PERGAMON

Vision Research 39 (1999) 559–568

Vision Research

On the relationship between the spatial channels for luminance and disparity processing

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

Department of Ophthalmology, McGill Vision Research Unit, McGill University, 687 Pine Avenue West, Room H4-14, Montréal, Québec, Canada

Received 12 May 1997; received in revised form 15 December 1997

Abstract

To determine the relationship between the spatial channels for luminance and shape-from-stereo-disparity processing we measured disparity modulation sensitivity as a function of disparity spatial frequency for sinusoidal modulations of a field of Gabor micropatterns of differing luminance spatial frequency. We first examine the effects of contrast, spatial bandwidth and element density and show that it is only the last of these which is critical for the shape of the disparity modulation threshold function. We show that the shape of this function depends on the luminance spatial frequency of the surface that is modulated in depth. Specifically, low corrugation frequencies enjoy a greater scale support from the early luminance spatial filters than do high corrugation frequencies. The results are consistent with higher spatial frequency disparity channels receiving a greater input from higher spatial frequency luminance channels.

Keywords: Disparity modulation threshold function; Luminance spatial frequency; Gabor

1. Introduction

A great deal of research over the past few decades has been directed towards an understanding of the properties and organization of the early filters (Blakemore & Campbell, 1969; Graham & Nachmias, 1971; Stromeyer & Julesz, 1971; Graham, 1989). One of the key issues in vision involves how the outputs of the early spatial and temporal filters are put together to extract spatial contours, motion, colour and stereo-depth. In terms of stereo-depth one must consider not only local disparity of isolated objects but also the shape of whole textured surfaces. In this study we ask, how is the information from early luminance spatial filters used to derive stereo shape?

Shape-from-stereo processing typically has been investigated using 'disparity gratings' (Tyler, 1974) in the form of noise patterns sinusoidally modulated in disparity. Disparity gratings provide a compelling percept of a 'depth corrugation' and are a useful tool for probing the mechanisms responsible for shape-from-stereo processing. They provide a way of relating the threshold for detecting depth corrugations to their corrugation frequency, producing a function known as the disparity modulation threshold function. This is the analog for stereopsis of the well-known luminance contrast sensitivity function. The disparity modulation threshold function is generally U-shaped, implying that thresholds are lowest at some intermediate corrugation frequency. Disparity gratings have been used in masking and summation studies to provide evidence for channels sensitive to different spatial frequency ranges of disparity modulation (Tyler, 1983; Cobo-Lewis & Yeh, 1994; Schumer & Ganz, 1979).

Disparity gratings are used here to address an important issue: What is the nature of the 'first stage' inputs to the mechanisms involved in shape-from-stereo processing? While these inputs are presumably local disparity detectors, not all such detectors are alike. There is an established body of evidence from both neurophysiology (DeAngelis, Ohzawa & Freeman, 1995) and psychophysics (Blakemore & Hague, 1972; Felton, Richards & Smith, 1972; Julesz & Miller, 1975; Schor, Wood & Ogawa, 1984; Yang & Blake, 1991; Smallman & McLeod, 1994) that local stereopsis mechanisms are spatially tuned. This has given rise to the notion of the 'size-disparity correlation' (Schor & Wood, 1983; Schor

* Corresponding author. Fax: +1 514 843 1691; e-mail: rhess@bradman.vision.mcgill.ca.

0042-6989/98/$ - see front matter
PII: S0042-6989(98)00127-8

et al., 1984), in which it is held that fine disparities are detected by fine-scale detectors and coarse disparities by coarse-scale detectors. The finding that local stereopsis mechanisms are spatially tuned raises the possibility that the mechanisms responsible for integrating local disparity information to provide information about whole-surface-shape also possess some degree of luminance scale selectivity. This issue becomes particularly pertinent in the light of the evidence for disparity modulation channels given above. If a size-disparity correlation at the level of individual disparity detectors exists, then it follows that a similar 'size-disparity-frequency' correlation might also exist, a result of individual disparity modulation channels being selective for their size-tuned disparity detector inputs. One of the expected behavioural manifestations of such a size-disparity-frequency correlation would be a shift in the position of the disparity modulation function with the luminance scale of the stimulus.

Two previous studies have attempted to determine the relationship between the shape of the disparity modulation function and luminance scale. The first, by Pulliam (1981), measured thresholds for detecting disparity modulation for vertically oriented luminance sine-wave gratings which were sinusoidally modulated in disparity along their vertical axis. Pulliam found that the U-shaped disparity modulation function appeared to shift rightwards with an increase in the luminance spatial frequency of the grating. This supports the notion of a size-disparity-frequency correlation. However, one cannot ignore the possibility that monocular vernier cues, which Pulliam showed to have lower thresholds than those for disparity modulation, were used to help solve the stereo task. A more recent study by Lee and Rogers (1997) used two-dimensional filtered noise stereograms, in which such monocular cues were effectively camouflaged. Their results led them to conclude, contrary to Pulliam, that "… luminance spatial frequency and disparity corrugation frequency are largely independent dimensions, with the exception of a slight interaction at high luminance frequencies and low corrugation frequencies". However, the disparity modulations in Lee and Roger's stimuli were introduced only after they were filtered to render them narrow-band-in-luminance. The necessary shearing of the filtered stimulus required to produce the two stereo-half-images introduced high luminance spatial frequencies into the stimulus, and one cannot rule out the possibility that these were detectable. If so, their presence may have obscured any dependence of thresholds on luminance scale at high disparity modulation frequencies.

