PERGAMON

Vision Research 40 (2000) 1157–1165

Vision Research

www.elsevier.com/locate/visres

Local luminance factors that determine the maximum disparity for seeing cyclopean surface shape

Lynn R. Ziegler *, Frederick A.A. Kingdom, Robert F. Hess

Department of Ophthalmology, McGill Vision Research, McGill University, 687 Pine Avenue West, H4-14, Montreal, Qué., Canada H3A 1A1

Received 27 April 1999; received in revised form 1 October 1999

Abstract

We measured the maximum disparity grating amplitude (d_{max}) for seeing cyclopean surface shape, using stereograms made from dense arrays of micropatterns, whose luminance characteristics were manipulated. In Experiment 1, we used disparity gratings made from Gabor micropatterns. D_{max} was found to vary inversely both with luminance spatial frequency and with Gabor size, but was constant for a constant bandwidth (frequency times size). To test whether this was due to changes in bandwidth per se or to changes in the number of local features, in Experiment 2 we manipulated the local feature content with a range of micropatterns that we termed 'edgels'. The results supported neither hypothesis. In Experiment 3 we varied the phases of the Fourier components of square wave edgels, thereby introducing more features, and we found that this did not change d_{max}. Taken together, our results show that d_{max} decreases with an increase in the number of local luminance cycles at each luminance scale. D_{max} is mainly limited by false target matching between similar components of the micropatterns. Stereopis, in terms of surface shape perception, is served only by first order mechanisms, and only by luminance filters that are broadband.

Keywords: Stereopsis; Binocular correspondence; Size-disparity correlation; Spatial frequency; Contrast envelope

1. Introduction

Disparity gratings are typically formed from random arrays of elements, such as dots or pixels, whose disparities are varied sinusoidally across the visual field (Tyler, 1974). Such gratings can provoke a compelling impression of shape, a corrugation in depth. When the disparity amplitude is slowly increased, a point is reached where the corrugation shape abruptly collapses. What factors determine this maximum disparity (d_{max}) for cyclopean shape? In a previous study (Ziegler, Hess & Kingdom, 2000) we examined how d_{max} is affected by global factors, such as disparity spatial frequency and the gradient of disparity. Here we examine how d_{max} is affected by local factors, specifically the luminance defined characteristics of the micropatterns that comprise our disparity gratings.

To our knowledge there is only one previous study that has examined how local luminance factors determine d_{max} for stereoscopic shape. Pulliam (1982) used luminance sinewave gratings that were vertically oriented and had a constant luminance spatial frequency. Stereo shape was introduced by sinusoidal modulation of the phase of the gratings in opposite directions in the two eyes' views. Because disparity varied along the vertical axis, viewers saw a horizontal corrugation in depth. Pulliam found that d_{max} was constant with respect to luminance spatial frequency. His gratings however were not cyclopean because they contained a monocular cue, vernier displacement.

Most studies have investigated d_{max} not for cyclopean shape, but for isolated micropatterns, or for luminance gratings that were planar (flat). The results from these studies are conflicting, and they suggest a number of factors that might also affect d_{max} for cyclopean shape. Schor, Wood and Ogawa (1984), for example, used single difference of Gaussian (DOG) elements and found that d_{max} varied inversely with luminance spatial frequency over part of the range tested. Over another part of the range however d_{max} was constant. Hess and Wilcox (1994) and Wilcox and Hess (1995) required

* Corresponding author. Fax: + 1-514-843-1691.

E-mail address: zieg@vision.mcgill.ca (L.R. Ziegler)

0042-6989/00/$ - see front matter
PII: S0042-6989(99)00213-8

observers to judge the depth, either near or far from the point of fixation, of single micropatterns. They concluded that d_{max} was determined by micropattern envelope size. Glennerster (1998) found that element density affected d_{max} in a study that also used a near/far task and a target consisting of a subset of random dots. Glennerster argued however that the effect may have been due to the average luminance spatial frequency changing with dot density. With dots it is impossible to determine which factor, density or spatial frequency, determines d_{max}.

Other studies have concluded that d_{max} may be determined by local features, such as bars or edges. Boothroyd and Blake (1984) and Mayhew and Frisby (1981) used a planar missing fundamental grating that provided an edge at the spatial frequency of the fundamental. Although their tasks did not involve shape perception, both studies concluded that binocular correspondence could occur at disparities beyond that possible for any single component of the missing fundamental stimuli when presented alone. This suggested that stimulus features such as edges may have a special role in establishing binocular correspondence and in determining d_{max}.

The above studies using either isolated elements or planar gratings suggest a number of possible factors determining d_{max}, namely spatial frequency, envelope size, and local features. Evidence has been reported previously however that the rules can be different for seeing cyclopean shape than for seeing the depth of an isolated element (Ziegler & Hess, 1999). Even if there were unanimous agreement as to what factors determine d_{max} for single element or planar grating tasks, those conclusions may not necessarily generalize to cyclopean shape perception.

In order to resolve this issue we examined which of the above local factors determined d_{max} for seeing cyclopean shape. To insure that judgments were based on cyclopean shape we used a criterion-free method that required observers to judge the orientation of a cyclopean disparity grating (Hess, Kingdom & Ziegler, 1999). In each experiment we used random arrays of specific micropatterns so that we could apply various combinations of the above mentioned factors.

2. Methods

2.1. Subjects

The three authors acted as observers. They had normal or corrected-to-normal visual acuity and normal stereopsis, and were experienced psychophysical observers.

