

Anisotropies in Stereo Depth Thresholds of Spatial Patterns and the Stereo Aperture Problem

by
Elizabeth Thorpe Davis,
Robert A. King, Brial Geunter,
& Alana Anoskey

GIT-GVU-92-15
July 1992

**Graphics, Visualization & Usability
Center**

**Georgia Institute of Technology
Atlanta GA 30332-0280**

Anisotropies in Stereo Depth Thresholds of Spatial Patterns and the Stereo Aperture Problem

**Elizabeth Thorpe Davis
Robert A. King
Brian Guenter
& Alana Anoskey**

**Georgia Institute of Technology
School of Psychology and
Graphics, Visualization, and Usability Center
Atlanta, Georgia 30332-0170**

**Keywords: Stereopsis, stereoacuities, anisotropies,
orientation, stereo aperture problem**

Does a pattern's orientation affect stereoacuity? Do vertical disparities play a role in stereopsis? Furthermore, is retinal disparity encoded by changes in spatial position or spatial phase? The latter implies a "stereo aperture problem" -- only disparities perpendicular to the pattern's orientation are encoded. We found stereoacuities are worse for oblique patterns than for vertical patterns, but are worst for horizontal patterns. Moreover, for oblique patterns (unlike vertical or horizontal patterns), vertical disparity can be as effective as horizontal disparity. Results for vertical and oblique patterns suggest spatial phase encoding of disparity. In contrast, results for horizontal patterns suggest spatial position encoding of disparity

INTRODUCTION

Two fundamental purposes of vision are to recognize objects and to determine their spatial layout in the environment. Stereopsis can provide essential information in both the recognition of objects and the perception of their spatial layout in either a real or a virtual environment.

Stereopsis results from the two slightly different views of the world our laterally-displaced eyes receive (Schor, 1987; Tyler, 1983). The binocular parallax that results from these two different views can provide information about distance or depth of an object relative to the fixation plane. Partially because our two frontal eyes are laterally displaced, some investigators argue that it is only horizontal retinal disparities that provide depth information (e.g., Cumming, Johnson, & Parker, 1991). Moreover, vertical retinal disparities, at best, may provide information about egocentric distance (e.g., Bishop, 1989; Gillam and Lawergren, 1983; Mayhew and Longuett-Higgins, 1982) but, at worse, can seriously degrade stereo depth thresholds.

Some Empirical Issues

Suppose horizontal retinal disparities do help stereo depth perception but vertical disparities can hurt it. This suggests stereoacuity should be worse for an oblique spatial pattern than for the same spatial pattern with a vertical orientation. There are a few reasons for this prediction.

First, oblique patterns contain sparser information about horizontal disparity than do vertical patterns. To illustrate this point, Figure 1 shows both an oblique pattern oriented at 45 degrees and a vertical pattern. Notice that the oblique pattern is shrunk in the vertical extent by a factor of 0.707, corresponding to a factor of $\sin(\theta)$, where θ is 90 degrees for a vertical pattern. Thus, fewer pixels of a CRT display will convey horizontal disparity information for the oblique pattern than for the vertical one. Also, notice that the oblique pattern is stretched in horizontal extent by a factor of 1.414, corresponding to a factor of $\csc(\theta)$. Let's suppose that signed zero-crossings or edges in the luminance patterns are the primitives used to match corresponding locations in the left and right images (e.g., Marr and Poggio, 1979). Thus, the primitives are more sparsely

distributed in the horizontal dimension for the oblique pattern than for the vertical one.

INSERT FIGURE 1 HERE

Second, for oblique patterns, the presence of a horizontal retinal disparity can also introduce a vertical disparity, and vice versa. Figure 2 illustrates this relation between horizontal and vertical disparities for a 45 degree obliquely-oriented pattern. For a vertical or horizontal pattern, however, the presence of a horizontal retinal disparity cannot introduce a vertical disparity; nor can the presence of a vertical disparity introduce a horizontal disparity.

INSERT FIGURE 2 HERE

We will determine if stereoacuity is indeed worse for oblique patterns than for vertical ones. If so, we will examine the effects of sparseness of horizontal disparities and the presence of vertical disparities on stereo depth thresholds for spatial patterns. Moreover, we will show how, in a visual display with both vertical, oblique, and horizontal contours, one can enhance processing of the oblique contours with little or no cost for processing of vertical or horizontal contours. In these issues, retinal disparities play a unique role in stereoscopic depth perception.

Some Theoretical Issues

Although some processing components are unique to stereopsis, monocular visual processing components also can affect local stereopsis. Models of local stereopsis assume that there are monocular inputs to a binocular mechanism at an early stage of binocular processing, such as occurs in simple cells of the striate cortex. The receptive field profile for each monocular input is tuned for spatial frequency, orientation, spatial position, and spatial phase.