That such a dependence would exist seems likely on a priori grounds. If local disparity detectors are size-tuned, this must inevitably place certain limitations on the surface-shapes that can be represented. A high spatial frequency disparity grating constructed from low luminance spatial frequencies, assuming that the disparity corrugations were themselves physically realizable, would be poorly represented in the output of an array of coarse-scale disparity detectors. The grating's fine-scale changes in disparity, which result from differential fine-scale gradations in luminance between the eyes, would presumably be smoothed out by such filters.

To summarise, there appear to be four possible linkages between luminance scale and disparity modulation processing, and these are illustrated in Fig. 1. The linkage may be indiscriminate (A), selective at only low disparity frequencies (B), selective at only high disparity frequencies (C) or selective at high and low disparity frequencies. The results of Pulliam (1981) support the linkage in D, whereas those of Lee and Rogers (1997) support the linkage in B.

We have re-examined the linkage between luminance spatial scale and the detection of disparity modulation. To avoid the problems inherent in generating narrow-band-in-luminance disparity gratings by filtering noise, we used a different approach: modulating the disparity of a field of randomly positioned Gabor micropatterns. This allows one to assess how the shape of the overall disparity modulation sensitivity function varies for micropatterns having different luminance spatial frequencies. This not only allows one to easily assess the role of luminance spatial frequency bandwidth but also to disassociate textural density from luminance spatial frequency. This latter factor has been previously shown to be of some importance in motion processing (Boulton & Baker, 1991; Eagle & Rogers, 1996) and is necessarily confounded in previous attempts to gauge the linkage between luminance spatial frequency and disparity spatial frequency (Pulliam, 1981; Lee & Rogers, 1997).

2. Methods

2.1. Apparatus

A graphic workstation (Silicon Graphics, Indigo^2) was programmed for all stimulus-generation, experiment control and data collection. A video projector (Electrohome, ECP-4100), capable of a 120 Hz vertical sweep frequency, back-projected to a large paper-plexiglas screen (Crist & Robinson, 1989). The raster dimensions (height $\times$ width) were 112×124 cm, at a resolution of 1024×1280 pixels respectively, so each pixel was about 1 mm wide. Stereo views were presented with the computer in stereo mode, using liquid-crystal (LCD) cross-polarizing shutter-glasses (StereoGraphics, CrystalEyes model CE-PC with E-1 emitter). The half-images were presented at 120 Hz,

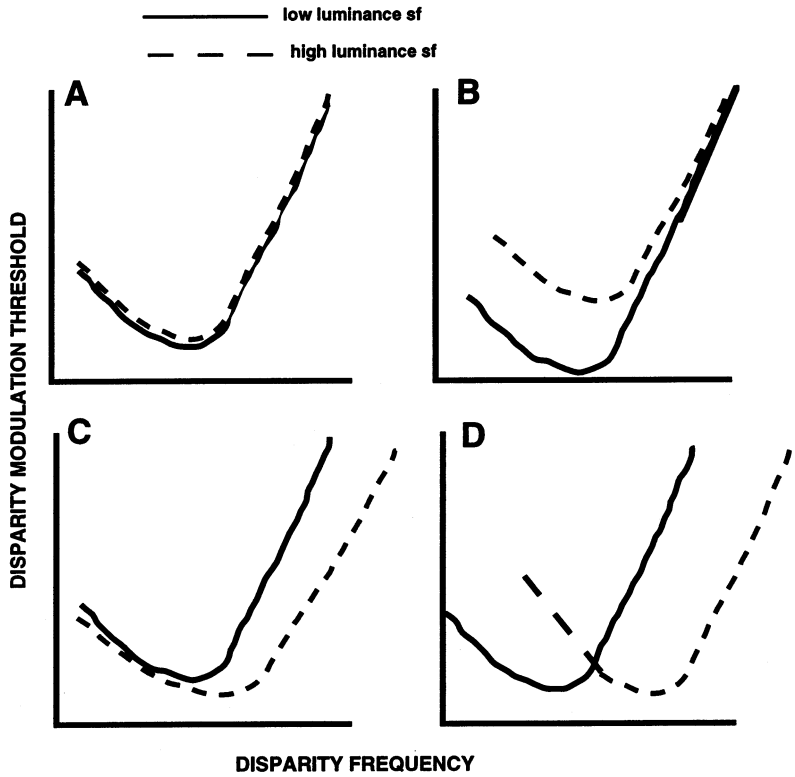

Fig. 1. Four possible types of linkages between low level luminance spatial channels and higher level disparity spatial channels. Previous research has suggested the linkage in B and D.

alternating in the odd and even raster scan lines. These alternations were synchronized with the lens transparencies so an eye saw its image at an effective frame rate of 60 Hz. The LCD glasses were easy to use and provided excellent image registration over the entire field, compared to alternatives such as mirror stereoscopes. A disadvantage is some cross-talk. With the low contrast and grey background used, any effects of cross-talk however would be negligible with respect to the factors manipulated here.

Before any experiments began, we tested the luminance linearity of the display using standard photometric techniques and the computer's built-in hardware gamma-correction to linearize it.

2.2. Stimuli

Each stimulus consisted of a large number of identical Gabor patterns (Graham, 1989) plotted at random on the screen. These were copies of a square template in main memory that had been generated according to:

$$L(x,y) = L_{\text{mean}} + L_{\text{amp}} \cdot \sin(2\pi f x) \cdot \exp\{-(x^2+y^2)/2\sigma^2\}$$

where $L(x, y)$ is thepixel grey level (at location relative to center), L_{mean} is the background grey level, L_{amp} is the carrier luminance amplitude, f is the carrier luminance spatial frequency (cpd), σ is the scale factor of Gaussian envelope (standard deviation).

