



Detecting disparity in two-dimensional patterns

Bart Farell *

Institute for Sensory Research, Syracuse University, 621 Skytop Road, Syracuse, NY 13244-5290, USA

Received 26 June 2001; received in revised form 26 August 2002

Abstract

One can measure the disparities between two retinal images in several different ways. Experiments were conducted to identify the measure that is invariant at the threshold for detecting the disparity of two-dimensional patterns. The patterns used were stereo plaids, which permit a partial dissociation between the disparity of the pattern and the disparities of its one-dimensional components. For plaids with near-horizontal disparities, thresholds are limited by a disparity phase shift equal to the threshold phase shift for single gratings. For non-horizontal disparities, thresholds are elevated, yet are still phase-limited. In no disparity direction are thresholds for detecting disparity determined by the spatial extent of the plaids' disparity. Effects of the number and the orientation of components with task-relevant disparities indicate that plaid thresholds are limited by the disparity of the plaid's one-dimensional components. No evidence was found that these components form any higher-order pattern that can be used in detecting disparity. Oblique and near-vertical disparities generate elevated thresholds at a stage beyond component disparity detection. This second stage combines component disparities, which are ambiguous about depth, into pattern disparities capable of supporting veridical depth perception.

© 2003 Elsevier Science Ltd. All rights reserved.

Keywords: Stereopsis; Disparity; Stereoacuity; Threshold vision; Component analysis

1. Introduction

Whether a stereo threshold is measured in seconds, minutes or even degrees depends on stimulus contrast, spatial frequency, orientation, retinal position, and duration, among other variables. Despite these multiple influences, threshold disparities display certain invariances over fairly wide ranges in stimulus parameters. In particular, changes in spatial frequency and orientation can, within limits, leave disparity thresholds unchanged. These invariances give hints about how stereo matches and detection decisions are made. But we might overlook them if we measure stereo thresholds only as horizontal spatial disparities.

Measuring horizontal spatial disparity thresholds for one-dimensional (1-D) patterns such as lines or gratings shows them to vary reciprocally with the sine of the angle the pattern makes with the horizontal (Blake, Camisa, & Antoinetti, 1976; Ebenholtz & Walchli, 1965; Farell & Ahuja, 1996; Morgan & Castet, 1997; Ogle,

1955). However, expressing the grating data as disparity phase angles shows disparity to be invariant at threshold, so long as the grating's orientation is not too close to horizontal. Similarly, if we measure horizontal spatial disparity thresholds for sinusoidal gratings we find thresholds decrease linearly as spatial frequency increases, at least up to moderate frequencies, 2.5 or 3.5 c/d or so (Badcock & Schor, 1985; Schor & Wood, 1983; Schor, Wood, & Ogawa, 1984; cf. Smallman & MacLeod, 1994). These thresholds, too, are constant when expressed as a disparity phase angle. Examples of invariance over orientation and spatial frequency appear in Fig. 1, which shows disparity detection threshold phase angles for sinusoidal luminance gratings displayed in large (7.8°) circular fields (Farell & Ahuja, 1996).

Closely connected with this disparity threshold phase constancy is the 'aperture problem' shared by motion and stereo. Two samples of a 1-D pattern, be they motion samples or binocular images, can be matched in any direction within 180° that is not parallel to the pattern's orientation. In the case of stereo it has been the convention since Wheatstone's time to take the matching direction to be horizontal or epipolar, the dominant direction of interocular perspective disparities. Recent

* Tel.: +1-315-443-9717/9714; fax: +1-315-443-1184.

E-mail address: bart_farell@isr.syr.edu (B. Farell).




Fig. 1. Grating disparity thresholds. Threshold phase angles for sinusoidal gratings are shown as a function of grating spatial frequency and orientation. The threshold for non-horizontal gratings is approximately constant across variation in spatial frequency and orientation, while those for horizontal gratings vary with frequency and are higher. Gratings had a contrast of 10% and appeared for 180 ms within 7.8° circular windows in two-interval forced choice trials. Error bars are ± 1 s.e.m.

psychophysical evidence, however, questions whether horizontal matching applies to 1-D stimuli (Farell, 1998; Farell & Ahuja, 1996; Morgan & Castet, 1997; van Ee & Schor, 2000). The notion that the stereo matching direction for these stimuli may be perpendicular to their orientation fits nicely with physiological evidence on the preferred disparity direction of cortical neurons (e.g., LeVay & Voigt, 1988; Maske, Yamane, & Bishop, 1986; Ohzawa & Freeman, 1986), including evidence showing small departures from isotropy (Anzai, Ohzawa, & Freeman, 1999; however, see Cumming, 2002).

The disparity direction usually used in studies of spatially two-dimensional (2-D) patterns is horizontal. However, these patterns can have a disparity along any direction. This direction is unambiguous; 2-D patterns are not subject to the aperture problem.¹ Nevertheless, stereo matching might disregard the 2-D spatial properties of such stimuli—the features that impart an unambiguous direction to its retinal correspondences—and operate instead on the pattern's 1-D components. Matching might then take place along multiple directions: every direction perpendicular to a detectable 1-D

¹ Periodic patterns, whether they are spatially 1-D or 2-D, are susceptible to matching ambiguities (e.g., the wallpaper effect) that are distinct from the classical aperture problem. In the case of 2-D patterns there can be multiple directions of periodicity, each providing a potential match at integral multiples of the period in that direction. I ignore these complications in this paper, concentrating instead on the simpler distinction made clear if one compares non-periodic 1-D patterns, which are intrinsically ambiguous as to matching direction, and non-periodic 2-D patterns, which are unambiguous.

component. Indeed, effects of disparity adaptation suggest just such a matching process (Farell, 1998). Because this evidence comes from suprathreshold disparities, it raises but does not answer the question of what determines disparity threshold for 2-D stimuli.