2.2. Apparatus

All stimuli were generated using a Silicon Graphics O2 and displayed on its monitor (Sony GDM-20E21). Observers wore LCD shutter-glasses (StereoGraphics Inc. CrystalEyes) synchronized to the alternating stereo half-images, at 60 Hz for each eye. The linearity of the monitor was measured using a photometer with a photometric head (United Display Technologies S370) and gamma corrected.

2.3. Viewing conditions

Observers sat at a viewing distance of 57 cm and the natural ambient illumination was low. The display subtended $28 \times 36°$ at a resolution of 1024×1280 pixels. Measured through a shutter-glass lens, mean luminance was 6.0 cd/m^2.

2.4. Stimuli

Stimulus duration was 117 ms (14 stereo frames), below vergence latencies (Rashbass & Westheimer, 1961; Stevenson, Cormack & Schor, 1994). Each stereo half-image consisted of a random array of identical micropattern elements. In the first experiment, these were Gabor micropatterns with carriers always in sine phase and with an average luminance equal to the background level. Pixel luminances were assigned with subpixel accuracy, and peak Michelson contrasts were 33% unless otherwise noted. In Experiment 1, for observer LZ, contrast was 66% for the narrowest bandwidth Gabors.

Each micropattern array was created by adding micropattern intensity levels, without their baseline (d.c.) component, to a buffer in the computer's main memory (details in Hess et al., 1999). This technique prevented spurious and potentially conflicting depth cues from micropattern occlusion, and avoided any artifacts from patch edges. The disparities of the micropatterns were modulated sinusoidally to produce a disparity grating. On each trial, the corrugation was oriented either left or right oblique, $\pm 26°$ from horizontal. These angles were chosen to be sufficient to allow observers to make orientation judgements. The phase of each disparity grating was assigned at random to be sure that responses were not based upon perceived depth at a single location.

In Experiment 1, over a range of disparity spatial (corrugation) frequencies, we measured d_{max} while independently varying each of three factors: density, Gabor size, and luminance spatial frequency. For each factor, three levels were used, a baseline level as well as greater and less than the baseline. When one factor was varied, the other factors were held constant at their baselines. The baseline for density was 1200 Gabors per screen, for luminance spatial frequency was 1.68 c/deg, and for Gabor size (Gaussian scale factor σ) was 0.36°.

2.5. Procedure

Mouse buttons were used to report perceived corrugation slant, left or right from vertical. Disparity amplitude began at 10 min and was adjusted automatically by a conventional staircase procedure, i.e. increased after two consecutive correct responses and decreased after every incorrect response, each by one-quarter octave (~ 19%). A staircase run terminated automatically after 12 reversals. The geometric mean of the last eight reversals yielded one estimate of d_{max}. Values reported here are geometric means of generally four or more such estimates.

3. Experiment 1: Gabor micropatterns

We measured the effect on d_{max} of luminance spatial frequency, Gabor density, and Gabor size, over a range of disparity spatial frequencies, with all other parameters held constant. The results are shown in Fig. 1. As can be seen, they were consistent across subjects and across disparity spatial frequencies. Fig. 1a shows that Gabor density had no effect on d_{max} over the range of disparity spatial frequencies, since the curves are not displaced. The remaining experiments in this study, unless noted, used stimuli with a constant density of 1200 micropatterns per display. Fig. 1b shows that, with Gabor luminance spatial frequency held constant, d_{max} varied inversely with Gabor size because the curves are, in this case, displaced. Fig. 1c shows that with Gabor size held constant, d_{max} varied inversely with luminance spatial frequency. We conclude from these results that d_{max} is determined both by micropattern size and by luminance spatial frequency, but not by micropattern density.

Our finding that Gabor size affects d_{max} might be because the Gabors were planar. That is, because each Gabor stereo-pair specified a single disparity, increasing Gabor size may have obscured the sinusoidal surface shape because of an averaging of depth when Gabor patches overlapped. If this were true, the decrease in d_{max} as Gabor size increased would be expected to become disproportionally larger at higher disparity spatial frequencies. That is, more depth averaging would occur as the Gabors extended between the peaks and troughs of the corrugations. In Fig. 1b, however, the