In all experiments but one, in which we used frequencies one octave below and above, the Gabors had one of two carrier frequencies: 0.42 and 1.68 cpd. Two sizes of Gabors were used, 16×16 and 64×64 pixels. The parameters used for each type of Gabor are in Table 1. Fig. 2 illustrates our stimulus (an obliquely oriented corrugation) for Gabors having either the same octave bandwidth but different carrier spatial frequencies (top versus middle) or the same size Gaussian envelope (hence the same coverage) but different carrier spatial frequencies (top versus bottom). In all three cases the Gabor locations are identical.

Before we began the experiments we found that with equal physical contrasts the large Gabors appeared perceptually more salient than the small ones, i.e. they had higher apparent contrast. Contrasts could either be set for equal physical contrast or for equal perceptual contrast. We decided to set contrast for equal perceptual contrast. The values used are based on a separate experiment to find the point of equal perceived contrast. To establish however that the different physical contrasts were not the cause of the differences in our results we ran a separate control experiment, discussed below.

We required a method of displaying a large number of Gabors on the screen while avoiding as much as possible the introduction of luminance artifacts. If Gabor patches were simply overlaid at random onto

Table 1
Parameters describing the three Gabor patches used in most of the experiments

Gabor type (size/freq)	Patch size (pixels)	Sigma (°)	Carrier frequency (cpd)	Michelson contrast
Large/low	64×64	0.72	0.42	0.14
Small/high	16×16	0.18	1.68	0.33
Large/high	64×64	0.72	1.68	0.25

Two others in one experiment were modifications of these as described in the text.
See Fig. 2 for illustration.

the screen, edge artifacts would appear since the edges of the patches would cover bright or dark parts of previously plotted patches. Conventional techniques to avoid overlap entirely, such as placing each patch randomly within the cell of a grid pattern, themselves introduce artifacts. To avoid this problem our algorithm used Gabors that had no 'DC' component before they were added to a screen buffer. Only after all the patches were accumulated was the background luminance added.

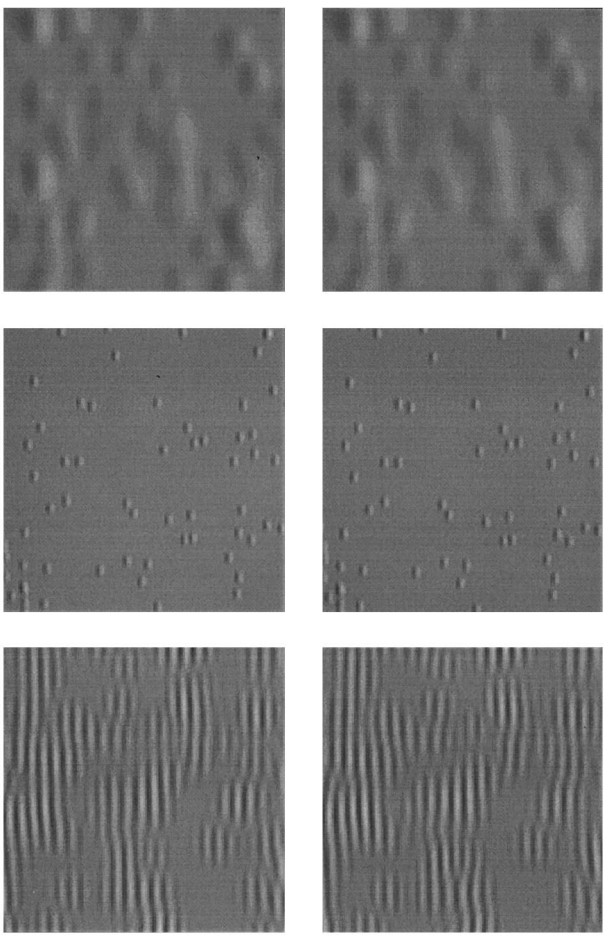

Fig. 2. Illustration of a small part of our stimulus field (an obliquely oriented corrugation) for Gabors having either the same bandwidth but different peak luminance spatial frequencies (top versus middle) or the same size Gaussian envelope (hence the same coverage) but different peak spatial frequencies (top versus bottom). In all three cases the Gabor density (number/picture) is constant.

Disparity was defined with subpixel resolution as follows. At the start of each experiment a large number of 'template' Gabors were generated in computer memory. Each template was a square patch of pixels within which a Gabor was positioned. The number of templates was equal to the number of subpixel disparity offsets. With disparity expressed in pixel units, the integral portion of that value provided the offset for a template's position on the screen, while the fractional portion provided the index to the set of templates.

Each of these disparities was calculated for a Gabor's random location on the screen, to produce the corrugation pattern, according to:

$$\delta(y') = \delta_{\max} \cdot \sin(2\pi F y' + \varphi)$$

where $\delta_{\max}$ is the disparity amplitude (adjusted during procedure), F is the modulation spatial frequency, y' is the vertical location on the screen, φ is 0–2π at random for each trial.

For our experiments with an oblique orientation of modulation, y' was a rotation of the actual location of a Gabor (x, y), according to:

$$y' = 0.707 \cdot (y \pm x)$$

where the $+$ or $-$ introduces the right or left orientation. Examples of the stimuli are provided in Fig. 2.

2.3. Observers

The three authors acted as subjects. The visual acuity of each subject was normal or corrected-to-normal, with prescribed eye-glasses worn during the experiments.

2.4. Viewing conditions

Except for differences in psychophysical methods and stimuli described here, conditions were essentially identical. Observers viewed the screen freely at a distance of 114 cm, where the field size was $52 \times 57°$.