Disparity thresholds for several different 1-D waveforms have been examined by Heckmann and Schor (1989). They asked, for example, whether two waveforms having the same spatial frequency content but different luminance gradients would have the same thresholds. The brief answer is 'Yes'. Disparity thresholds for compound gratings varied with the spatial frequency and contrast of the sinusoidal components, not with the relative phase of these components and therefore not with the shape of the grating's luminance profiles or the steepness of their luminance gradients. As one would expect from a component analysis of disparity, Heckmann and Schor (1989) found that thresholds for compound gratings equaled the threshold for the component whose disparity the observer was most sensitive to. Binocular fusion, too, depends on shared spatial frequency content between the stereo half-images, rather than on similar luminance profiles (Levinson & Blake, 1979). While the Marr and Poggio (1979) independent-channels notion of stereo matching receives support from these data, their proposed coarse-to-fine channel interactions have not been confirmed (Rohaly & Wilson, 1993; Smallman & MacLeod, 1997; but see Farell, Li, & McKee, 2002). However, Mayhew and Frisby (1978), using 2-D patterns (filtered random-dot stereograms), did find evidence for channel interactions in the identification of cyclopean figures. Figures depicted within each of two widely spaced spatial-frequency bands were identified better than those depicted within either band alone. But as Heckmann and Schor (1989) point out, this interaction might not involve stereo matching or disparity detection. Cyclopean identification presumably takes place later than disparity detection and combines inputs from separate disparity channels. Perhaps for similar reasons, the perceived depth of compound gratings does not show consistent dependence on the disparity of sinusoidal components (Boothroyd & Blake, 1984).

The use the visual system makes of 1-D components in detecting the disparity of 2-D patterns not so clear. The shape of local 2-D features and the steepness of luminance gradients will depend on how the pattern's 1-D components are combined, as in the case of complex 1-D patterns. But 2-D patterns have properties that 1-D patterns do not have. For example, the disparity direction of a 2-D pattern may differ from the perpendicular disparity direction of any of its 1-D components. We know that perpendicular disparity direction has little effect on grating thresholds, except for directions near vertical (Fig. 1). Therefore the elevated detection thresholds found for 2-D patterns with near-vertical

disparities (Morgan & Castet, 1997; Nielsen & Poggio, 1984; Stevenson & Schor, 1997; Westheimer, 1984) cannot readily be explained by appealing to the observer’s sensitivity to the disparity of the components: A 2-D pattern with near-vertical disparity might contain 1-D components all of whose perpendicular disparities are near-horizontal.

We expect disparity thresholds for 2-D patterns to display an invariance across stimulus scales, just as they do for 1-D patterns. The form that thresholds take on depends on whether they are limited by a constant noise or a noise that scales with the stimulus, and whether the effective stimulus is the 1-D component or the 2-D pattern. So, for 2-D patterns with a particular disparity direction—horizontal, say—threshold invariance could take the form of a constant horizontal spatial disparity (e.g., 2 min of arc), a fixed proportion of the horizontal extent of the pattern or of the pattern’s periodicity (i.e., a constant pattern phase angle), or a fixed proportion of the components’ periodicity (i.e., a constant component phase angle). In order to find out what limits disparity detection in 2-D patterns, I measured thresholds for stereo plaids: patterns made up of a pair of differently oriented 1-D components, each having a different disparity (Adelson & Movshon, 1984; Farell & Ahuja,

1996). When the components are static, sinusoidal, and briefly presented, they are not seen as a pair of transparent gratings, but rather as cohering in depth and possessing the 2-D spatial structure of a plaid (Farell, 1998; Farell & Li, 2002).

The disparity of a stereo plaid—that is, the disparity of its 2-D spatial structure or “features” such as grating intersections—cannot be derived directly from the disparities of the components, for it depends on the components’ orientations as well as on their disparities. As shown schematically in Fig. 2, changing the orientation of one of the components, even a zero-disparity component, can change both the direction and magnitude of the disparity of the plaid. Plaid disparity can be specified by the same intersection-of-constraints solution that has been applied to 2-D motion (Adelson & Movshon, 1982; Farell, 1998). A plaid’s disparity direction ϕ and magnitude ρ can be expressed as functions of the component disparities:

$$\phi = \theta_1 - \arctan \left\{ \frac{A_1 * \cos(\theta_2 - \theta_1) - A_2}{A_1 * \sin(\theta_2 - \theta_1)} \right\}, \quad (1)$$

$$\rho = \left| \frac{A_1}{\cos(\phi - \theta_1)} \right|, \quad (2)$$




Fig. 2. Changes in plaid disparity with changes in orientation of zero-disparity component. Lines of differing orientations in left column represent gratings depicting the components of the stereo plaids in the second and third columns. One grating (the target) has non-zero disparity, indicated by the offset between the left (L) and right (R) eyes’ images; the other grating (mask) is imaged at corresponding retinal location on both eyes (LR). The arrow shows the disparity of the 2-D features (taken as the lines’ intersections) directed from the left eye’s image to the right eye’s image. Both the direction and the magnitude of disparity vary with mask orientation. Panels (a) and (c) depict depth-consistent plaids—the horizontal component of the plaid’s disparity is in the same direction, left or right, as that of the target grating’s disparity—while panel (b) depicts a depth-reversed plaid. Examples of each of the three cases can be seen by fusing the stereograms occupying the middle and right columns.

where θ_1 and θ_2 are component disparity directions normal to the gratings' orientations, and Δ_1 (non-zero) and Δ_2 are the respective spatial disparity magnitudes. Note that a plaid's disparity can have a horizontal component whose sign differs from that of each of the constituent gratings. As a result, two gratings individually seen behind the fixation plane can combine into a plaid seen in front of the fixation plane (Farell, 1998).

By varying the orientation of the components we can manipulate the direction and magnitude of the disparity of the plaids' 2-D features (whatever these may be); disparity thresholds, if influenced by the plaid's 2-D spatial properties, should vary accordingly. In a recent study, van Ee, Anderson, and Farid (2001) looked at something similar. They measured disparity thresholds for bars or letters presented as either adjacent or partially overlapping pairs. They found no effect of the disparity of junction-features formed by the occluding contours, even when these were large compared to the disparities of the bars or letters. However, the interpretation of these results is unclear. The issue of whether the bars and letters were perceived as segregated in depth or as coherent was not answered. This makes the status of the junctions—whether they are what may be called first-order or second-order features—uncertain. And the 2-D direction of the junction disparities was confounded with bar orientation and letter disparity magnitude. The contrast of these results with the data presented here will be taken up in Section 4.