Some important issues for stereopsis are the receptive field profiles of the monocular inputs and how the monocular inputs are combined to form a binocular mechanism. Here we will focus only on the front-end of a local stereopsis model: We will emphasize the importance of the receptive field profiles of monocular inputs. For instance, suppose both monocular receptive field profiles are

symmetric and centered on corresponding retinal positions. The resulting binocular mechanism is optimally stimulated by a pattern presented at fixation (zero-disparity); this binocular mechanism could not distinguish a pattern positioned slightly in front of fixation (crossed disparity) from one slightly behind (uncrossed disparity). To discriminate a pattern presented at crossed or uncrossed disparity, either (1) the monocular receptive fields must be spatially offset so that their centers do not fall on corresponding retinal positions or (2) the two monocular receptive fields must differ in spatial phase. Each of these two options has different implications for stereopsis.

INSERT FIGURE 3 HERE

Spatial Position Hypothesis. Suppose that the receptive field profiles for the left and right monocular inputs to a binocular mechanism have the same shape, as suggested by McKee, Levi & Bowne (1990). Also, suppose that the centers of the receptive field profiles for the left and right inputs are not at corresponding points of the two retinae. (See Figure 3.) Thus, the difference in spatial positions of the monocular inputs determines the optimal retinal disparity for that binocular mechanism. The difference in spatial positions of the two monocular inputs can be along the horizontal direction or along any other direction. So, the orientation tuning of the mechanism does not necessarily constrain the retinal disparity.

INSERT FIGURE 4 HERE

Spatial Phase Hypothesis. Suppose instead that receptive field profiles for left and right monocular inputs to the binocular mechanism are in different spatial phases, as suggested by Freeman and his colleagues (Freeman & Ohzawa, 1990; Ohzawa, DeAngelis & Freeman, 1990; DeAngelis, Ohzawa & Freeman, 1991). For instance, receptive field profiles of the two monocular inputs could form a quadrature pair, with one symmetric and the other antisymmetric. Also, suppose that the centers of the receptive field profiles for the left and right monocular inputs are at corresponding points on the two retinae. (See Figure 4.) Notice that such an organization introduces what we call the "stereo aperture problem." That is, the binocular mechanism can only encode retinal disparities perpendicular to the mechanism's optimal orientation. For the oblique pattern shown in Figure 4, this corresponds to the retinal disparity labelled "nearest neighbor." A nearest-neighbor matching

algorithm may be used to match corresponding locations of the left and right images, where the matching primitives may be signed luminance zero-crossings in the spatial patterns.

We will show here whether a pattern's orientation affects the precision and use of horizontal and vertical disparity information, as suggested above.

METHODS

Stimuli and Apparatus

Stimuli were spatially-localized, sixth derivatives of Gaussian (D6) luminance patterns, based on Hermite polynomials (Swanson, Wilson, and Giese, 1984; Wilson, McFarlane, and Phillips, 1983). These patterns are symmetric, with a bright, blurred central bar surrounded by darker and lighter flanks. The D6 patterns are well localized in the spatial frequency domain as well as in the spatial domain; the half-amplitude bandwidth is 1.0 octave. Most stimuli had a peak spatial frequency of either 2 cpd or 4 cpd and a contrast of 1.00 or 0.945, respectively. In one condition of Experiment 1, the peak spatial frequency of an oblique comparison stimulus was either 2.828 cpd or 5.656 cpd, with a contrast of 0.96 or 0.87, respectively. In a control experiment, similar to Experiment 1, all comparison stimuli had a contrast of 0.87. Contrasts and mean luminance were calibrated with a Pritchard photometer.

Each stimulus display consisted of a comparison stimulus presented above the standard stimulus. The standard was always vertical and presented at zero disparity. The comparison stimuli were vertically or obliquely oriented in Experiments 1 and 2; in Experiment 3, comparison stimuli could also be horizontally oriented. Moreover, comparison stimuli were presented at crossed-, zero-, and uncrossed-disparity. Stimulus presentations were 500 msec for two subjects (S1 and S2) in Experiments 1 and 2. Otherwise, stimulus presentations were terminated by pressing a button. A Macintosh IIx computer controlled the stimulus presentations and a keyboard was used to collect responses.

The patterns were presented on the face of an Apple Hi-Resolution Monochrome visual display monitor with a mean luminance of 28 foot-Lamberts (ft-L). So that one inch on the display would correspond to one degree of visual angle, the visual display was positioned 57.3 inches from the subject's eyes. The display was viewed dichoptically through a modified Wheatstone stereoscope with first-surface mirrors. Fixation targets were continuously presented to maintain fixation. The subject's head was held in position with a chinrest.

Procedures

One important issue for these experiments is if stereoacuity is worse for oblique patterns than for vertical ones. If so, why? Another issue is what the relevant roles of horizontal and vertical disparities in human stereopsis are for patterns of different orientations. Each of the three experiments was designed to study different aspects of these issues.

In Experiment 1 we studied the effects of sparse horizontal disparity information in oblique patterns, as compared to vertical patterns. As described previously, oblique patterns are shrunk in vertical extent and stretched in horizontal extent. This results in sparser horizontal disparity information for the oblique patterns than for vertical patterns.