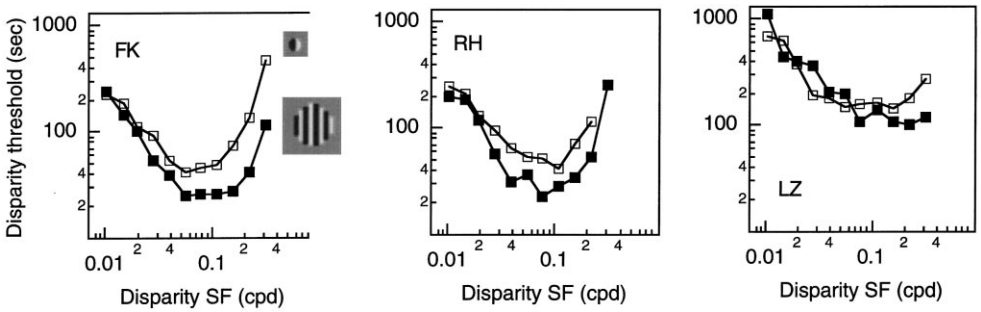

Fig. 3. Effect of bandwidth with 2AFC task. Disparity modulation threshold is plotted against disparity corrugation spatial frequency for three subjects. Both Gabors had luminance $sf = 1.68$ cd. Gabors had σ of 0.18 and 0.72° (see icons). The density was 1200 gabor patches per screen in both conditions. Graphs show that bandwidth has little effect, provided all else is constant.

2.5. Method of adjustment procedure

Our first two experiments used a horizontally oriented modulation of the disparity, with disparity amplitudes adjusted manually. In the first experiment, phase was randomized. In the second, phase was held constant at zero degrees (sine phase), because the disparity modulation was windowed by a Gaussian envelope that would otherwise give the maximum disparity on the screen a random component. The windowing was set with a sigma equal to 0.5 cycles of the modulation and was only used in the second experiment.

Each adjustment sequence began with a disparity spatial frequency selected randomly from a list. The initial disparity amplitude (δ_{max}) was sufficiently high that it was easy to perceive the 3-D surface corrugation. Observers pressed the right button to raise the disparity amplitude and the left button to lower it (each by 25%). Each stimulus contained a fresh randomization of Gabor locations as well as a random phase of disparity modulation to prevent observers judging depth based upon a single location on the screen, as may have happened in other studies (Lee & Rogers, 1997). Each stimulus was visible for 1–3 s and after a button-press it was replaced after about 1 s by a new stimulus, during which time the screen was at the mean luminance. At a disparity amplitude just sufficient for the modulation to be perceived the middle button was pressed to record the threshold estimate. Then the next adjustment began. To allow breaks, sequences were divided into blocks including one adjustment for each of the disparity spatial frequencies on the list. Eight adjustments were made for each modulation spatial frequency. Thresholds are reported as geometric means.

2.6. 2AFC Procedure

In all other experiments (except those whose data are displayed in Fig. 6) observers judged the orientation of the disparity modulation pattern as either slanted left or right 45° from the vertical (see Fig. 2 for illustration) by pressing the respective mouse buttons. Orientation was presented at random. A standard 'two-down, one-up' staircase procedure set the disparity amplitude for each trial whereby one mistake increased it, and it was lowered by two consecutive correct responses. At the start of each run the disparity was high to insure that the modulation was perceived. Changes were by a factor of 25%, although to reach threshold quickly this was doubled prior to the first reversal in a run. The procedure terminated after 12 reversals and the threshold was estimated as the geometric mean of the disparity amplitude at the last eight reversals. We report the average of these estimates from at least three staircase runs for each disparity spatial frequency.

2.7. Monocular control

For our stimuli with oblique corrugations, the vertical disparity component introduced a small monocular density cue. That is, disparities shifted the Gabors so there were fewer at the peaks and troughs of the modulation. We measured monocular performance with the same stimuli and found that, in the worse case (highest density of the small Gabors) thresholds were a factor of 3–4 higher than the highest disparity thresholds measured and two orders of magnitude worse than the lowest disparity thresholds measured. Thus monocular cues were unlikely to have contributed to the disparity thresholds reported here.

3. Results

Before one can address the main issue raised in this paper, namely the linkage between luminance spatial scale and disparity modulation channels, it is necessary to decide on the appropriate parameters for such a comparison. For example, since the bandwidth of Gabor micropatterns (the full width at half height of the Fourier amplitude spectrum) varies inversely with both their frequency and their size (Graham, 1989), as

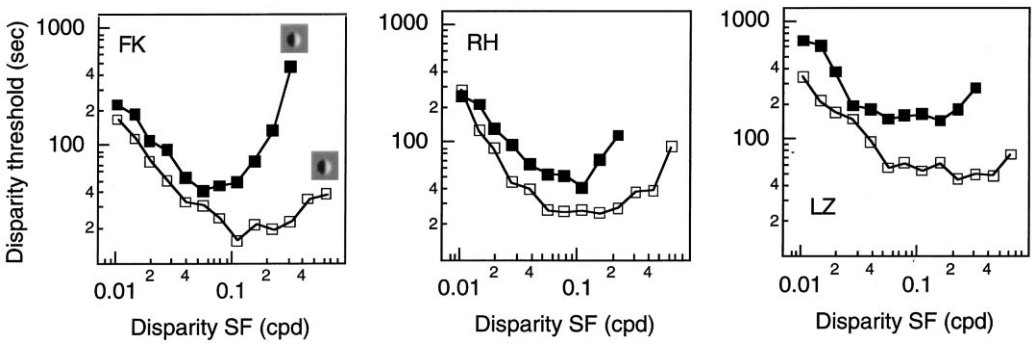

Fig. 4. Effect of density with 2AFC task. Disparity modulation threshold is plotted against disparity corrugation spatial frequency for three subjects. Gabor patches used had luminance $sf = 1.68$ cd and $\sigma = 0.18°$. Filled squares = 1200 gabor patches per screen. Open squares = 19200 Gabor patches per screen.