The first experiment measured disparity thresholds for plaids with strictly horizontal disparities and compared them to thresholds for each of the plaids' component gratings. For these stimuli, a horizontal shift of the carrier plaid pattern brings the left and right images into register. The plaid and its components necessarily had the same horizontal spatial disparity. The maximum horizontal extent of the plaid's periodic elements (formed by the intersection of one cycle of each grating) is limited by the horizontal period of the component grating whose orientation is closest to vertical, yet it generally differs from the *perpendicular* period of either of the components. The dependence of disparity threshold on these periodic extents in this and the following experiment should indicate whether detection is limited by phase disparity noise or spatial disparity noise and whether its source is in the plaids' 1-D components or their 2-D features.

The second experiment used plaids containing a grating whose disparity was always zero. This 'masking' grating is uninformative by itself for the task of detecting disparity but it determines the plaid's disparity in two-space: The plaid's disparity has a direction that is parallel to the mask's orientation and a spatial magnitude that varies with the relative orientations of the mask and the target grating, in accordance with Eqs. (1) and (2). Examples are shown in Fig. 2, where the target

grating has a fixed disparity direction and magnitude, while the plaid's disparity, indicated by the arrow, varies with the orientation of the zero-disparity mask. The plaid's disparity magnitude increases without bounds as the difference in component orientations approaches zero. If one assumes that observers are sensitive to the disparity of the plaid's 2-D spatial structure, then one would expect that, for a particular plaid disparity direction, thresholds should vary with the relative orientations of the two gratings so as to conserve either the plaid's spatial offset or its phase offset. One would also expect disparity thresholds to be elevated when the masking grating is oriented vertically, for in this case the disparity of the plaid is vertical. However, if one assumes that observers are sensitive only to component disparities, threshold should be independent of the plaid's parameters. In this case, the threshold for a plaid, expressed as the disparity of the target component, should be similar to the threshold for the target grating presented separately, rising only when the target grating approaches horizontal (see Fig. 1).

Results of these and associated experiments show that a pattern's 2-D spatial properties, in particular its disparity direction, do influence disparity threshold and that this influence of disparity direction cannot be reduced to effects of the patterns' 1-D components. However, disparity direction seems not to affect stereo matching. Rather, it affects the disparity threshold for combining independently stereo-matched 1-D components into 2-D feature disparities. Only those components whose disparity is at or above their own threshold for detection support the detection of 2-D pattern disparity.

2. Methods

Observers were shown a stimulus in each of two intervals. Only a difference in disparity distinguished the two stimuli in a task-relevant way. The observer's task was to discriminate the interval displaying the zero-disparity stimulus and the interval displaying the non-zero-disparity stimulus. No explicit judgment of depth was required.

2.1. Stimuli

The plaids consisted of two superimposed sinusoidal luminance modulations having the same spatial frequency, the same contrast, and different orientations. In the first experiment the two component gratings had the same horizontal spatial disparity—necessarily also the plaid's horizontal disparity—but differed in their disparity phase angle because of their orientation difference. In the second experiment only one component grating, the target, was presented with non-zero disparity; the other grating, the mask, was always presented

with zero disparity. In this case the disparity of the resulting plaid, given by Eqs. (1) and (2), has a direction parallel to the orientation of the mask grating and a magnitude equal to the target grating's disparity in this direction.

Of interest are effects of both the relative and the absolute orientations of the plaids' component gratings. These orientations took on various values, generally in steps of 15° , in both experiments. The orientations of the gratings were fixed within each 80-trial run. The data are typically plotted as a function of the orientation of one component, with the orientation of the other as a parameter.

Left and right half-images were displayed on the two sides of a luminance-calibrated monitor controlled by a Macintosh computer. Viewing was through a modified Wheatstone stereoscope at distances of 57, 74, or 93 cm, depending on the monitor used. Pixels extended 2 arcmin on a side. The gratings were visible within circular windows 7.8° in diameter displayed on a screen measuring approximately 21° (horizontal) \times 16.0° (vertical). The windows had a disparity of zero and were centered on black fixation squares, 6 arcmin on a side, which were continuously visible throughout the run of trials; the only non-zero disparities were interocular carrier phase shifts. The spatial frequency of the gratings was 1.0 c/d and their contrast was 0.10 for the data reported here. Mean screen luminance, which was also the mean luminance of each grating, was approximately 20 cd/m^2 .

2.2. Procedure

A trial consisted of two stimulus intervals, each 180 ms in duration and separated in time by 0.5 s. The stimuli were identical across the two intervals except for the disparity and absolute phase of the gratings. In one interval both gratings had zero disparity. In the other interval both gratings had the same positive horizontal spatial disparity (Experiment 1) or one grating (the target grating) had positive disparity and the other grating (the mask grating) had zero disparity (Experiment 2). The absolute phases of the gratings were randomized in every interval identically for the two eyes, thus translating the plaid unpredictably within the window between intervals and eliminating potential monocular position cues, without affecting disparity.