In Experiment 2 we studied the effects of potential vertical disparities in oblique patterns. As described previously, for oblique patterns the presence of a horizontal retinal disparity can introduce a vertical disparity, and vice versa. For vertical or horizontal patterns, however, there is no confounding of horizontal and vertical retinal disparities.

In Experiment 3 we studied the ability to use vertical or horizontal disparities for depth perception when patterns differed in orientation. According to the Spatial Phase Hypothesis, there should be a stereo aperture problem. That is, only retinal disparities perpendicular to the pattern's orientation will be encoded by the appropriate stereoscopic mechanisms. According to the Spatial Position Hypothesis, however, there is no stereo aperture problem. That is, neither the pattern's orientation nor the stereoscopic mechanism's optimal orientation should constrain the necessary direction of retinal disparities. Experiment 3 was designed to examine the effect of the pattern's orientation on the ability to detect vertical or horizontal disparities and to explore the potential stereo aperture problem.

Experiment 1 -- Sparseness of Horizontal Disparities in Oblique Patterns.¹ There were four conditions: (i) vertical

¹ For the 4 cpd patterns, the conditions for Experiments 1 and 2 were conducted within the same sessions. Thus, there were five blocks of

comparison stimuli; (ii) +45 degree oblique comparison stimuli; (iii) +45 degree oblique comparison stimuli that were elongated by a factor of 1.414; and (iv) +45 degree oblique comparison stimuli whose peak spatial frequency was increased by a factor of 1.414. Except for comparison stimuli in condition (iv), all stimuli within a session had the same peak spatial frequency of 2 cpd (or of 4 cpd). Four blocks of trials, one for each condition, were randomized within a session to counterbalance any order effects, et cetera. There were five sessions for each standard spatial frequency.

Experiment 2 -- Potential Vertical Disparities in Oblique Patterns.¹ There were three conditions: (i) vertical comparison stimuli whose horizontal disparities were varied; (ii) vertical comparison stimuli whose vertical and horizontal disparities were varied in equal amounts; and (iii) +45 degree oblique comparison stimuli whose horizontal disparities were varied. Within a session all stimuli had the same peak spatial frequency (viz, 2 cpd or 4 cpd). Three blocks of trials, one for each condition, were randomized within a session. There were five sessions for each standard spatial frequency.

Experiment 3 -- Vertical versus Horizontal Disparities in Patterns of Different Orientations. In half the conditions (i-iii) only the horizontal disparity of the comparison patterns was varied and, in the other half (iv-vi), only the vertical disparity was varied. The comparison stimuli were vertical in two conditions (i and iv); they were obliquely oriented at +45 degrees in two other conditions (ii and v); and they were horizontal in the two remaining conditions (iii and vi). All stimuli had a peak spatial frequency of 2 cpd. Six blocks of trials, one for each condition, were randomized within a session. There were five sessions for each subject.

Procedural Details. Spatial two-alternative, forced-choice (2AFC) trials were presented using the method of constant stimuli. Subjects initiated a trial by pressing a button. Within a trial two stimulus patterns were presented simultaneously: a comparison stimulus above the vertical standard stimulus. The subject's task was to report if the comparison stimulus appeared in front of or behind the standard. The standard stimulus was always presented in

trials, one for each condition, presented in randomized order within a given session. There were five sessions per subject.

the same depth plane as the fixation target. (The lateral spatial location of the comparison stimulus, with respect to the standard stimulus, was jittered randomly from trial to trial, within ± 8 minutes of arc. Thus, monocular cues could not disambiguate whether the comparison stimulus was in front of or behind the standard.)

The retinal disparity of the comparison stimulus was chosen, without replacement, from a set of stimuli that usually spanned a range from 50% to 100% correct response, as determined from pilot studies. On almost half of the trials, chosen at random, the comparison stimulus was presented at crossed disparity; on almost half of the trials the comparison was presented at uncrossed disparity; and on the remainder of the trials the comparison was presented at zero disparity. If the subject thought the comparison stimulus was in front of the standard, the number "4" was pressed on the keypad. If the subject thought the comparison was behind the standard, the number "1" was pressed. When the subject was unsure of the depth of the stimulus, he or she was told to guess and, on average, to make half of those guesses "behind" the standard. Feedback was given for incorrect responses.

A psychometric function was generated from a total of 100 trials per data point for each condition for each subject. We used seven to ten data points to determine a given psychometric function. A maximum-likelihood curve-fitting procedure was used to fit a Quick function to the data (Harvey, 1986; Swanson and Birch, 1992). The curve-fitting procedure estimated the stereoacuity threshold retinal disparity that corresponded to 75% correct response; it also calculated the 95% confidence intervals for the threshold estimate.

Subjects

Three subjects participated in these experiments. Two subjects (S1 and S2) were naive about the purposes of the experiments during testing; all three subjects were naive about their results during testing. One subject is an emmetrope (S2) and the other two are corrected myopes; all subjects had a best-corrected distance acuity of 20/20. All three subjects had been screened for visual acuity with a Lovie-Bailey eyechart and for stereopsis with a Titmus stereo test prior to experimentation. None of the subjects had any astigmatism or other uncorrected visual problem.