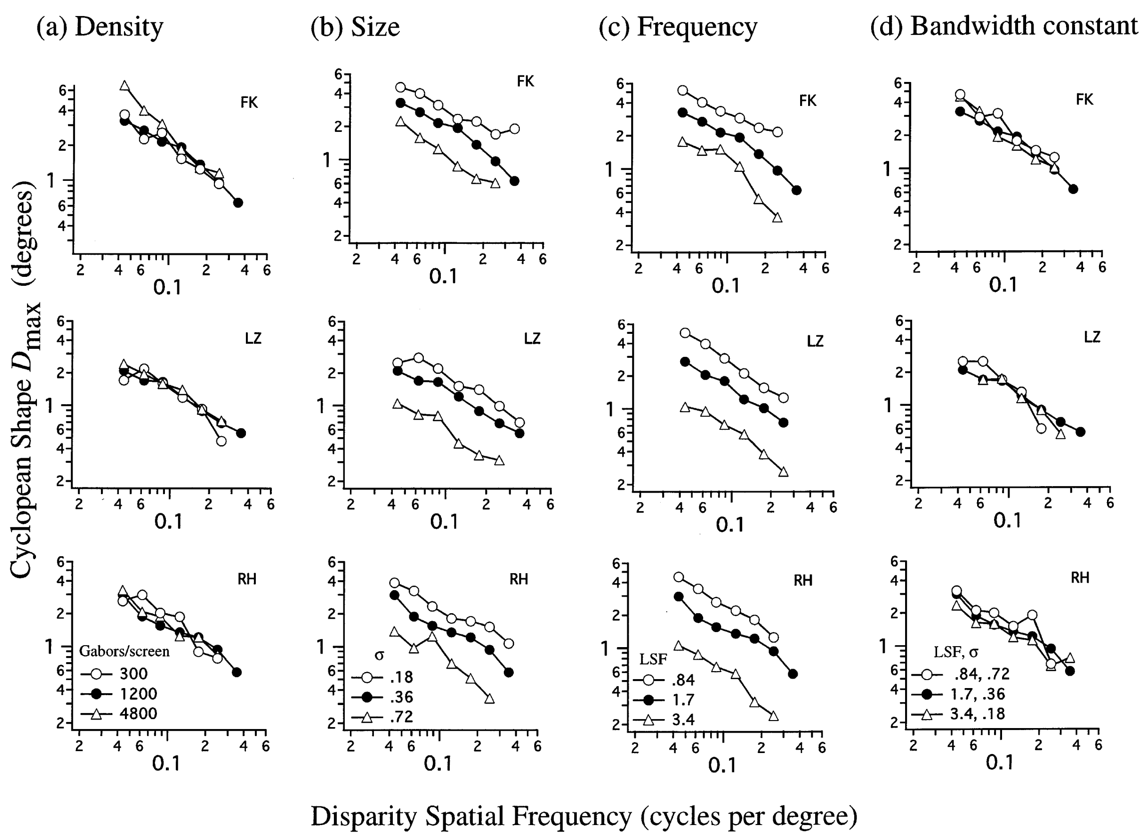

Fig. 1. D_{max} for cyclopean shape perception as a function of disparity spatial frequency (depth corrugation frequency) for three values of each factor: (a) density (Gabors per screen); (b) Gabor contrast envelope size (σ in degrees); (c) Gabor luminance spatial frequency (c/deg); and (d) Gabor luminance spatial frequency as in (c) except envelope size was also varied for a constant size-frequency product, or bandwidth (Experiment 1).

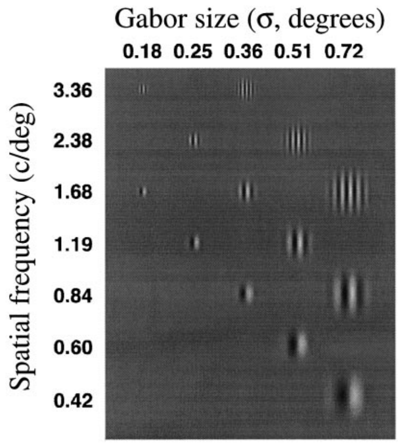

Fig. 2. Gabors used in the second part of Experiment 1. Along each (long) diagonal they are scaled replicas of each other, i.e. have the same bandwidth. In our experiment, the apparent size of these Gabors ranged from 0.5° (for σ 0.18°) to 2°.

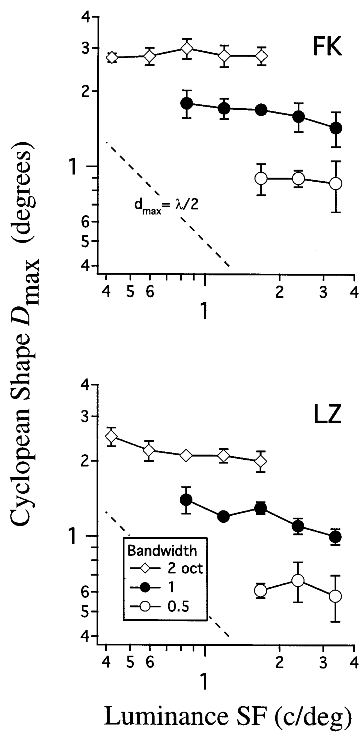

Fig. 3. D_{max} for cyclopean shape perception as a function of luminance spatial frequency, for three Gabor bandwidths. Symbols correspond with the Gabors shown in Fig. 2. Lines connect symbols for Gabors with the same size-frequency product. These iso-bandwidth curves correspond to the long diagonals in Fig. 2. For a particular luminance spatial frequency, Gabor contrast envelope size increases top to bottom (between curves). Disparity spatial frequency was held constant at 0.125 c/deg. Within each graph, the broken line represents the theoretical half-cycle limit.

functions of d_{max} with disparity spatial frequency for different Gabor sizes are parallel. This implies that the effect of Gabor size is not the result of Gabor planarity.

The effects on d_{max} of both Gabor size (Fig. 1b) and luminance spatial frequency (Fig. 1c) were equivalent, and this suggests that perhaps d_{max} is a function of the product of Gabor size and frequency. To test this possibility, we varied both frequency and size together, so that their product remained the same. That is, we used Gabors that were scaled replicas of each other and had the same bandwidth. The results are shown in Fig. 1d, and, as can be seen, the curves are not displaced, indicating that d_{max} did not vary when the Gabor size-frequency product was held constant.