their frequency is varied, what should be held constant: bandwidth or size? Should different luminance micropatterns have the same density (micropatterns per screen) or coverage[1]? Finally should different luminance micropatterns be displayed at the same physical contrast or the same perceived contrast? To resolve these questions we compared disparity modulation threshold functions for fields of Gabor micropatterns having different bandwidths, densities and contrasts. Fig. 3 shows the effect of bandwidth/coverage of Gabor micropatterns for the 2AFC task. Here we compare disparity modulation thresholds for micropatterns with the same carrier frequency (1.68 cd) and density (1200 micropatterns/screen) but different sizes ($\sigma = 0.18$ and 0.72°). Although there is some variation in the overall shape of the disparity modulation threshold function, the quadrupling of micropattern bandwidth has very little effect on the overall shape. Coverage and bandwidth are therefore not key parameters for the results presented here.

One interesting aspect of our results is that by using a much larger stimulus field size than previous studies (nine times the area), we have been able to extend our measurements to much lower corrugation frequencies (a factor of ten). This is crucial for the assessment of whether performance is better for low spatial frequency corrugations constructed from low (as opposed to high) spatial frequency Gabors. As a result of our large field size, we find a peak in the disparity modulation function (around 0.08 cpd for the high spatial frequency Gabors) about a factor of four lower than previous studies (Tyler, 1974). The position of the peak in previous studies may have been the result of inadequate summation at low corrugation frequencies due to restricted field size.

Fig. 4 illustrates the effect of density. The Gabor micropatterns all had the same luminance spatial frequency (1.68 cpd) and size ($\sigma = 0.18°$). The filled squares are for a field of 1200 micropatterns/screen and the unfilled squares are for 19200 micropatterns/screen. The peak of these disparity modulation threshold functions is strongly dependent on micropattern density. Thresholds at higher frequencies of disparity modulation are lowered at the higher micropattern density.

The effect of the contrast of the micropatterns is shown in Fig. 5 for one subject, for two micropattern spatial frequencies a factor of four apart (namely, $sf = 1.68$ cpd, $\sigma = 0.18°$; $sf = 0.42$ cpd, $\sigma = 0.72°$.) The density was constant at 1200 micropatterns/screen. The slope of the contrast dependence is shallow (approximately square root) and independent of luminance spatial frequency. Therefore the absolute contrast of the luminance micropatterns does not have a strong influence on disparity threshold. Nor from these results would it seem that contrast is critical for the comparison of results between micropatterns of different luminance spatial frequency. Therefore in all subsequent experiments where results for different luminance spatial frequency micropatterns are compared we set the contrast of the micropatterns to perceptual equality.

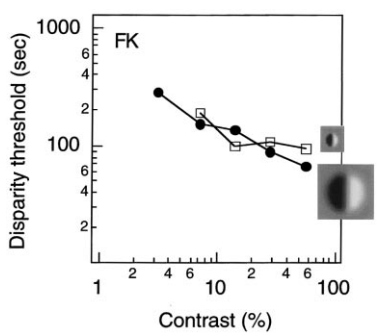

Fig. 5. Effect of contrast with 2AFC data for FK only. Disparity modulation threshold is plotted against stimulus contrast. Conditions were open squares: luminance $sf = 1.68$ cpd, $\sigma = 0.18$ °, 19200 per screen. Filled circles: luminance $sf = 0.42$ cpd, $\sigma = 0.72$ °, 1200 per screen (see icons). Disparity sf was 0.014 cpd. The fact that the slope of the contrast dependence is shallow (approximately square root) for both conditions suggests that our decision to perceptually match the contrast of gabors of different frequency was not crucial.

[1] The proportion of the screen area paved with micropatterns.

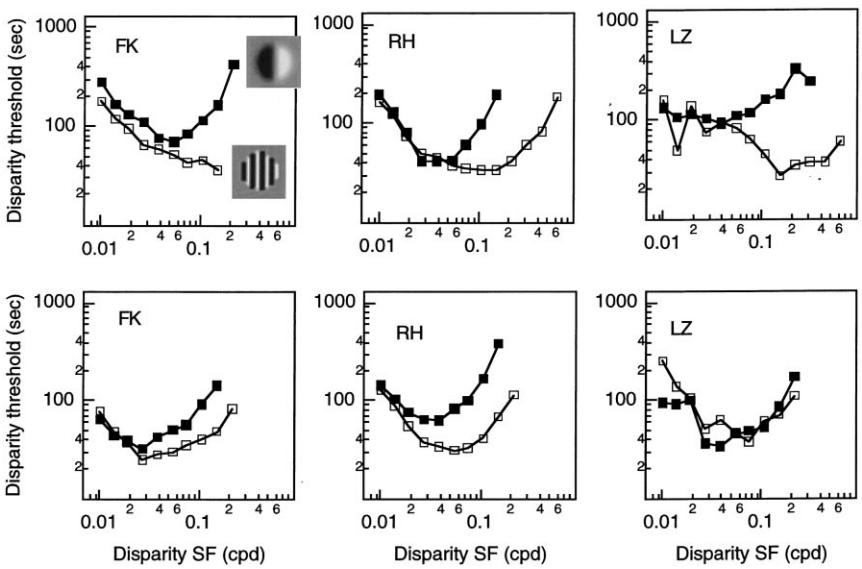

Fig. 6. Method of adjustment task. The stimulus had either a density of 600 Gabor patches per screen, and a disparity modulation that was non-Gaussian-enveloped (top row) or a density of 1200 and a Gaussian-enveloped disparity modulation with a constant number of cycles (bottom row). All Gabor patches have a σ of 0.72°). The Gabor patches have luminance spatial frequencies of 0.42 and 1.68 cpd (see icons).