RESULTS

Some Empirical Issues

Overall, depth thresholds are larger for obliquely-oriented comparison patterns than for vertical comparison patterns when only the horizontal disparities are manipulated. That is, we found an apparent "oblique effect" for stereoacuity. The regular oblique and regular vertical thresholds in Figures 5 and 6 show that this difference in stereoacuties is significant for each observer for both 2 cpd and 4 cpd patterns.

Experiment 1 -- Sparseness of Horizontal Disparities in Oblique Patterns. The oblique effect for stereoacuity cannot be explained by the sparseness of horizontal disparity information in the oblique patterns. Remember that the sparseness of horizontal disparities in oblique patterns is caused by two factors: (1) the vertical extent is shrunk and (2) the horizontal extent is stretched. Yet, compensation for either of these factors did not significantly improve stereoacuity across the three subjects.

INSERT FIGURE 5 HERE

First, to compensate for the shrunk vertical extent, we had elongated the oblique pattern by a factor of 1.414. Figure 5 shows that elongation of the oblique pattern only sometimes improved stereoacuity for two subjects (S1 and S3), but did not help the other subject. In general, stereoacuity thresholds were no better for the longer oblique patterns than for the oblique patterns of regular length.

Second, to compensate for the stretched horizontal extent, we increased the peak spatial frequency of the oblique pattern by a factor of 1.414. Figure 5 shows that increasing the peak spatial frequency of the oblique pattern improved stereoacuity for one subject (S1), but not for the other two subjects. In fact, stereoacuity thresholds often were worse for the thinner oblique patterns than for the regular oblique patterns.

Some investigators have found stereoacuity is worse at lower contrasts than at higher contrasts (e.g., Halpern and Blake, 1988; Legge and Gu, 1989). Thus, the slightly lower contrast of the thinner oblique pattern could have a confounding effect on our results. To

check for this, in a control experiment for S3, the physical contrasts were carefully equated for the thinner and regular oblique patterns as well as for the vertical pattern. Yet, stereoacuity for the thinner oblique pattern still was significantly worse than for the regular oblique pattern.

Experiment 2 - Potential Vertical Disparities in Oblique Patterns. The oblique effect probably cannot be explained by the relatively small vertical disparities present in the oblique patterns. Remember that the left and right images are laterally displaced to produce binocular parallax. For oblique patterns, unlike vertical ones, this lateral displacement can produce vertical retinal disparities in addition to the intended horizontal disparities.

INSERT FIGURE 6 HERE

In a pilot study we had found that stereoacuity was seriously degraded for vertical patterns with a constant, relatively large vertical disparity of 0.4 degree, although the pattern still appeared to be fused. In the present studies, however, the amount of vertical disparity was relatively small (less than 1' of arc). Figure 6 shows that in the present study the small amount of vertical disparity in the modified vertical patterns had no effect for two subjects (S1 and S2) and caused a small but insignificant increase in threshold for the third subject (S3).

Some Theoretical Issues

Some theoretical considerations are whether retinal disparity is encoded by spatial position or by spatial phase, as described in the introduction.

Spatial phase encoding of disparities implies there is a stereo aperture problem. That is, the binocular mechanism only encodes retinal disparities located perpendicular to the optimal orientation of that mechanism. We also assume stereo depth threshold is determined by a binocular mechanism whose optimal orientation is similar to that of the spatial pattern. For instance, an obliquely oriented binocular mechanism would encode the stereo depth threshold of an oblique pattern. Thus, the matching correspondence points in the left and right images would be "nearest-neighbors," as shown in Figure 2.

We have calculated the nearest-neighbor stereo depth thresholds for the regular oblique patterns. These are shown in Table 1. The nearest-neighbor stereo depth thresholds for vertical patterns require no calculation; they correspond to the pattern's horizontal retinal disparity at stereo depth threshold.

INSERT TABLE 1 HERE

In general, the nearest-neighbor stereo depth thresholds of a given subject in a given condition are very similar for the vertical and oblique patterns. Moreover, the 95% confidence interval for the stereo depth threshold of the vertical pattern usually brackets the nearest-neighbor stereo depth threshold of the relevant oblique pattern. Conversely, the 95% confidence interval for the nearest-neighbor stereo depth threshold of an oblique pattern usually brackets the stereo depth threshold of the relevant vertical pattern. This suggests there is a stereo aperture problem, as predicted from the spatial phase hypothesis. (A stereo aperture problem is not necessarily predicted from the spatial position hypothesis.)

Experiment 3 - Vertical versus Horizontal Disparities in Spatial Patterns of Different Orientations. The results of Experiment 3 also addresses the potential stereo aperture problem and the use of a nearest-neighbor matching algorithm.

Suppose only disparities perpendicular to the orientation of the pattern are encoded. Thus, for a horizontal pattern only vertical disparities are encoded. Similarly, for a vertical pattern only horizontal disparities are encoded. Furthermore, for both vertical and horizontal patterns, if the vertical disparity is manipulated there is no horizontal disparity introduced and vice versa.