To verify the generality of these results, we repeated our measurements using more combinations of frequency (0.42–3.36 c/deg) and size ($\sigma = 0.18$–$0.72°$), staggered in half-octave steps. These Gabors are shown in Fig. 2. Fixed values were used for corrugation frequency (0.125 c/deg) and density (1200 elements), while other conditions remained the same. The results confirm that an equal size-frequency product results in the same d_{max}, as shown in Fig. 3. When plotted with separate curves for each size-frequency product, there was again little effect of luminance spatial frequency on d_{max}.

These results all indicate that for disparity gratings made from Gabors, d_{max} is determined neither by frequency nor size alone, but by their product. This implies that bandwidth is the critical factor, but why should bandwidth determine d_{max}? Indeed, other things besides bandwidth vary with the product of Gabor size and frequency, for example, the number of cycles, or the number of luminance-defined features. Perhaps the increase in the number of cycles that occurs with a reduction in bandwidth makes binocular correspondence more difficult due to the *false target problem*, i.e. matching errors when similar features lie adjacent within a micropattern. We therefore decided to test whether d_{max} for cyclopean shape was determined by bandwidth per se or by the number of local features in the micropatterns.

4. Experiment 2: bandwidth or number of features?

4.1. Stimuli

The micropatterns used in this experiment are shown in Fig. 4. They were all windowed by the same Gaussian envelope with $\sigma = 0.36°$, and were based upon a fundamental frequency (1F) of 0.60 c/deg. These micropatterns were used to provide various combinations of (1) bandwidths and (2) numbers of local features, such as bars and edges. One of the micropatterns, containing a vertically oriented luminance step edge, we called a ‘hard’ edge element (HE) or ‘edgel’ (the Fourier spectrum consisted approximately of a fundamental frequency and its odd harmonics). It had the same number of edges but a larger bandwidth than the 1F Gabor. We also used a missing fundamental (MF) edgel, created by subtracting the 1F Fourier component from the HE. This edgel had a narrower bandwidth than the HE, but more edges. Finally, we used a Gabor with a frequency equal to the third harmonic (3F), that is, the second component of the HE. This had the narrowest bandwidth as well as the largest number of edges. Each edgel had a Michelson contrast of 33%, except for the MF edgel which was 54%, to be sure that the 5th harmonic (5F) component was visible.

We measured d_{max}-shape using each of these micropatterns over a range of disparity spatial frequencies. Depending on whether d_{max} was determined by (1) bandwidth, or (2) the number of features, two different patterns of results would be expected. If bandwidth alone determined d_{max} then d_{max} would be expected to decrease in the order HE > MF > 1F > 3F. On the other hand, if the number of features alone determined d_{max}, then the expected order would be (1F = HE) > MF > 3F. In particular, the number of features hypothesis predicts 1F > MF, while the bandwidth hypothesis predicts the opposite. Both predict MF > 3F.

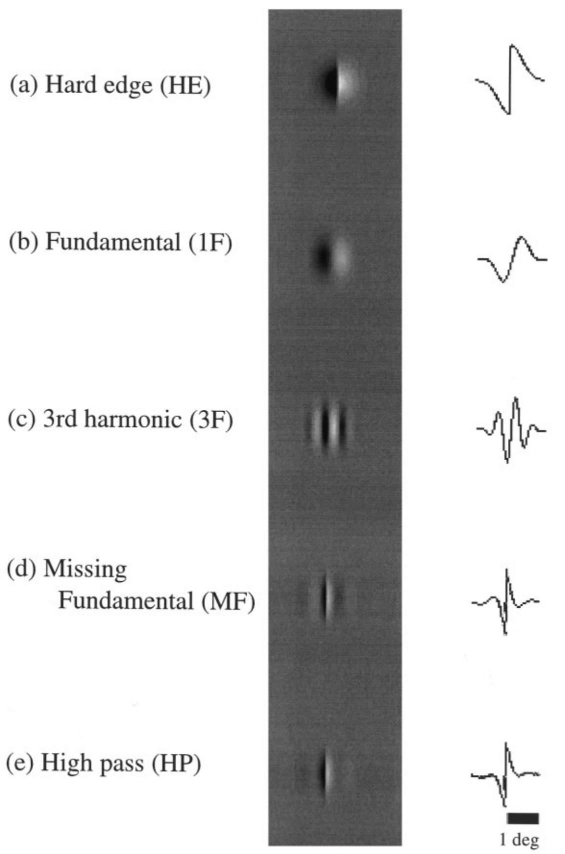

Fig. 4. The micropatterns and their luminance profiles as used in Experiment 2: (a) HE, a single luminance 'hard edgel'; (b) 1F, its fundamental component; (c) 3F, its second component (3rd harmonic); (d) MF, an edge with a missing fundamental; and (e) HP, high pass filtered edgel (see text). The size of each micropattern was the same (σ 0.36°). These images and those in Fig. 6 were created via the subroutines used in the experiments, and their profiles drawn using a graphics routine from their pixel values.

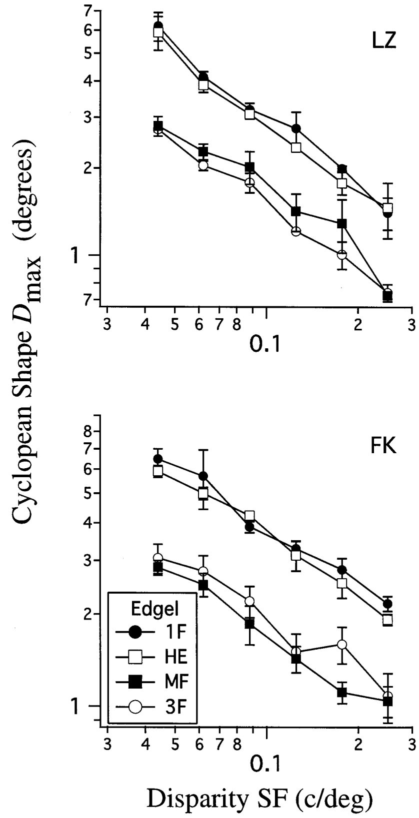

Fig. 5. The results of Experiment 2. D_{max} was smaller for the 3F and missing fundamental (MF) than for the HE and 1F micropatterns.