Though the target's disparity, when non-zero, had a positive horizontal component, observers did not necessarily see depth in Experiment 2 as "far". Superimposed sinusoids of similar frequency cohere in depth despite differences in component disparities. The resulting plaid can have a negative horizontal disparity component and "near" depth even if neither of its component gratings do (Farell, 1997, 1998; Farell & Li, 2002). The situation is diagrammed in Fig. 2, where the horizontal components of disparity for 1-D component

Polarity of Plaid Horizontal Disparities

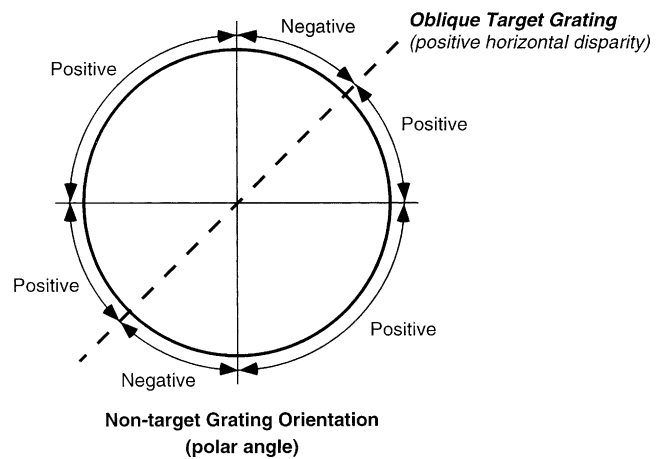


Fig. 3. Sign of horizontal disparity of a plaid pattern as a function of the orientation of a zero-disparity component grating. The plot's polar angle gives the orientation of the zero-disparity grating. The obliquely oriented target grating is shown as a dashed line. The target's disparity, measured perpendicularly to its orientation, has a horizontal component that is assumed to be positive. Then the horizontal disparity component of the resulting depth-coherent plaid will be either positive or a negative, as shown for the different ranges of mask orientation. Masks and targets with orientations on the same side of vertical, with the mask the more vertical of the two, combine to yield depth-reversed plaids, i.e., plaids appearing on the opposite side of the fixation plane from the target grating viewed separately. Plaids having disparities close to vertical typically appear partially diplopic and at the fixation plane or with ambiguous depth.

and 2-D pattern have the same signs in panel a and opposite signs in panel b. Such reversals occur in the second experiment when the mask grating is more vertical than the target grating and the orientations of both lie within the same quadrant, regions illustrated in Fig. 3. Observers were told in advance which conditions were likely to produce "near" plaid depth and which "far" plaid depth.

The task was to identify the interval displaying non-zero disparity. Observers were made aware of how the stimuli were constructed. In particular, for Experiment 1 they were told, in effect, that disparity signals were carried both by component gratings and by the plaid as a whole, and that a decision could in principle be based on any of these signals. Similarly, for Experiment 2 they were told that only one grating would have disparity, that the task was to detect the disparity of this grating, and that a decision may have to be based on properties of the plaid as a whole because of depth coherence. No decisions based on the perceived depth of the stimulus were required, though such decisions were generally what the observers reported making. Responses were made by clicking labeled buttons that appeared on-screen 0.5 s after the end of the second stimulus interval. A subsequent click initiated the following trial after a short delay. The importance of maintaining fixation throughout the trial was stressed.

Two threshold estimates were collected on each run of trials using randomly interleaved QUEST staircases (King-Smith, Grigsby, Vingrys, Benes, & Supowit, 1994; Watson & Pelli, 1983) with a criterion of 82% correct responses.

2.3. Observers

Four observers, three of them (including the author) highly experienced in stereo experiments, were run in the experiments. All had normal or corrected-to-normal acuity. Their stereo vision was also normal, yet there was nearly a fourfold range in their grating stereoacuties, due to the high thresholds of the least experienced observer. This difference did not appear to affect the results of interest, which were threshold elevations.

3. Results

3.1. Experiment 1: horizontal disparities

In the first experiment the plaid's disparity direction was horizontal. Thus the horizontal spatial disparities of the two component gratings were equal and the same as that of the plaid. The components' disparity phase angles differed, however, because of their orientation difference. Samples of these plaids appear in Fig. 4a and disparity thresholds for two observers are shown in Fig. 4b and c. Here, in order to measure the effect of component orientation, thresholds are given as the disparity phase angle of the more vertical of the two component gratings, designated the 'target', as a function of the orientation of the other, 'non-target', grating. Fig. 4b and c shows that target phase disparity at threshold is approximately 10° , varying little with the orientation of either the target or the non-target component grating. A similar pattern was found in the data of the other two observers.

The invariance of threshold across component orientations is shown in Fig. 5 in two ways for two different observers. Fig. 5a shows thresholds for the plaid's target component as a function of the target's orientation and the non-target's orientation. Single-grating thresholds are also shown for comparison. For a fixed target orientation (90° , the target orientation paired with the largest number of non-target orientations), the abscissa denotes non-target orientation; for a fixed non-target orientation (15°), it gives target orientation; and for single gratings, it gives their orientations. All three of these orientation measures yield flat threshold functions with similar absolute values. Fig. 5b shows thresholds for each target grating orientation (on the abscissa) averaged over all the non-target orientations; also shown are single-grating disparity thresholds. Again, threshold disparity phase angle for plaids is invariant

across non-horizontal component orientations, duplicating the invariance found with single gratings. In absolute value, too, thresholds for gratings in plaids are very similar to those for solitary gratings.

At threshold for detection, then, the disparity phase angle measured on the target grating shows little or no effect of the 2-D spatial properties of the plaid in which the grating is embedded; these 2-D properties vary with the relative orientations of the components, whereas threshold does not. Nor does threshold show an effect of the unambiguous disparity direction given by the plaid, nor, indeed, of the presence of the non-target grating that gives rise to the 2-D pattern. Instead, the threshold horizontal disparity for plaids conserves the threshold disparity phase angle of the more-vertical of the two component gratings. Phase disparity constancy gives a $1/\sin\theta$ relation between horizontal spatial disparity thresholds and target grating orientation, which the data approximate, as seen in Fig. 6. Thus, as in the case of gratings, there is no fixed spatial offset that must be exceeded for detection of disparity of plaids and, presumably, other 2-D patterns. Instead, threshold is limited by a disparity phase angle.²

However, the proper interpretation of this finding is unclear. If we measure the extent of the periodic elements of the plaid along any direction, we find that it equals the period of one of the plaid's components in this direction. (This component is the more vertical of the two under the conditions of this experiment, where both components have the same spatial frequency and the direction of concern is horizontal.) Therefore, a disparity along this direction will represent the same phase shift—i.e., same proportion of extent—for the plaid's 'blobs' as it does for the component. So, results showing disparity phase constancy tell us only that threshold disparity is limited by phase; they do not tell us whether the limiting interocular phase shift is that of the plaid's 1-D components or that of its 2-D features. We do know from Fig. 5 that the threshold phase shift for detecting disparity in plaids is numerically the same as that for detecting disparity in single gratings, though they are not constrained to be the same. This is consistent with a component analysis, but is not inconsistent with a feature analysis. We will return this issue when discussing the results of the next experiment.