For oblique patterns, however, the nearest-neighbor matching correspondence points are at an oblique disparity (90 degrees from the orientation of the pattern). Furthermore, for oblique patterns, if the vertical disparity is manipulated, then there is a horizontal disparity introduced and vice versa. Thus, for oblique patterns oriented at +45 degrees, a vertical disparity of a fixed amount should have the same effect on stereo depth threshold as a horizontal disparity of the same fixed amount.

INSERT FIGURE 7 HERE

Horizontal disparity is more effective than is vertical disparity, as shown in Figure 7, for spatial patterns that are either vertical or horizontal. That is, stereo depth thresholds usually are lower for these patterns when horizontal disparity is manipulated than when vertical disparity is manipulated. (In fact, when only vertical disparity is manipulated, the estimated stereo depth thresholds for vertical and horizontal patterns are outside the range of tested retinal disparities. That is, the estimated stereoacuties shown in Figure 7 for these patterns are extrapolations. Moreover, these extrapolated threshold estimates are very noisy, as indicated by the large 95% confidence intervals.)

So, the results for horizontal patterns are contrary to the predictions based on the spatial phase hypothesis and the implied stereo aperture problem. Instead, for horizontal patterns, stereoacuity thresholds are lower when horizontal retinal disparity is manipulated than when vertical retinal disparity is manipulated. These results are consistent, however, with the spatial position hypothesis.

In contrast, results for the oblique patterns are consistent with predictions based on the spatial phase hypothesis and the implied stereo aperture problem. Results with 45 degree oblique patterns show that horizontal disparity is *not* more effective than is vertical disparity. For patterns oriented at +45 degrees, vertical disparity of a fixed amount has approximately the same effect on stereo depth threshold as does a horizontal disparity of the same fixed amount.

DISCUSSION

Some Empirical Issues

Our results show that there is an apparent "oblique" effect for stereoacuity. If horizontal disparity is manipulated to produce binocular parallax, the stereo depth threshold for an oblique pattern requires a larger horizontal retinal disparity than is required for a vertical pattern. Neither sparseness of horizontal disparity information nor the presence of vertical disparities in oblique patterns could adequately account for the oblique effect. That is, shrinkage in the vertical dimension, stretching in the horizontal dimension, and the presence of small amounts of vertical disparities in oblique patterns seemed to have little or no effect on stereoacuity.

The difference in stereoacuity for vertical versus oblique patterns, however, could be accounted for by a stereo aperture problem. That is, only the disparities located perpendicular to the pattern's orientation matter. When the "nearest-neighbor" stereoacuity was calculated for oblique patterns, it seemed to approximately match that of the vertical patterns, as shown in Table 1.

It may seem that horizontal disparity only matters for vertical patterns. Yet, horizontal disparity does play an important role in stereopsis, as shown by the results of Experiment 3. Those results indicated that, in general, stereoacuity was better if only horizontal disparities were present than if only vertical disparities were present. That is, for both vertical and horizontal patterns manipulation of horizontal disparity resulted in smaller depth thresholds than did manipulation of vertical disparity.

For oblique patterns (unlike vertical or horizontal patterns) one can translate vertical disparities into horizontal disparities and vice versa. (See Figure 2.) Thus, it is not surprising that for an oblique pattern of +45 degrees, stereo depth threshold measured by manipulating vertical disparity is approximately the same as that measured by manipulating horizontal disparity.

Everyday experience provides obliquely-oriented receptive fields of binocular neurons the opportunity to use both vertical and

horizontal disparities for stereo depth cues. Remember that for either vertical or horizontal patterns, the presence of a horizontal disparity does not introduce a vertical disparity, nor vice versa. Thus, for vertical or horizontal oriented binocular receptive fields, vertical disparity may not be very useful in providing depth cues.

Some Theoretical Issues

It does seem that for vertical and oblique patterns, the nearest-neighbor matching algorithm is used to match corresponding primitives from the left and right images. These results for vertical and oblique patterns are consistent with disparity encoded by spatial phase differences in monocular inputs, as suggested by DeAngelis et al. (1991) and by our stereo aperture problem.

But the nearest-neighbor matching algorithm is not used for horizontal patterns; our data suggest that binocular, horizontally-oriented mechanisms are less sensitive to depth information than are vertical or oblique mechanisms. Moreover, because horizontal disparity is more helpful than is vertical disparity, the results for horizontal patterns are consistent with disparity encoded by spatial position. If disparity is encoded by spatial position of the monocular inputs to a binocular mechanism, then disparity is not necessarily constrained by the stereo aperture problem. That is, the difference in spatial positions of the two monocular inputs may be along the horizontal direction (or any oblique direction).

This is a hybrid interpretation utilizing the concepts that (1) horizontal disparities are important; (2) binocular, oriented mechanisms encode disparity information from spatial phase information for mechanisms that are vertically or obliquely oriented; and (3) binocular, oriented mechanisms encode information from spatial position information for mechanisms that are horizontally oriented.