4.2. Results

The results, shown in Fig. 5, were consistent between observers and across corrugation frequencies. D_{max} was about the same for the HE and 1F edgels. The additional odd harmonics (3F, 5F, etc.) within the HE appear to have had little effect on d_{max}. This result implies that the number of features, not bandwidth, is the critical factor determining d_{max}. However this conclusion is not supported by the results with the MF and 3F edgels. Even though they contained different numbers of features, their d_{max} was the same, about 1/3 of that of the HE and 1F edgels. To be certain that this was not because the contrast of the 5F component was below threshold, we repeated measurements for one observer (LZ) at 100% contrast for the MF edgel, and found similar results.

To confirm that the difference in d_{max} between the 1F and MF edgels was not due to the relatively small number of cycles in the MF edgel, we repeated our comparison with larger ($\sigma = 0.72°$) micropatterns. Half as many (600) were used, and contrast was 88% but otherwise conditions were identical. One observer (LZ) was tested with three staircases per condition. At a modulation frequency of 0.125 c/deg, d_{max} for the MF edgel was 0.59° (s.e. 0.04), again much lower than the value for the 1F edgel of 1.6° (s.e. 0.05). It was closer to

d_{max} for the 3F edgel, 0.44° (s.e. 0.01). We tried a number of combinations of stimulus conditions, including 100% contrast and longer (200 ms) exposure times, yet found d_{max} for the MF was always much lower than d_{max} for the 1F alone.

In conclusion, the results of this experiment are consistent with neither of the initial two hypotheses. D_{max} was not determined by micropattern bandwidth alone, nor by the number of local features alone, though, of the two, the number-of-features hypothesis appears better supported. The results however led us to propose a new hypothesis, that d_{max} is determined by the number of cycles of the lowest frequency component of a micropattern. That would predict all of our findings, specifically that d_{max} is the same for the 1F and HE, that d_{max} is the same for the MF and 3F, and that the latter set of values are smaller than the former. The next experiment was aimed at directly testing between this hypothesis and the number-of-features hypothesis.

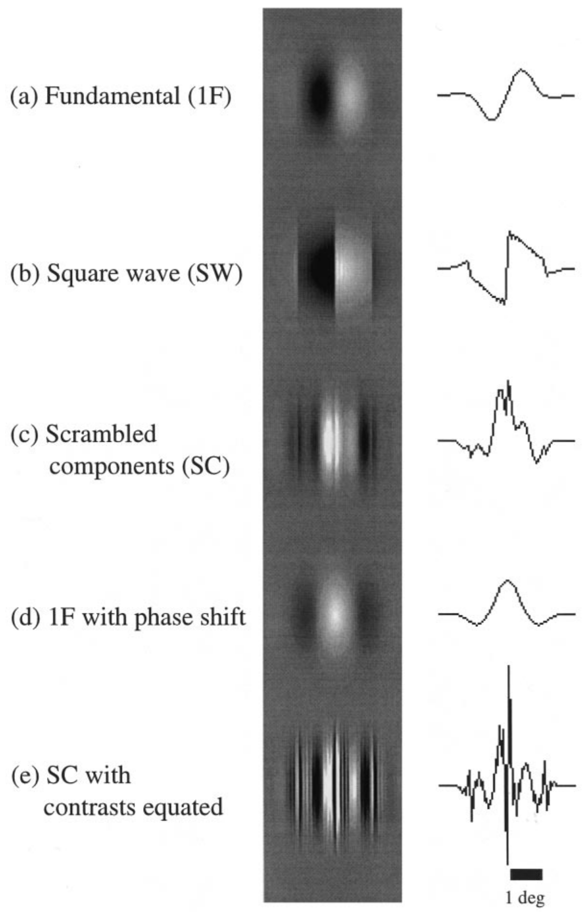

Fig. 6. Micropatterns used in Experiment 3: (a) 1F or fundamental; (b) SW, square wave with the same period; (c) SC, the same square wave with phase-scrambled components; (d) control condition, the fundamental shifted 90°; and (e) the same as SC except the components have equal contrasts. All envelope sizes were the same ($\sigma = 0.72°$).

5. Experiment 3: varying component phases

This experiment was designed to test whether d_{max} is determined by (1) the number of local features, or (2) the number of cycles at the lowest spatial frequency. To do this, we varied the phase relationships among the Fourier components of an edgel consisting of approximately one cycle of a square wave. The components are the odd harmonics, whose amplitudes decrease with frequency, and whose relative phases are all the same (zero). Altering the phases can vary the number of local features without changing the number of cycles of any of the Fourier components. There are two possible outcomes of this procedure. If the number of features alone determines d_{max}, then changing the relative phases and thereby introducing more features should result in a smaller d_{max}. On the other hand, if the number of cycles of the lowest frequency component alone determines d_{max}, then altering their phases should have no effect.