² If threshold phase angle of the non-target grating were plotted in Figs. 4 and 5, the results would be very different. This is because the more vertical the grating of a particular spatial frequency, the greater the increment in disparity phase angle for a given increment in horizontal spatial disparity. The target grating, being more vertical in this experiment, always had a larger disparity phase angle than the non-target grating and therefore determined threshold. In other words, when the target disparity is at its threshold, the non-target disparity is below threshold.

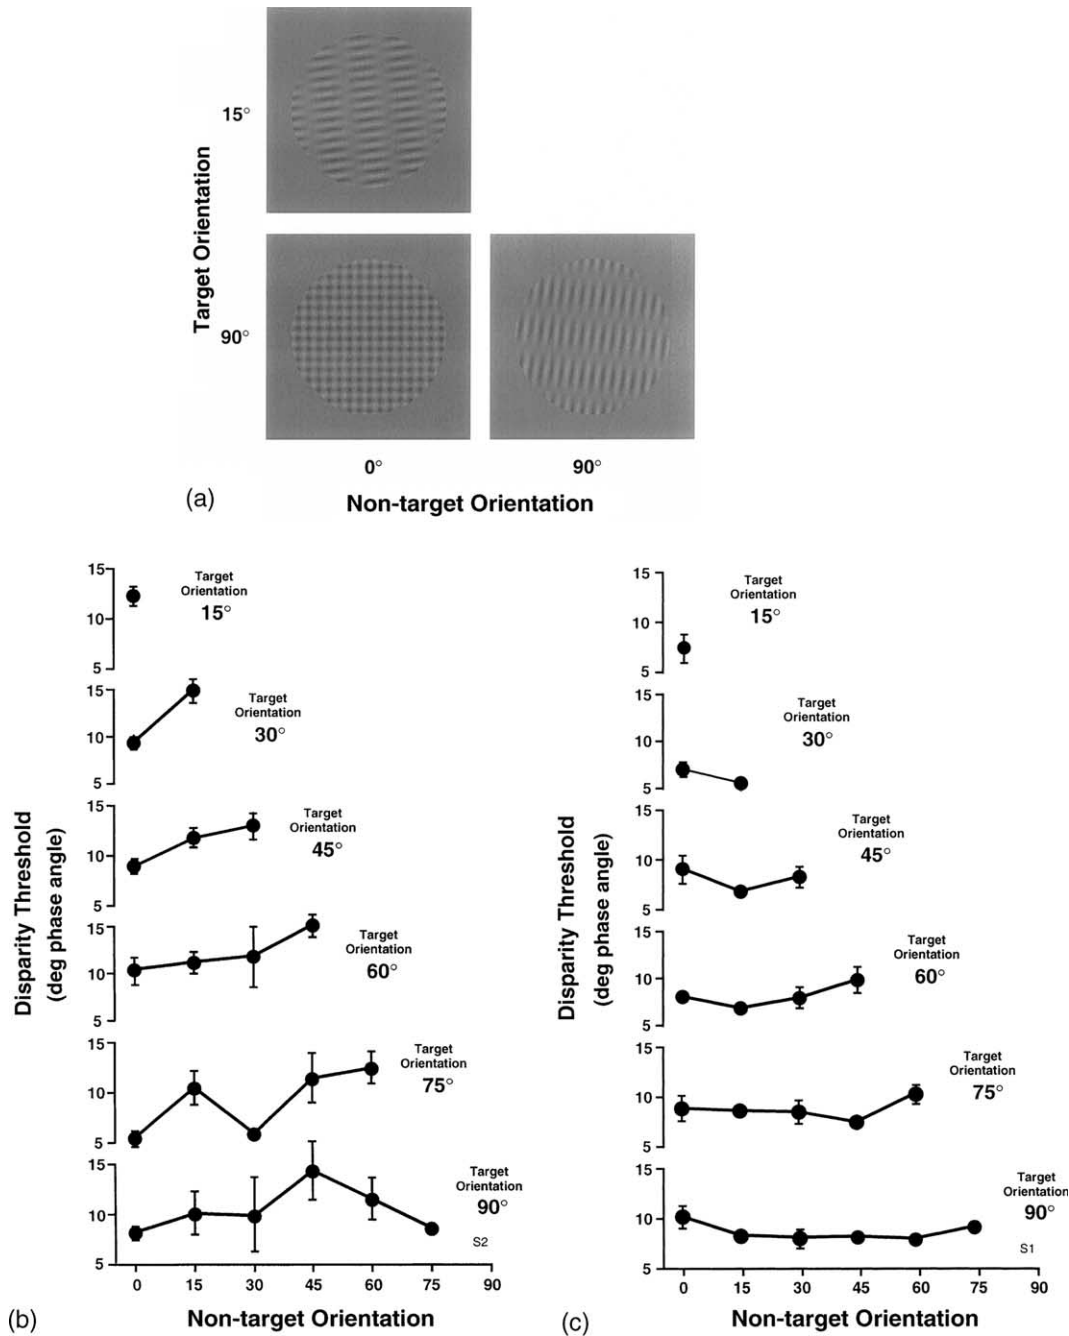


Fig. 4. (a) Plaids with component orientations $15^\circ + 0^\circ$, $90^\circ + 0^\circ$, and $90^\circ + 75^\circ$, corresponding to the corner data points in panels (b) and (c). Other component orientations produce plaids with intermediate intersection orientations and aspect ratios. For clarity the number of grating cycles within the circular window is twice the number used in the data reported here. The disparity direction for all plaids in the experiment was horizontal. (b, c) Threshold disparity for the plaids in Experiment 1, measured as a phase angle on the plaid's target (most-nearly vertical) component grating, for two observers. Thresholds are plotted on separate ordinates for each target orientation as a function of the orientation of the less-vertical of the gratings. Error bars are ± 1 s.e.m.

