant channels adds noise to the decision process (Tyler & Chen, 2000). When only one of M channels is informative, and all the channels are independently noisy, the observer should use a MAX rule: choose the interval that gave the greatest response taken across all M channels. The model of Fig. 6 is the special case where $M = 1$ (no irrelevant channels). Fig. 7 shows that with just one irrelevant channel and one informative channel, the simple noisy decision model outlined here accounts well for the slight (<2 dB) pedestal effect when phase was known, and the more robust (4–5 dB) pedestal effect obtained when phase was uncertain.

We also ran a version of the model where each bipolar channel was replaced by two monopolar channels (cf. Klein, 1985) whose mean response to a signal was always positive. One channel was sensitive only to 0 phase while the other was sensitive only to 180 phase. The MAX decision rule remained the same as before. When the model observer monitored just the relevant channels (dashed curves in Fig. 8), facilitation was predicted when phase was uncertain, but the degree of facilitation was under-estimated. When internal uncertainty was added, analogous to Fig. 7, by assuming equal numbers of relevant and irrelevant channels (solid curves in Fig. 8), facilitation increased but the observed difference between phase-known and phase-uncertain was under-estimated.

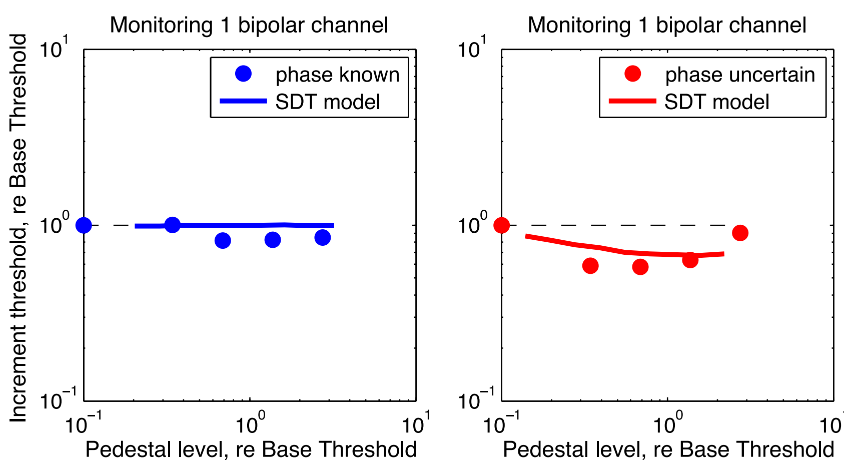

Fig. 6. SDT model predictions (curves) for the pedestal experiment, compared with experimental results (symbols; group means from Fig. 4, right-hand panels). Thresholds and pedestal levels are expressed as multiples of the baseline detection threshold (no pedestal; leftmost data point). Left: phase known. Right: phase uncertain.

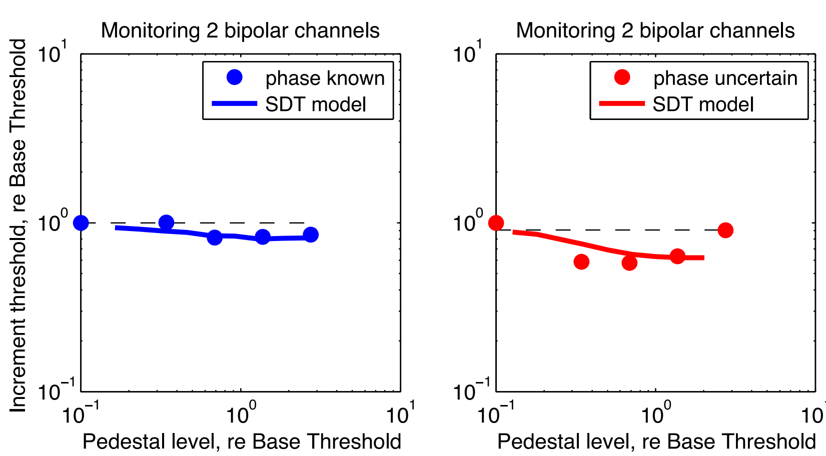

Fig. 7. As Fig. 6, but an extended model in which the observer monitors one informative channel (as in Fig. 6) and one non-informative channel that carries only noise. This additional uncertainty in the observing process produces slight facilitation even when phase is known (left panel).