5.1. Stimuli

For this experiment we used micropatterns in which the odd harmonic components of a square wave were either in-phase (SW, see Fig. 6b), or assigned relative phases that resulted in multiple edges ('scrambled components', SC, Fig. 6c). For the SC micropatterns, the phases were $+90°$ for the fundamental, $-180°$ for the 3rd harmonic, and 0° for the other harmonics (0° is sine phase with origin at patch center). As a control for the effect of shifting the 1F component in the SC micropattern, we also measured d_{max} for the fundamental alone, shifted $+90°$ (Fig. 6d). All micropatterns were the same size ($\sigma = 0.72°$) and we used one disparity spatial frequency (0.125 c/deg). The contrast of the 1F component, from which contrasts of the other components were determined, was 33%.

5.2. Results

The fundamental (1F) and the square wave (SW), as well as the shifted 1F control edgel, resulted in the same d_{max} as shown in Fig. 7. D_{max} for the SC micropattern, in which the square wave component phases were scrambled, was also the same. These results indicate clearly that d_{max} is not determined by the number of visible local features. The results are consistent with the suggestion that d_{max} is determined by the number of cycles of the lowest frequency component.

What, however, if we increased the contrasts of the higher harmonics, which had a $1/f$, or fractal relation, in our edgels? If our hypothesis were strictly true, this should not affect d_{max}. For this manipulation we used the same scrambled phases as in SC above, except the components had equal amplitudes (Fig. 6e). This new

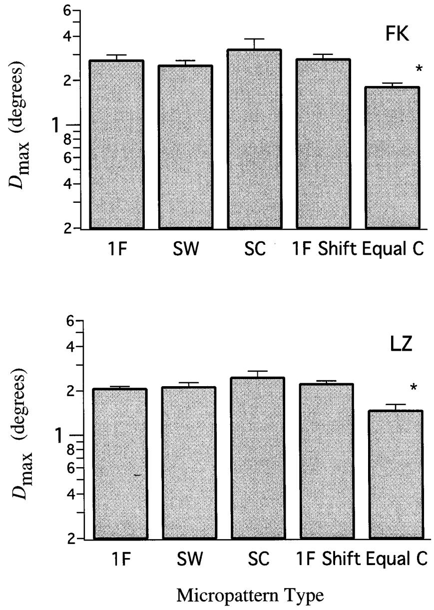

Fig. 7. The results of Experiment 3. D_{max} was the same for the fundamental (1F), the square wave (SW), and the square wave with phase scrambled components (SC), as well as for the control (phase shifted 1F). D_{max} was reduced however when the contrasts of all of the components of the SC micropattern were made equal.

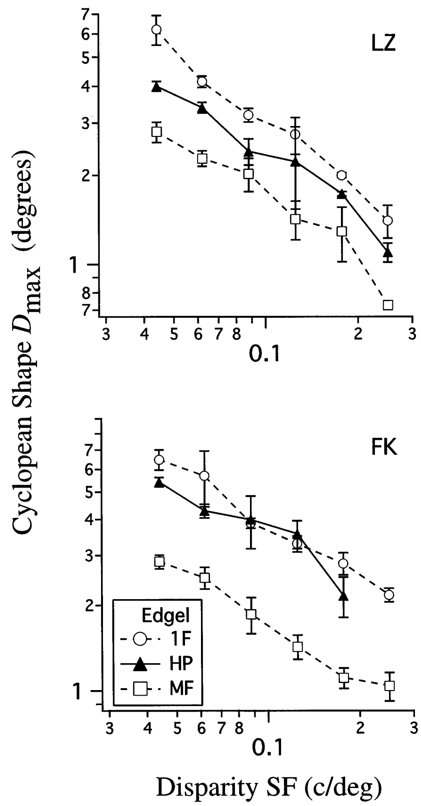

Fig. 8. Adding only 20% of the fundamental (1F) component to the missing fundamental (MF) edgel raised d_{max} toward the 1F levels. Data for the 1F and MF micropatterns are included for comparison.

micropattern, with the contrasts of the high frequency components of the SW edgel increased, resulted in a smaller d_{max} ('Equal C', Fig. 7). T-tests between every pair of conditions for both observers showed differences ($P < 0.05$) only between the equal contrasts condition and the others. This implies d_{max} is not always determined completely by the lowest luminance spatial frequency component, nor by the component that has the fewest number of cycles. Rather, a component's contribution may be reduced by another component of sufficient contrast.

To examine further how the lowest frequency component affects d_{max} when different spatial scales are combined, we measured d_{max} with a micropattern identical to the missing fundamental used in Experiment 2, except that we left 20% of the 1F component. In other words, we removed only 1.0 times the fundamental from the hard edgel (we had originally subtracted the entire 1F Fourier component; its amplitude is 1.273 times the amplitude of the edge). This new high-pass filtered edgel (HP, Fig. 4e) was only slightly different in appearance from the MF edgel (Fig. 4d). Yet this small amount of energy at the lower scale had a large effect on d_{max}, as shown in Fig. 8, raising it consistently to midway between the 1F and 3F levels for observer LZ. For observer FK it was raised midway, or in some cases completely, to the 1F level. This complements our previous finding, where raising the contrasts of the high frequency components lowered d_{max}.

In conclusion, our results are consistent with the hypothesis that, for any particular spatial scale when presented alone, d_{max} varies inversely with the number of local cycles. When a stimulus contains multiple spatial scales d_{max} depends upon a combination of these contributions at each scale and is determined to some extent by the components' relative contrasts.