3.2. Experiment 2: varying 2-D disparity directions

Only one of the plaid's components, the target grating, had a variable disparity in Experiment 2. The masking grating, whose disparity was always zero, was uninformative as to the disparity of the target. Of interest is the mask's effect on threshold through interac-

tions with the target grating. The results are best seen in comparison with expectations based on two types of threshold constancy; these are given next.

Threshold disparities can be expressed in two ways: as phase disparity thresholds or as spatial disparity thresholds. It is immaterial whether thresholds of each type are measured on the plaid's 1-D components or on

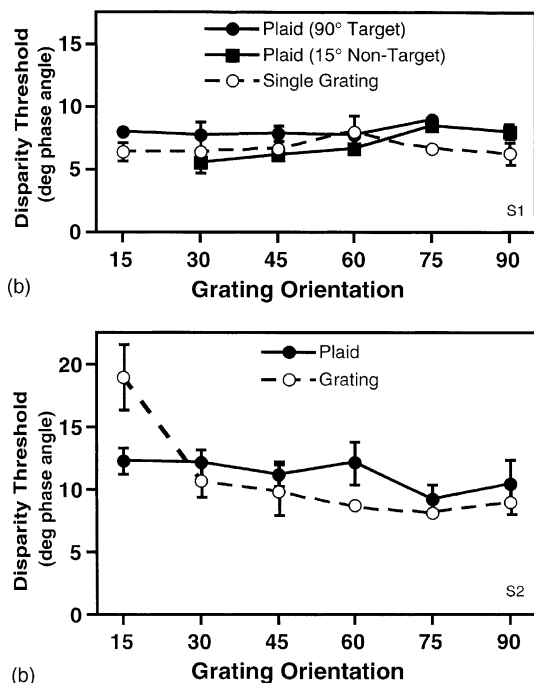


Fig. 5. Threshold phase disparities for horizontally offset plaids in Experiment 1. (a) Threshold for plaids with 90° targets and for plaids with 15° non-targets as a function, respectively, of the orientation of the plaids' non-target grating and of its target grating. Threshold phase angle for single gratings is also shown. Observer: S1. (b) Thresholds for plaids with target orientations given on the abscissa, averaged across the non-target orientations. Threshold phase angles are derived from those of Fig. 4. Single-grating thresholds are shown for comparison. Observer: S2. Error bars are ±1 s.e.m.

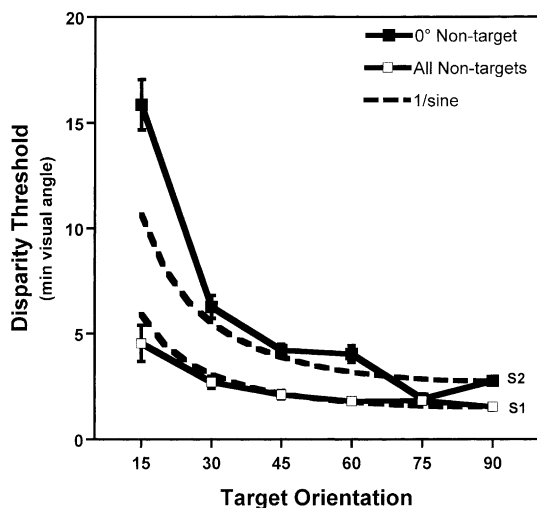


Fig. 6. Threshold spatial disparities of plaids in Experiment 1, as a function of target orientation. Data are plotted two ways. For observer S1, thresholds are averaged over all non-target orientations. For observer S2, thresholds are for 0° non-targets, the only non-target orientation paired with all target orientations. Comparison function shows reciprocal of sine of the target orientation, pinned to the threshold at 90°. Error bars are ±1 s.e.m.

its 2-D periodic elements; numerically these thresholds will be the same when they are measured in the same

direction (this direction, given by Eq. (1), is parallel to the orientation of the mask grating). Predicted thresholds are shown in Fig. 7 for a target grating with an orientation of 45°; other target orientations simply rotate the curves. Panel a shows predicted phase thresholds and panel b shows predicted spatial thresholds. Each panel gives thresholds expected under the assumption of threshold phase-constancy and under the assumption of threshold spatial constancy. In these polar plots, the radial distance from the origin gives threshold elevation (the ratio of thresholds for the plaid's target component and for a single grating with the same orientation) in accordance to Eq. (2); the polar angle gives the orientation of the masking grating (and the disparity direction of the plaid).

Suppose that in order to be detected, the plaid's disparity has to equal or exceed a fixed spatial extent, regardless of the disparity direction. The phase disparity of the target grating needed to give the plaid's disparity this criterion extent will depend on the relative orientations of the two component gratings. The required disparity will be greatest when the orientations are orthogonal and will shrink as the orientations converge. This is because a target grating with a small phase disparity creates a plaid with a large spatial disparity when combined with a similarly oriented mask grating. Predicted disparities appear as Fig. 8 threshold elevation curve when plotted on the phase disparity graph of Fig. 7a. No threshold elevation is predicted, only facilitation. A very different prediction arises if the threshold plaid disparity is a fixed proportion of the spatial extent of the plaid's periodic elements in the direction of disparity—a fixed phase offset. This is shown by the parallel lines in the spatial-disparity threshold elevation plot of Fig. 7b. The shape of the periodic elements change with the relative orientation of the components, growing ever longer as the orientations converge. To attain the plaid's threshold phase disparity, the spatial extent of the target grating's disparity must increase as the orientation of the mask grating approaches that of the target. In this case there is no threshold facilitation predicted for spatial disparities, only threshold elevation. The other functions—continuous circles with radius of 1—show plaid thresholds displaying phase disparity constancy (Fig. 7a) and spatial disparity constancy (Fig. 7b) when plotted in a phase metric and in a spatial metric, respectively.

Observed thresholds are shown as phase disparities in Fig. 8 and as spatial disparities in Fig. 9. Data points are plotted twice, reflected about the origin, with separate axes for normalized and absolute thresholds.³ The

³ Threshold phase angles for single 0° gratings are somewhat higher than those for gratings of other orientations (see Fig. 1). The absolute threshold axes in Fig. 9 apply only to the 30°, 45°, and 90° targets and are approximate, due to small variations in the thresholds for single gratings with these orientations.