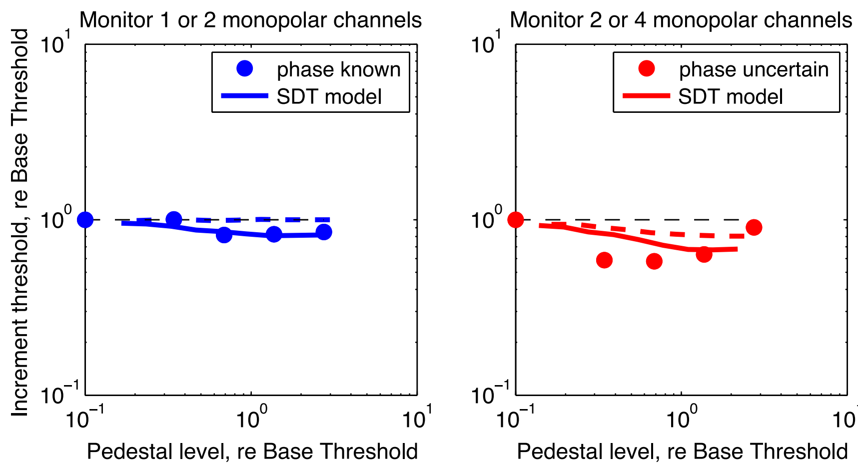

Fig. 8. Similar to Figs. 6 and 7, but for model with monopolar channels that respond to one or other phase but not both. Dashed curves: model observer monitors only the 1 (left) or 2 (right) task-relevant channels. Solid curves: monitor 1 relevant & 1 irrelevant channel (left) or 2 relevant & 2 irrelevant channels (right). See text for details.

These are small experimental effects, and so some caution is required, but in general the model based on bipolar channels (Figs. 6 and 7) seemed better able to capture the amount of facilitation and its dependence on phase uncertainty.

4.2. Facilitation and masking

Only one other study, to our knowledge, has found facilitation by disparity pedestals. Farell, Li, & McKee (2004) studied disparity increment thresholds for planar images (gratings and random textures), using several psychophysical procedures. For gratings, they often found small 'dips' (facilitation) at non-zero pedestal disparities up to about 3x threshold. Their results are at least partly consistent with our account, because most of their procedures involved uncertainty, either about the sign of the pedestal, or the sign of the test increment/decrement. For random-dot textures (their Fig. 4) no dip was evident, but nor would it be expected because the pedestal disparities were all 10–100 times threshold – in the 'masking' region.

For corrugated surfaces, modulation thresholds in the masking region rise as a power function of pedestal amplitude, with an exponent of about 0.5–0.7 (Lunn & Morgan, 1997, their Fig. 2; also Figs. 2 and 3 here). This behaviour is very similar to contrast discrimination of gratings in the luminance domain, and so analogies between the visual processing of luminance and disparity (Lunn & Morgan, 1995) remain interesting. Models to account for the rise of disparity increment thresholds with pedestal disparity, reviewed by Farell et al. (2004), typically involve proposals about the receptive-field structure and disparity-tuning of binocular neurons and, at present, are expressed in very different terms from models for the analogous contrast masking. A general understanding of the basis for DM sensitivity is beginning to emerge, however, from combined physiological, psychophysical and theoretical analyses of V1 neurons and their responses to corrugated textures (Banks, Gepshtein, & Landy, 2004; Nienborg, Bridge, Parker, & Cumming, 2004).

Wright & Ledgeway (2004) were, like us, interested in the possible summation of luminance shading and stereo cues to depth in corrugated gratings. They found that luminance gratings never masked the detection of disparity modulation (DM), and in some conditions (especially where luminance and disparity gratings had the same spatial frequency, 0.4 c/deg) the luminance grating facilitated the DM threshold. They argued that this was occurring not at the level of depth cue combination, but rather that the luminance grating in a known phase reduced spatial uncertainty in the DM task. When luminance phase was randomized from trial to trial, it did not facilitate DM detection. Evidence for the spatial cueing argument came from an experiment in which the luminance grating was replaced by thin lines of texture, again spaced at 0.4 c/deg. These lines carried neither shading nor stereo depth cues. They had no influence when they lay at the zero-crossings of disparity modulation, but they facilitated DM detection a little when they lay at the disparity peaks and troughs. The most robust facilitation, however, came when the lines lay only at the disparity troughs. In this condition the DM phase is known, while in the other two cases there remained a 0/180° phase uncertainty. Thus these DM facilitation effects, like ours, can be largely explained by reduction of phase uncertainty, albeit by a different route.