6. Discussion

Our results show that d_{max} for seeing cyclopean shape is not determined by luminance spatial frequency. Nor is it determined by any other local factor alone, whether size, bandwidth, number of features, or edge content. We can best summarize our results by this general rule: D_{max} for cyclopean shape depends upon the contributions at different luminance spatial scales, and each of the individual contributions varies inversely with the number of cycles at that scale, with some weighting by component contrasts. This rule predicts that d_{max} should be independent of spatial frequency per se.

This rule explains the results of each experiment. In Experiment 1 there was only one spatial scale in each Gabor, and d_{max} was found to be inversely proportional to the number of cycles, as the rule predicts. The rule also works correctly for Experiment 2, where both the hard edgel and the fundamental had one cycle at the spatial scale with the highest contrast (the 1F scale) and gave the same d_{max}. The 3rd harmonic and the missing fundamental edgels which both contained three cycles at the scale of the 3rd harmonic, produced the same d_{max}, lower than that for the fundamental. In Experiment 3, all micropatterns had the same number of cycles at each scale, and all but one gave the same d_{max}.

A role for relative contrast was made clear when we raised the contrasts of the high frequency components of the scrambled phase stimulus in Experiment 3, and this resulted in a smaller d_{max}. The exact role however of relative contrast in determining d_{max}, is uncertain at this time. Some possibilities include masking by high frequencies as well as an averaging of component contributions.

6.1. Relations to previous studies

In Experiment 2 we found that d_{max} for disparity gratings made from missing fundamental (MF) micropatterns was not equal to that of the fundamental but to that of the 3rd harmonic. This finding for stereo shape is different from those of two other reports of experiments involving planar stimuli at a single disparity. Mayhew and Frisby (1981) provided a demonstration where missing fundamental gratings provided the opposite depth percept to 3rd harmonic gratings, suggesting luminance features (the edges) were used. A similar conclusion was reached by Boothroyd and Blake (1984) who asked observers to adjust the disparity of a 1F luminance grating until its perceived depth matched that of an adjacent missing fundamental (MF) grating. They concluded that correspondence could be made to the composite information or edges in the missing fundamental. In contrast, here we found no evidence that features, such as edges, have a special role in the perception of stereo shape (see also Ziegler & Hess, 1999). Furthermore, Schor et al. (1984) measured d_{max} using isolated micropatterns (DOGs, similar to Gabors) having envelope sizes that varied inversely with center frequency, so bandwidth was also held constant (1.75 octaves). They found that, over the range of luminance spatial frequencies used here, d_{max} was constant at high frequencies and varied inversely with frequency at low frequencies. In Experiment 1 however we found that, when Gabor bandwidth was constant, d_{max} for cyclopean shape was constant over a range of luminance spatial frequencies (Fig. 3). Wilcox and Hess (1995) used single micropatterns in a local stereo task and concluded that d_{max} *increases* with increasing envelope size, the opposite of our results (Fig. 1b and Fig. 3). Our results provide further evidence that, for establishing correspondence, contrast envelopes do not play a general role; it has been demonstrated that only first order (linear) mechanisms are involved in stereo shape perception (Ziegler & Hess, 1999). D_{max} for seeing stereo shape appears governed by a different set of rules than that for seeing the depth of a single micropattern. In the later case, near/far tasks, typically used, may be more sensitive (Ziegler & Hess, 1999) to particular components or features. For example, sensitivity to low frequencies may explain some depth seen during near/far tasks, even at disparities where a Gabor's main components result in diplopia (Ziegler & Hess, 1997). Even here, d_{max} was sensitive to low frequencies that did not produce correspondence noise (Fig. 8).

6.2. Why the number of cycles?

Why is the number of micropattern cycles so important for cyclopean shape d_{max}? Binocular correspondence is generally modeled by a mathematical cross-correlation function that operates locally between the two eyes' views. Our results imply that the visual system is generally unable to find the best overall match between two Gabors according to such a model. Rather, they show (Fig. 3) that d_{max} decreases with an increase in the number of cycles per Gabor. It seems unlikely therefore that the early luminance filters supporting stereo shape perception are narrowband, with each filter sensitive to many local luminance cycles.

On the contrary, our finding that d_{max} was reduced as the number of cycles within a Gabor increased could be better explained if our narrowband Gabors had stimulated many adjacent, or perhaps overlapping, broadband filters. This would have resulted in incorrect or ambiguous assignment of correspondence between similar portions of the half-images (*false targetting*). The resulting mixtures of correct and incorrect depth signals (*correspondence noise*) could have reduced d_{max}. The cause appears to have been local, between the intended Gabor half-images, since d_{max} was clearly not reduced by an increase in Gabor density (Fig. 1a; nor with up to 9600 Gabors per screen in Ziegler et al., 2000). That is, dense Gabor half-images at random did not alone reduce d_{max}. We conclude that any model of binocular correspondence should explain the results of Experiment 1, specifically, why correspondence could not be established between the overall positions of the narrowband Gabors.

6.3. Why does luminance spatial frequency NOT determine d_{max}?

Our finding that luminance spatial frequency, with bandwidth constant, had no effect on d_{max} (Fig. 3) may appear surprising, because stereopsis has been considered to exhibit 'size-disparity correlation' (Marr & Poggio, 1979; Smallman & MacLeod, 1994). By that hypothesis, stereopsis, at least veridically, fails when disparities exceed a half-cycle limit, i.e. half of the period of the lowest luminance spatial frequency to which a filter responds. Is it possible to reconcile our results and the notion of a size-disparity correlation?