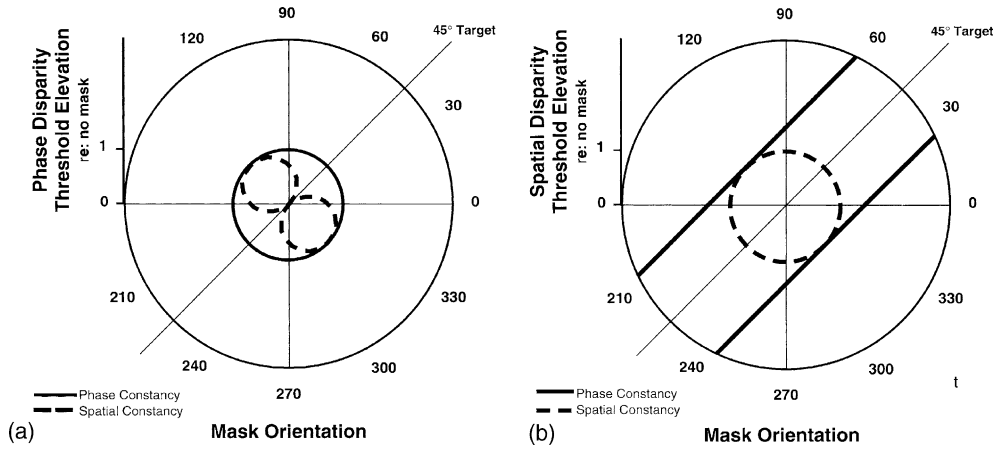


Fig. 7. Predicted threshold elevation for a target grating with an orientation of 45° as a function of the orientation of the zero-disparity masking grating. Phase thresholds (a) and spatial thresholds (b) are shown under the assumption of threshold phase-constancy and threshold spatial-constancy. Radial distance gives threshold elevation (with respect to single-grating thresholds) and polar angle gives mask orientation and plaid disparity direction. No threshold elevation is predicted for phase disparity (all radial values ≤ 1) under either constancy assumption and no threshold facilitation is predicted for spatial disparity (all radial values ≥ 1).

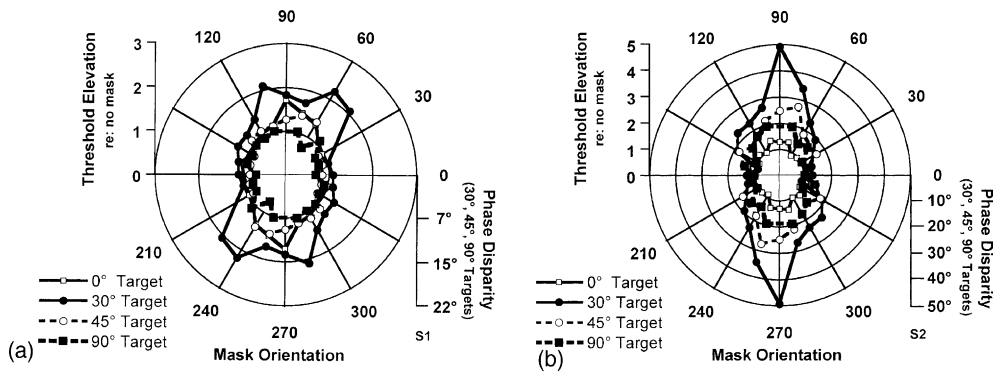


Fig. 8. Threshold phase disparities for Experiment 2, measured on the plaid's target grating. Polar angle gives the orientation of the zero-disparity mask grating and radial distance from center gives threshold. Separate curves are plotted for the four target orientations used in the experiment. Left scale gives threshold elevations, measured relative to thresholds for target gratings presented in the absence of a mask grating; a value of 1 indicates no elevation. The right scale gives the approximate phase shift at threshold of the 30°, 45°, and 90° targets. The phase shift for the 0° target requires multiplication of this scale by 2.8 for observer S1 (a) and by 2.3 for observer S2 (b). Mean phase s.e.m. for S1 is 0.63° and for S2, 1.72°.

various target orientations—0°, 30°, 45°, and 90°—are denoted by different symbols. The thresholds of Figs. 8 and 9 are for the two observers displaying the greatest and the least effect of mask orientation.

There are several variables affecting the thresholds for plaids and a fair amount of quantitative variation between the data of different observers. Behind this variation are two principle results that speak to the determinants of disparity detection thresholds: Thresholds for plaids are no smaller than thresholds for single gratings, and they increase as the orientation of the mask approaches the vertical. As discussed next, the first of these findings indicates that threshold is not determined by the spatial *magnitude* of the plaid's disparity; the second indicates that threshold is determined, at least in part, by the *direction* of the plaid's disparity.

3.2.1. No decreased in threshold phase disparity due to mask

If stereo threshold represents a constant spatial disparity, then threshold for a plaid, expressed as phase disparity, will decrease as the components approach one another in orientation (Fig. 7a). When measured as a phase angle on the target grating, this threshold should reach a maximum (threshold elevation of 1) when the component orientations are orthogonal and a minimum when component orientations are most similar, within 15° in our case. Fig. 8 shows instead that target grating's disparity at threshold is minimal when the mask orientation is around 0° (horizontal), regardless of the orientation of the target; threshold does not vary systematically with the components' relative orientations. This minimum equals the threshold when the target