In our subsequent comparisons of disparity modulation thresholds for fields of micropatterns we ensured that all luminance Gabors had the same size Gaussian envelopes (i.e. areas) and densities (number of micropatterns/screen). Our initial measurements used the method of adjustment both for unwindowed sinusoidal modulations in which the number of displayed cycles varied directly with the modulation frequency and for Gaussian windowed sinusoidal modulations of disparity in which the number of displayed cycles was fixed. Results for both of these stimulus conditions are shown in Fig. 6. The filled squares represent results for a field of Gabors of low spatial frequency ($sf = 0.42$ cpd; $\sigma = 0.72°$), the unfilled squares represent a field of Gabors of higher spatial frequency ($sf = 1.68$ cpd; $\sigma = 0.72°$). The results on the top row are for a density of 600 Gabors/screen with unwindowed disparity modulations and on the bottom row for 1200 Gabors/screen with Gaussian-windowed disparity modulations having a constant number of cycles ($\sigma = 0.5$ cycles). The shape of the disparity modulation threshold function clearly depends on the peak spatial frequency of the underlying Gabors. While threshold performance at low disparity frequencies is little affected by the change in luminance spatial frequency, threshold performance at high disparity spatial frequencies is improved for higher luminance frequencies. This results in a shift in the peak of the disparity threshold function to higher disparity frequencies for fields of Gabors of higher luminance spatial frequency. This shape change in the disparity modulation sensitivity function does not depend on whether the number of displayed cycles of disparity modulation are held constant or allowed to vary with disparity frequency.

Fig. 7 illustrates this same interaction between luminance spatial frequency and disparity spatial frequency for a wider range of Gabor luminance spatial frequencies (Gaussian blob, 0.42, 1.68, 3.36 cpd) using a 2AFC method in which the orientation of an obliquely oriented sinusoidal disparity modulation was detected. All disparity frequency modulations were displayed using a field of randomly positioned Gabors of the same density and Gaussian standard deviation. In the Gaussian blob condition both light and dark blobs were randomly mixed. Higher luminance spatial frequencies result in a shift in the peak of the disparity modulation threshold function to higher frequencies. A factor of eight change in the luminance spatial frequency results in approximately a factor of two shift in the peak of the disparity threshold function.

4. Discussion

The use of a modulated field of spatially bandpass micropatterns affords a number of advantages over the use of spatially filtered white noise stimuli. First it provides the opportunity of disassociating element density from peak spatial frequency. We show that element density rather than element coverage is the important factor. Second, the luminance spatial frequency spectrum does not change as a consequence of introducing a disparity into the stimulus. Third, equal stimulation of different luminance spatial frequency channels can be provided, so essential for any comparison of their relative contribution to stereo sensitivity.

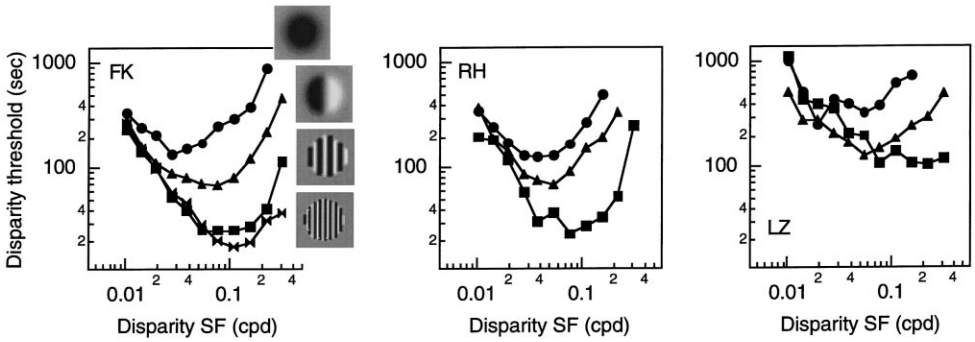

Fig. 7. 2AFC task. Disparity modulation threshold is plotted against disparity corrugation spatial frequency for three subjects. All stimuli had 1200 Gabor patches per screen, and all Gabor patches had a S.D. of 0.72° Gabor luminance *sf*'s were Gaussian blob, 0.42 and 1.68 (see icons). In the Gaussian blob condition both dark and bright blobs were randomly mixed. Only FK could do the $sf = 3.36$ cd condition.

There are two factors which could in principle affect the shape of the disparity modulation function for our micropattern-based stimuli, apart from the luminance spatial frequency of the micropatterns. The first is micropattern density, the second envelope size. Our stimuli consisted of finite numbers of randomly positioned micropatterns, and because of sampling limitations the high spatial frequencies of disparity modulation were necessarily less well represented than the lows. Moreover, for a given micropattern density, a change in envelope size will result in a different overall coverage of the display. Furthermore, our micropatterns, while randomly positioned and therefore potentially able to represent all disparities along the corrugated surface, were nevertheless themselves 'flat', i.e. did not follow the curvature of the surface itself. This could arguably introduce a degree of disparity quantization, dependent, like coverage, on the envelope size of the micropatterns. To investigate the effects of both these factors, we measured performance using micropatterns with the same luminance spatial frequency but with different micropattern densities and envelope sizes.

Our results demonstrate how a 16-fold increase in micropattern density significantly affects the shape of the disparity modulation function, by selectively reducing thresholds for the relatively high spatial frequencies. This is exactly what one would expect on the basis of sampling considerations, independent of whether one is dealing with Gabors or random dots. Furthermore, we demonstrate that micropatterns with a 4-fold difference in envelope space constant (a 16-fold increase in micropattern area and hence coverage) produce only a slight change in the overall shape of the disparity modulation function. This dependence is likely due to the greater contrast energy of the larger size micropatterns. Both sets of results argue for removing both density and envelope size as potential confounds when determining the effect of micropattern luminance spatial frequency on the shape of the disparity modulation function. Hence for the main experiment we decided to keep micropattern density and envelope size constant.