Our hybrid interpretation is consistent with some recent sensory physiological studies of binocular simple cells in the cat (De Angelis, Ohzawa, & Freeman, 1991). They reported that the left and right monocular receptive field profiles are matched for binocular simple cells tuned to horizontal orientation. For binocular simple cells tuned to vertical, however, the left and right monocular receptive field profiles usually are different; for vertically tuned binocular mechanisms, the difference in monocular spatial phases

can be up to 180 degrees. Thus, horizontally tuned binocular simple cells would use spatial position encoding of disparity, but vertically tuned binocular simple cells could use spatial phase encoding of disparity.

Our hybrid interpretation is not consistent with a recent psychophysical study by Liu and colleagues (Liu, Tyler, Schor, & Ramachandran, 1992; Liu, personal communication). They reported human observers are more efficient in encoding stereo depth for vertical patterns based on spatial position disparity than based on spatial phase disparity. In that psychophysical study they used Gabor patterns of 4 cpd and higher, with a contrast of 10%. Schor and Wood (1983) found that for high spatial frequency patterns stereoacuity is approximately a constant (in seconds of arc); these results are consistent with spatial position encoding of disparity for high spatial frequency patterns and with Liu et al.'s (1992) results. Schor and Wood also reported that for lower spatial frequency patterns (2.4 cpd and below), stereoacuity is approximately a constant proportion of the pattern's spatial period; these results are consistent with spatial phase encoding of disparity for lower spatial frequency patterns and with our results for 2 cpd patterns. Moreover, the contrast of Liu et al.'s patterns is much lower than the contrasts used in our studies. Perhaps there are binocular mechanisms that are sensitive at lower contrasts, but saturated at higher contrasts; perhaps these binocular mechanisms encode disparity by changes in spatial position. Conversely, perhaps there are binocular mechanisms that are insensitive at lower contrasts, but respond at higher contrasts; perhaps these binocular mechanisms encode disparity by changes in spatial phase.

Some Implications for Stereoscopic Displays

These results also have implications for the design of stereoscopic displays and of virtual environments. First, oblique contours can require larger horizontal retinal disparities for depth perception than do vertical contours. Second, if the observer's head is free to move, tilting the head can change the orientations of contours presented on the visual display. One needs to compensate for both of these effects and their potential problems. A possible solution is to introduce some vertical disparity into the display as well as the necessary horizontal disparity. These small vertical disparities will not affect depth thresholds for vertical or horizontal contours in the proximal stimulus, but will enhance depth

perception for oblique contours. Furthermore, even tilting the head 90 degrees will not eliminate horizontal retinal disparities in the modified visual display. Thus, an optimal combination of both vertical and horizontal disparities in the distal stimulus can overcome both potential problems with the proximal stimulus.

ACKNOWLEDGEMENTS

We thank Shane McWhorter for comments on an earlier version of this manuscript. We also appreciate helpful comments from Izumi Ohzawa, who pointed us to a very relevant article on cat physiology. Sample programs for stimulus generation and experimental control, written in Lightspeed Pascal, were supplied by Hugh Wilson, and were adapted for our experiments. Bill Swanson provided us with programs for fitting the Quick psychometric functions. Most of these results have been reported at the 1992 annual meeting of the Association for Research in Vision and Ophthalmology in Sarasota, Florida, and at the 1992 SPIE/IS&T's Electronic Imaging Science and Technology's Human Vision, Visual Processing, and Digital Display III Conference in San Jose, California.

**Stereo Depth Thresholds
Table 1**

<u>Vertical Pattern & Oblique Pattern Disparities</u>			
	Horizontal (& Nearest Neighbor)	Horizontal	Nearest Neighbor
<u>Subject 1</u>			
2 cpd	20.0" (17.1"-23.0")	25.1" (22.0"-28.3")	17.8" (15.6"-20.0")
	20.0" (16.8"-24.2")	31.6" (27.9"-37.7")	22.4" (19.7"-26.6")
	10" (8.6"-12.9")	15.8" (14.1"-16.6")	11.2" (10.0"-11.8")
4 cpd	15.8" (13.6"-19.0")	31.6" (27.9"-36.1")	22.4" (19.7"-25.6")
<u>Subject 2</u>			
2 cpd	6.3" (5.0"-7.9")	15.8" (13.4"-19.1")	11.2" (9.5"-13.5")
	15.8" (12.9"-19.0")	20.0" (16.4"-25.2")	14.1" (11.6"-17.8")
	10.0" (8.9"-13.2")	15.8" (13.2"-18.4")	11.2" (9.3"-13.0")
<u>Subject 3</u>			
2 cpd	39.8" (29.7"-47.2")	50.1" (35.1"-65.0")	35.4" (24.8"-46.0")
	20.0" (17.7"-28.3")	25.1" (22.4"-32.3")	17.8" (15.8"-22.8")
	20.0" (17.4"-28.3")	31.6" (28.4"-43.8")	22.4" (20.0"-30.9")
4 cpd	15.8" (10.2"-18.7")	25.1" (22.4"-33.1")	17.8" (15.9"-23.4")

Note 1: The values in parentheses correspond to the 95% confidence limits. The relevant comparisons are to determine (1) if the stereo depth threshold for the vertical pattern falls within the 95% confidence interval for the oblique nearest-neighbor stereo depth threshold, and conversely, (2) if the nearest-neighbor stereo depth threshold for the oblique pattern falls within the 95% confidence interval for the vertical stereo depth threshold.