Consider the implications of disparity gradients, rather than single element disparities. Whenever we used a range of disparity spatial frequencies (Figs. 1, 5 and 8), d_{max} consistently decreased with increasing disparity spatial frequency. This effect has been explained by a disparity gradient limit (Tyler, 1974; Burt & Julesz, 1980). A disparity gradient however can be computed in different ways depending on the filter size. It can be computed from two pairs of adjacent high frequency filters, with each pair processing small disparities, or by two, more widely separated, pairs of low frequency filters responding to proportionally larger disparities. Thus combining a disparity gradient limit with size-disparity correlation might explain why we found d_{max} for stereo shape was unaffected by changes in luminance spatial frequency.

On the other hand, there are a number of reasons to question this as a complete explanation. First, with our stimuli we always experienced the shape increasing in depth amplitude until a disparity was reached where the shape abruptly disappeared. A size-disparity correlation would predict that disparities beyond the half-cycle limit would result in shape distortion, due to opposite directions of perceived depth than intended (aliasing), which we never saw. Second, size-disparity correlation itself has been questioned, as depth was reported at disparities well beyond the half-cycle limit (Pulliam, 1982) even using single elements (Schor et al., 1983). Third, the disparity gradient explanation for our results may require filters that are 'tiled' in the visual field, with lower frequency filters spaced farther apart. The gradient explanation, however, seems plausible, given the known physiology, and may not necessarily be inconsistent with extant models of early cortical stages (Ohzawa, DeAngelis & Freeman, 1996).

Our results here, as well as how global factors determine d_{max} (Tyler, 1974; Ziegler et al., 2000), might be explained by a model with at least two stages (Hess et al., 1999). A bank of luminance filters, tuned to a number of spatial frequencies, would serve to *register* local luminance disparity. The broadband tuning of those filters might be one of several reasons why, when their outputs combine at a higher stage to provide for the veridical perception of stereo shape, d_{max} was found insensitive to changes in luminance spatial frequency.

Acknowledgements

These results were presented at the annual meeting of the Association for Research in Vision and Ophthalmology, Fort Lauderdale, FL, in May, 1999. The work was supported by grants to RH and FK from the Natural Sciences and Engineering Research Council (NSERC) and the Medical Research Council (MRC) of Canada.

References

Boothroyd, K., & Blake, R. (1984). Stereopsis from disparity of complex grating patterns. *Vision Research*, *24*(10), 1205–1222.

Burt, P., & Julesz, B. (1980). Modifications of the classical notion of Panum's fusional area. *Perception*, *9*(6), 671–682.

Glennerster, A. (1998). d_{max} for stereopsis and motion in random dot displays. *Vision Research*, *38*(6), 925–935.

Hess, R. F., & Wilcox, L. M. (1994). Linear and non-linear filtering in stereopsis. *Vision Research*, *34*(18), 2431–2438.

Hess, R. F., Kingdom, F. A. A., & Ziegler, L. R (1999). On the relationship between the spatial channels for luminance and disparity processing. *Vision Research*, *39*(3), 559–568.

Marr, D., & Poggio, T. (1979). A computational theory of human stereo vision. *Proceedings of the Royal Society of London Series B*, *204*, 301–328.

Mayhew, J. E. W., & Frisby, J. P. (1981). Psychophysical and computational studies towards a theory of human stereopsis. *Artificial Intelligence*, *17*, 349–385.

Ohzawa, I., DeAngelis, G. C., & Freeman, R. D. (1996). Encoding of binocular disparity by simple cells in the cat's visual cortex. *Journal of Neurophysiology*, *75*(5), 1779–1805.

Pulliam, K. (1982). Spatial frequency analysis of three-dimensional vision. In K. L. S. Ketty, *Visual simulation and image realism II*, (pp. 71–77).Proceedings No. 303, SPIE.

Rashbass, C., & Westheimer, G. (1961). Disjunctive eye movements. *Journal of Physiology* (*London*), *159*, 339–360.

Schor, C. M., Wood, I., & Ogawa, J. (1984). Binocular sensory fusion is limited by spatial resolution. *Vision Research*, *24*(7), 661–665.

Smallman, H. S., & MacLeod, D. I. A. (1994). Size-disparity correlation in stereopsis at contrast threshold. *Journal of the Optic Society of America A*, *11*(8), 2169–2183.

Stevenson, S. B., Cormack, L. K., & Schor, C. M. (1994). The effect of stimulus contrast and interocular correlation on disparity vergence. *Vision Research*, *34*, 383–396.

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

Wilcox, L. M., & Hess, R. F. (1995). D_{max} for stereopsis depends on size, not spatial frequency content. *Vision Research*, *35*(8), 1061–1069.

Ziegler, L. R., & Hess, R. F. (1997). Depth perception during diplopia is direct. *Perception*, *26*, 1225–1230.

Ziegler, L. R., & Hess, R. F. (1999). Stereoscopic depth but not shape from second-order stimuli. *Vision Research*, *39*(8), 1491–1507.

Ziegler, L. R., Hess, R. F., & Kingdom, F. A. A. (2000). Global factors that determine the maximum disparity for seeing cyclopean surface shape. *Vision Research*, *40*(5), 493–502.