Our results show a small but clear dependence of the shape of the disparity modulation function on luminance spatial frequency. Higher luminance spatial frequencies result in enhanced threshold performance for higher disparity spatial frequencies. This finding runs contrary to the results of one previous study. We did not find any reduction in disparity sensitivity at low corrugation frequencies for stimuli composed of high luminance spatial frequencies (Lee & Rogers, 1997). While it is not clear what could account for this difference, our method does have the advantage of not introducing spurious luminance spatial frequency components as a function of increasing disparity. Furthermore, by using a global orientation judgment together with randomization of the spatial phase of the corrugation we ensured that performance was based on the global percept of a corrugated surface rather than the local disparity at any particular point. Finally, our range of low corrugation frequencies was much more extensive, extending to a factor of ten lower corrugation frequencies, because our screen size was much larger (nine times in area). The relationship that we find suggests that low corrugation frequencies have much wider scale support amongst the early spatial filters than do higher corrugation frequencies. In other words our threshold sensitivity at low corrugation frequencies is no better when these corrugations are constructed from low as opposed to high luminance spatial frequencies. This runs contrary to the conclusions of Pulliam (1981).

4.1. Physics of stimuli

The results that we find at high disparity frequencies is not an inevitable consequence of the design of the stimuli. There is no reason in terms of the physics of the stimuli why there should be any difference between the different conditions represented in Fig. 7. The

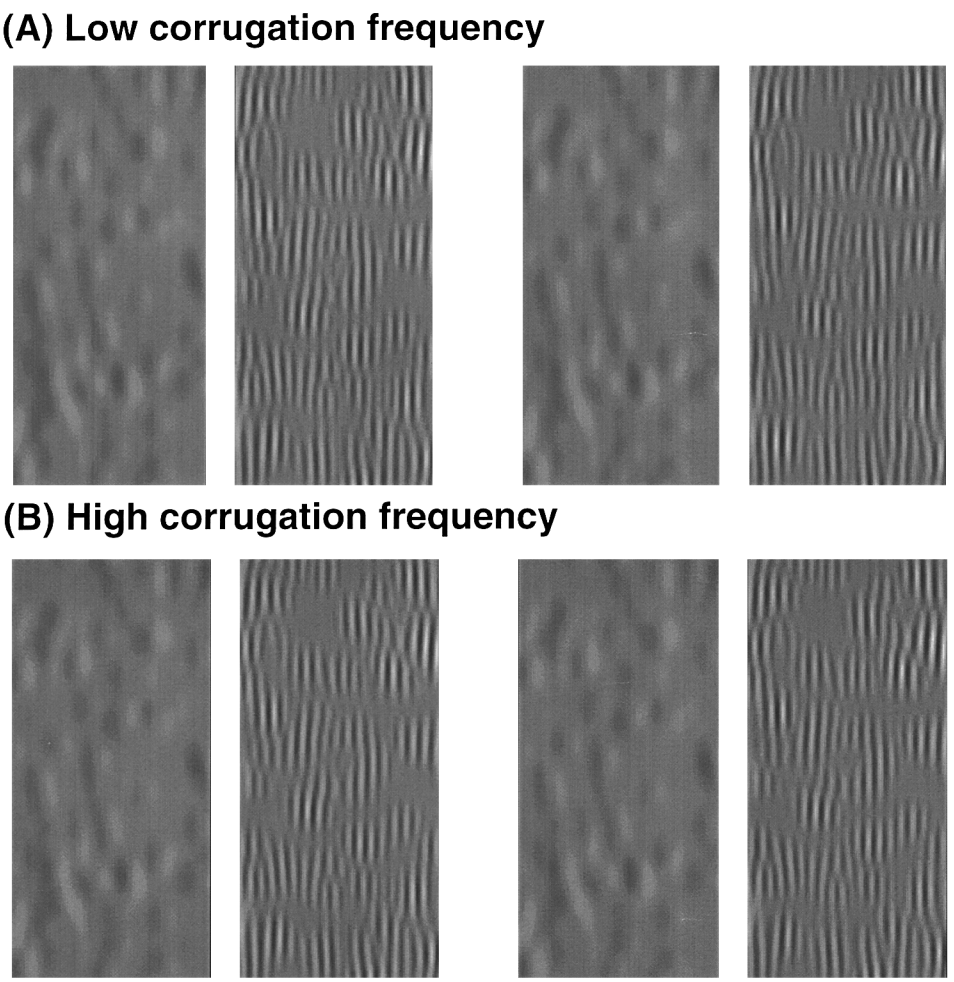

Fig. 8. Demonstration of our main finding. Stereograms are presented for a low (A) and high frequency (B) sinewave corrugation composed of high (left side) and low spatial frequency Gabor (right side). The Gabor size, Gabor perceptual contrast and Gabor densities have been kept constant. The disparity amplitude is 450 s arc for the low corrugaton frequency and 118 s arc for the high frequency corrugation. The low frequency corrugation (A) is easily seen for both fields of low and high Gabor, however the high frequency (B) corrugation is less apparent if at all with low frequency Gabor. Low frequency corrrugation enjoys a greater scale support from the early filters than do high frequency corrugation.