REFERENCES

- Cumming, B.G., Johnston, E.B., & Parker, A.J. (1991) Vertical disparities and perception of three-dimensional shape. Nature, 349, 411-413.
- Bishop, P.O. (1989). Vertical disparity, egocentric distance and stereoscopic depth constancy: A new interpretation. Proceedings of the Royal Society, London, B, 237, 445-469.
- De Angelis, G.C., Ohzawa, I. & Freeman, R.D. (1991). Depth is encoded in the visual cortex by a specialized receptive field structure. Nature, 352, 156-159.
- Freeman, R.D. & Ohzawa, I. (1990). On the neurophysiological organization of binocular vision. Vision Research, 30, 1661-1676.
- Gillam, B. & Lawergren, B. (1983). The induced effect, vertical disparity, and stereoscopic theory. Perception and Psychophysics, 34, 121-130.
- Halpern, D.L. & Blake, R. (1988). How contrast affects stereoacuity. Perception, 17, 483-495.
- Harvey, (1986). Efficient estimation of sensory thresholds. Behavior Research Methods, Instruments, and Computers, 18, 623-632.
- Legge, G.E. & Gu, Y. (1989). Stereopsis and contrast. Vision Research, 29, 989-1004.
- Liu, L., Tyler, C.W., Schor, C.M. and Ramachandran, V.S. (1992) Position disparity is more efficient in encoding depth than phase disparity. Investigative Ophthalmology & Visual Science Supplement, 33, 1373.
- Marr, D. & Poggio, T. (1979) A computational theory of human stereo vision. Proc. R. Soc. Lond. B, 204, 301-328.
- McKee, S.P., Levi, D.M. & Bowne, S.F. (1990). The imprecision of stereopsis. Vision Research, 30, 1763-1779.

- Ohzawa, I., DeAngelis, G.C. & Freeman, R.D. (1990) Stereoscopic depth discrimination in the visual cortex: Neurons ideally suited as disparity detectors. Science, 249, 1037-1041.
- Schor, C.M. (1987) Spatial factors limiting stereopsis and fusion. Optic News, 13, 14-17.
- Schor, C.M. & Wood, I. (1983) Disparity range for local stereopsis as a function of luminance spatial frequency. Vision Research, 23, 1649-1654.
- Swanson, W.H. & Birch, E.E. (1992). Extracting thresholds from noisy psychophysical data. Perception and Psychophysics, 51, 409-422.
- Swanson, W.H., Wilson, H.R. & Giese, S. (1984). Contrast matching data predicted from contrast increment thresholds. Vision Research, 24, 63-76.
- Tyler, C.W. (1983). Sensory processing of binocular disparity. In Schor, C.M. & Ciuffreda, K.J. (eds), Vergence eye movements: Basic and clinical aspects (pp. 199-295). London: Butterworths.
- Wilson, H.R., McFarlane, D.K. & Phillips, G.C. (1983). Spatial tuning of orientation selective units estimated by oblique masking. Vision Research, 23, 873-882.

FIGURE LEGENDS

Figure 1. An oblique pattern oriented at 45 degrees is stretched in the horizontal dimension by a factor of 1.414 and shrunk in the vertical dimension by a factor of 0.707, as compared to a vertical pattern. The amount of stretching and shrinkage depends upon the orientation of the oblique pattern. If a vertical pattern is assigned an orientation of 90 degrees, an oblique pattern of orientation θ is stretched in the horizontal dimension by a factor of $\csc(\theta)$ and shrunk in the vertical dimension by a factor of $\sin(\theta)$.

Figure 2. For oblique patterns, binocular parallax creates not only the intended horizontal retinal disparities but also potential vertical disparities. In the figure, the left and right images are horizontally displaced to create uncrossed disparity. Notice that a pattern oriented at 45 degrees has potential vertical disparities equal in magnitude to the horizontal disparities in the display. The vector labelled "nearest neighbor" corresponds to the smallest retinal disparity between the two oblique patterns; this disparity is oriented perpendicular to the pattern's orientation. The nearest neighbor disparity is smaller than the horizontal disparity by a factor of 0.707 for an oblique pattern oriented at 45 degrees. For a vertical pattern, however, the nearest neighbor disparity would be equal to the horizontal retinal disparity. Furthermore, for a vertical or horizontal pattern, introducing horizontal retinal disparities does not cause potential vertical disparities.