5. Conclusions

In the detection and discrimination of stereo corrugations, we observed facilitation of performance by weak pedestals (the 'dipper' effect), but only when the phase of corrugation was uncertain from trial to trial. When phase was known, the pedestal had little or no effect on the disparity modulation threshold. The latter result confirms the findings of Lunn & Morgan (1997). Simulations that assumed a linear response to disparity and a noisy decision process showed that external, stimulus uncertainty was sufficient to account for most of the observed facilitation, but some degree of internal uncertainty – monitoring a sensory channel that is non-informative – might also be present. On this view, the pedestal does not facilitate performance through some sensory process, such as nonlinear transduction. Rather, removing the pedestal hampers performance, but only when phase is uncertain, because the observer can no longer use the sign of response to guide the decision. It follows that the visual mapping from binocular disparity to a depth response is probably linear, with no threshold-like nonlinearity, at least for small disparities up to about ±20 s arc.

Acknowledgements

We thank the Engineering and Physical Sciences Research Council (U.K.) for grant support (GR/S07254/01, GR/S07261/01).

References

Banks, M. S., Gepshtein, S., & Landy, M. S. (2004). Why is spatial stereoresolution so low? *Journal of Neuroscience, 24*(9), 2077–2089.
Bradshaw, M. F., & Rogers, B. J. (1999). Sensitivity to horizontal and vertical corrugations defined by binocular disparity. *Vision Research, 39*(18), 3049–3056.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10*(4), 433–436.

Farell, B., Li, S., & McKee, S. P. (2004). Disparity increment thresholds for gratings. *Journal of Vision, 4*(3), 156–168.

Foley, J. M. (1994). Human luminance pattern-vision mechanisms – masking experiments require a new model. *Journal of the Optical Society of America A, 11*(6), 1710–1719.

Foley, J. M., & Chen, C. C. (1999). Pattern detection in the presence of maskers that differ in spatial phase and temporal offset: Threshold measurements and a model. *Vision Research, 39*(23), 3855–3872.

Georgeson, M. A., & Georgeson, J. M. (1987). Facilitation and masking of briefly presented gratings - time-course and contrast dependence. *Vision Research, 27*(3), 369–379.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley & Sons.

Klein, S. A. (1985). Double-judgment psychophysics – problems and solutions. *Journal of the Optical Society of America A, 2*(9), 1560–1585.

Legge, G. E., & Foley, J. M. (1980). Contrast masking in human vision. *Journal of the Optical Society of America, 70*(12), 1458–1471.

Lunn, P. D., & Morgan, M. J. (1995). The analogy between stereo depth and brightness – a reexamination. *Perception, 24*(8), 901–904.

Lunn, P. D., & Morgan, M. J. (1997). Discrimination of the spatial derivatives of horizontal binocular disparity. *Journal of the Optical Society of America A, 14*(2), 360–371.

Nachmias, J., & Sansbury, R. V. (1974). Grating contrast: Discrimination may be better than detection. *Vision Research, 14*, 1039–1042.

Nienborg, H., Bridge, H., Parker, A. J., & Cumming, B. G. (2004). Receptive field size in V1 neurons limits acuity for perceiving disparity modulation. *Journal of Neuroscience, 24*(9), 2065–2076.

Pelli, D. G. (1985). Uncertainty explains many aspects of visual contrast detection and discrimination. *Journal of the Optical Society of America A, 2*(9), 1508–1532.

Prince, S. J. D., & Rogers, B. J. (1998). Sensitivity to disparity corrugations in peripheral vision. *Vision Research, 38*(17), 2533–2537.

Schofield, A. J., & Georgeson, M. A. (1999). Sensitivity to modulations of luminance and contrast in visual white noise: Separate mechanisms with similar behaviour. *Vision Research, 39*(16), 2697–2716.

Schofield, A. J., Hesse, G., Rock, P. B., & Georgeson, M. A. (2006). Local luminance amplitude modulates the interpretation of shape-from-shading in textured surfaces. *Vision Research, 46*(20), 3462–3482.

Schumer, R. A., & Julesz, B. (1984). Binocular disparity modulation sensitivity to disparities offset from the plane of fixation. *Vision Research, 24*(6), 533–542.

Tyler, C. W. (1974). Depth perception in disparity gratings. *Nature, 251*(5471), 140–142.

Tyler, C. W., & Chen, C. C. (2000). Signal detection theory in the 2afc paradigm: Attention, channel uncertainty and probability summation. *Vision Research, 40*(22), 3121–3144.

Tyler, C. W., & Kontsevich, L. L. (2001). Stereoprocessing of cyclopean depth images: Horizontally elongated summation fields. *Vision Research, 41*(17), 2235–2243.

Wichmann, F. A., & Hill, N. J. (2001). The psychometric function. II. Bootstrap-based confidence intervals and sampling. *Perception & Psychophysics, 63*(8), 1314–1329.

Wright, M., & Ledgeway, T. (2004). Interaction between luminance gratings and disparity gratings. *Spatial Vision, 17*(1-2), 51–74.