Gabor patches are flat, and regardless of their center frequency they cover the same area. Thus, in principle, they each supply one depth estimate, and since the sampling is constant across frequency for the Gabors, so too will be their representation of surface depth. Thus the differences in the curves in Fig. 7 must be a consequence of the visual system. Having said that, there are however a number of somewhat trivial physiological explanations for why sensitivity at high corrugation frequencies might be worse when represented by low spatial frequency Gabors. These include spatial summation, sampling, and signal-to-noise considerations. It is for this reason that we do not dwell on this aspect of our data. Rather, we emphasize our finding that the sensitivity at low corrugation frequencies is independent of the spatial frequency of the luminance Gabors used to represent it. This finding runs contrary to the only two previously published reports (Pulliam, 1981; Lee & Rogers, 1997).

It is tempting to conclude that while high and low luminance spatial frequency tuned filters feed into low spatial frequency disparity tuned channels, higher spatial frequency disparity tuned channels may receive a greater input from higher luminance spatial frequency filters. However we would need to be confident that we did not, for some reason, disadvantage threshold sensitivity at low corrugation frequencies and thereby fail to see enhanced performance when using low luminance spatial frequency Gabors. Our stimulus field was larger than used in previous studies and we obtained very similar results for corrugation stimuli displayed with a constant number of corrugation cycles (Fig. 6, bottom row). Thus the thresholds that we measure at corrugation frequencies a factor of ten lower than previously reported are not differentially limited by spatial summation. Our results, while suggesting that low corrugation frequencies enjoy a greater scale support from the early spatial channels, say nothing about whether this

input remains segregated at the level of the disparity channels, an issue on which we are currently working.

4.2. Nature of the matching primitive

We have assumed in the previous discussion that each Gabor regardless of it's number of cycles provides one depth estimate. What if the matching primitive was at the level of the zero-crossing within each Gabor patch? Could it be that the high spatial frequency Gabors having more zero-crossings advantaged sensitivity and as a consequence obscure any 'hidden' weakness that high luminance spatial frequencies might have in supporting low corrugation frequencies? The results in Fig. 3 show the effect of varying sigma and according to the above line of reasoning, the number of assumed matching primitives. Since there is no significant enhancement of sensitivity at low corrugation frequencies with the larger sigma, the above argument does not hold. The enhanced sensitivity at high corrugation frequencies in Fig. 3 is likely a consequence of its higher contrast energy at its peak luminance spatial frequency.

Our main result is demonstrated in Fig. 8 where stereograms are presented for a low (A) and high frequency (B) obliquely oriented sinewave corrugation composed of high (left side) and low spatial frequency Gabors (right side). The Gabor size, Gabor perceptual contrast and Gabor densities have been kept constant. The disparity amplitude is set at 118 s arc. The low frequency corrugation (A) is very apparent for both the fields of low and high frequency Gabors, however the high frequency (B) corrugation is much less apparent, if at all, with the low frequency Gabors.

Acknowledgements

This work was supported by Canadian NSERC and MRC grants to RFH and FAAK.

References

Blakemore, C., & Campbell, F. W. (1969). On the existence of neurons in the human visual system selectively sensitive to the orientation and size of retinal images. *Journal of Physiology*, *203*, 237–260.

Blakemore, C., & Hague, B. (1972). Evidence for disparity detecting neurones in the human visual system. *Journal of Physiology*, *225*, 437–455.

Boulton, J. C., & Baker, C. L. (1991). Motion detection is dependent on spatial frequency not size. *Vision Research*, *31*(1), 77–87.

Cobo-Lewis, A. B., & Yeh, Y. Y. (1994). Selectivity of cyclopean masking for the spatial frequency of binocular disparity modulation. *Vision Research*, *34*(5), 607–620.

Crist, C., & Robinson, D. L. (1989). A large-field screen with even texture for vision research. *Visual Neuroscience*, *2*, 415–417.

DeAngelis, G. C., Ohzawa, I., & Freeman, R. D. (1995). Neuronal mechanisms undelying stereopsis: how do simple cells in the visual cortex encode binocular disparity. *Perception*, *24*(1), 3–31.

Eagle, R., & Rogers, B. (1996). Motion detection is limited by element density not spatial frequency. *Vision Research*, *36*, 545–558.

Felton, T. B., Richards, W., & Smith, R. A. (1972). Disparity processing of spatial frequencies in man. *Journal of Physiology*, *225*, 349–362.

Graham, N. V. S., & Nachmias, J. (1971). Detection of grating patterns containing two spatial frequencies. *Vision Research*, *11*, 251–259.

Graham, N. V. S. (1989). *Visual pattern analyzers*. New York: Oxford University Press.

Julesz, B., & Miller, J. E. (1975). Independent spatial-frequency-tuned channels in binocular fusion and rivalry. *Perception*, *4*, 125–143.

Lee, B., & Rogers, B. (1997). Disparity modulation sensitivity for narrow-band-filter stereograms. *Vision Research*, *37*, 1769–1778.

Pulliam, K. (1981). Spatial frequency analysis of three-dimensional vision. Visual simulation and image realism II. *SPIE*, *303*, 71–77.

Schor, C. M., & Wood, I. (1983). Disparity range for local stereopsis as a function of luminance spatial frequency. *Vision Research*, *23*, 1649–1654.

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

Schumer, R. A., & Ganz, L. (1979). Independent stereoscopic channels for different extents of spatial pooling. *Vision Research*, *19*(12), 1303–1314.

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

Stromeyer, C. F., & Julesz, B. (1971). Spatial frequency masking in vision: critical bands and spread of masking. *Journal of the Optical Society of America A*, *62*, 1221–1232.

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

Tyler, C. W. (1983). Sensory processing of binocular disparity. In C. W. Schor, & K. J. Ciuffreda, *Vergence eye movements*: *basic and clinical aspects*. London: Butterworths, 199–295.

Yang, Y., & Blake, R. (1991). Spatial frequency tuning of human stereopsis. *Vision Research*, *31*, 1177–1189.