Figure 3: Spatial Position Disparity Encoding. The figure shows the receptive fields of the left and right monocular inputs to a binocular mechanism. Both monocular receptive fields have the same shape (e.g., symmetric). Retinal disparity is encoded by the difference in spatial positions of the two monocular inputs -- they do not fall on corresponding points of the two retinæ. The monocular inputs shown here would encode uncrossed disparity for a distal stimulus farther away than the fixation plane.

Figure 4: Spatial Phase Disparity Encoding. The figure shows the receptive fields of left and right monocular inputs to a

binocular mechanism that uses spatial phase encoding of disparity. The receptive field profiles for the left and right monocular inputs are different. In the example shown, the right monocular receptive field is symmetric and the left monocular receptive field is anti-symmetric (viz, the two monocular inputs form a quadrature pair). Retinal disparity is encoded by the difference in the spatial phases of the two monocular inputs -- the centers of the monocular receptive fields do fall on corresponding points of the two retinae. The monocular inputs shown here would encode uncrossed disparity for a distal stimulus farther away than the fixation plane. Spatial phase encoding of disparity causes the "stereo aperture problem": The binocular mechanism can only encode retinal disparities perpendicular to the mechanism's optimal orientation. This would correspond to the "nearest-neighbor" disparity shown in Figure 2.



Figure 5: Sparseness of Horizontal Retinal Disparities.

Stereoacuity is shown by solid bars for vertical D6 luminance patterns with a peak spatial frequency of 2 cpd (panel a) or of 4 cpd (panel b). Data bars with wide diagonal stripes show stereoacuities for 45 degree oblique D6 luminance patterns of regular length and with a peak spatial frequency of 2 cpd (panel a) or of 4 cpd (panel b). To compensate for shrinkage in the vertical dimension, we elongated the oblique pattern by a factor of 1.414 -- data are shown by the bars with the stippled pattern for the long oblique patterns. To compensate for stretching in the horizontal dimension, we increased the peak spatial frequency of the oblique pattern by a factor of 1.414 -- data are shown by the bars with thin diagonal strips for the thin oblique patterns. The error bars show the 95% confidence limits for each data bar.

Figure 6: Presence of Vertical Retinal Disparities. All


patterns are D6 luminance patterns with a peak spatial frequency either of 2 cpd (panel a) or of 4 cpd (panel b). Stereoacuity is shown by solid bars for vertical patterns that have only horizontal retinal disparity. Bars with vertical stripes show data for vertical patterns that have both horizontal and vertical retinal disparities. Bars with diagonal stripes show data for 45 degree oblique patterns with horizontal disparity manipulated. Error bars show the 95% confidence intervals for each data bar.

Figure 7: Comparisons of Vertical and Horizontal Disparities for Spatial Patterns at Different Orientations. Stereoacuities for 2 cpd D6 luminance spatial patterns are shown in this figure. The three sets of dark bars show data for spatial patterns in which horizontal disparity was manipulated. The three sets of lighter bars show data for spatial patterns in which vertical disparity was manipulated. The solid bars show stereoacuities for vertical stimuli; the bars with diagonal stripes show stereoacuities for 45 degree oblique stimuli; and the stippled bars show stereoacuities for horizontal stimuli. Error bars show the 95% confidence intervals for each data bar. The horizontal line through the data bars show the upper limit of retinal disparities tested.




Left Image

Right Image




Left Monocular RF

Right Monocular RF






Right Monocular RF

Left Monocular RF





2 CPD Luminance Stereo Patterns Sparseness of Horizontal Disparities

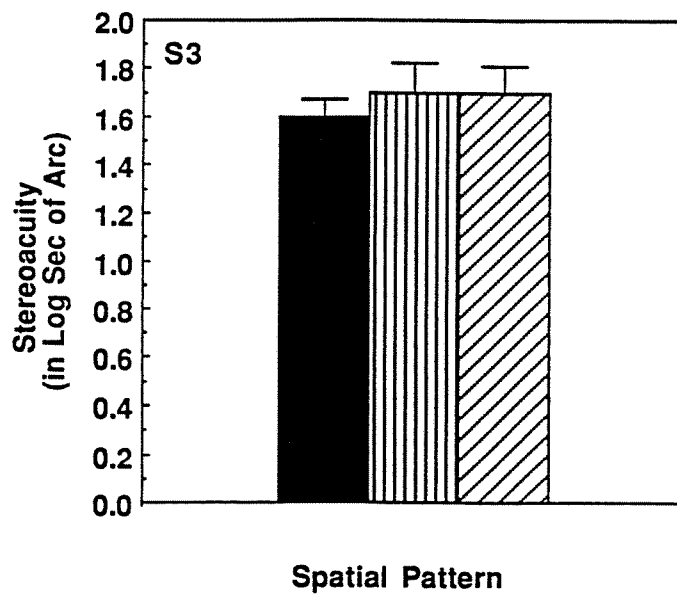




Spatial Pattern



4 CPD Luminance Stereo Patterns Sparseness of Horizontal Disparities



2 CPD Luminance Stereo Patterns Vertical Disparity



4 CPD Luminance Stereo Patterns Vertical Disparity



Comparison of Vertical & Horizontal Disparities for Patterns at Different Orientations