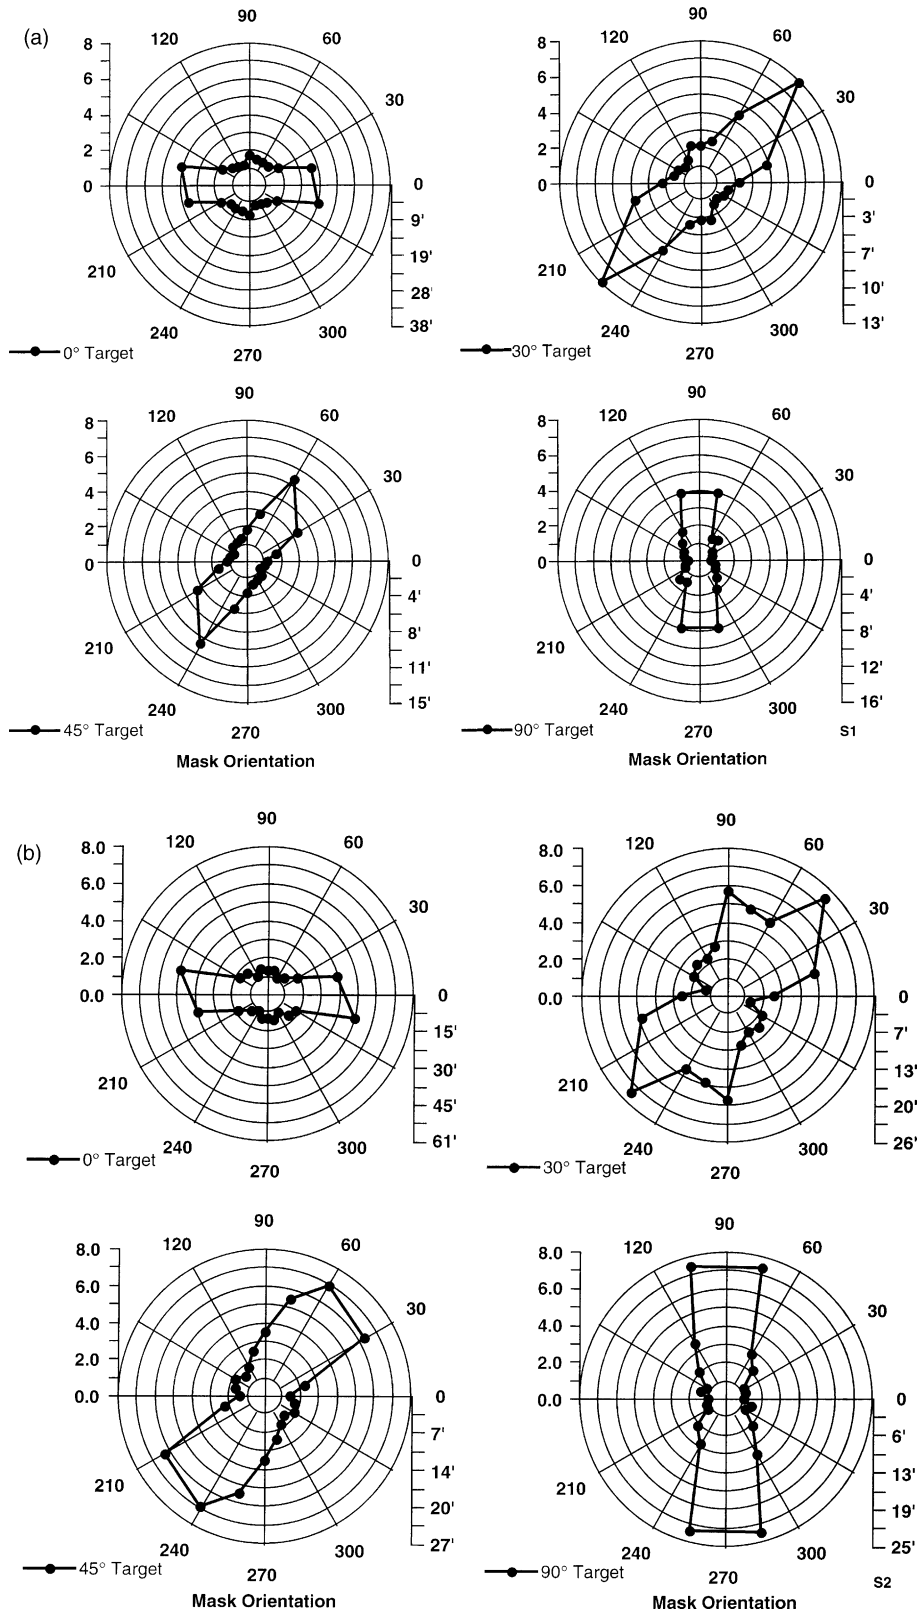


Fig. 9. Threshold plaid spatial disparities in Experiment 2, for the four target orientations. Thresholds are plotted as in Fig. 8, but separately for each target orientation. The left scale gives thresholds spatial disparities relative to the spatial offset of the threshold perpendicular disparity of the plaid's target grating, measured in the absence of a masking grating. The right scale gives spatial offset in visual angle. In all cases, the curve's direction of elongation is close to the orientation of the target grating. (a) Observer S1; (b) observer S2.

grating is presented alone, yielding a value of 1 on the threshold-elevation axis.

That the spatial extent of plaid disparities is not conserved at threshold is shown from another perspective by the approximate left–right and top–bottom symmetry of the curves of Fig. 8. The symmetry indicates that threshold is similar whether the mask grating is oriented at θ° or at $-\theta^\circ$. The target grating, if oblique, must be closer in orientation either to θ° or to $-\theta^\circ$. The closer the orientations, the greater the plaid’s spatial disparity. For targets with oblique orientations (30° and 45°), the plaid’s spatial disparity is greater for mask orientations between 0° and 90° than between 0° and -90° (as seen in Fig. 2a and c). The symmetry of Fig. 8 shows this difference in the disparity magnitude has little or no influence on threshold.

In Fig. 9 plaid thresholds are plotted in spatial units. A constant spatial disparity would plot as a circle. There is no such constancy at threshold. Thresholds that vary with the deviation of the disparity direction from the horizontal would plot with left–right symmetry. For 0° and 90° targets, the threshold functions are bilaterally symmetrical, but only because the target gratings are themselves bilaterally symmetrical; for 30° and 45° , the functions are very asymmetrical. The commonality across all target orientations is that spatial thresholds increase dramatically as component orientations become more similar. This is expected from thresholds that are limited by a constant phase disparity, as seen in the theoretical functions of Fig. 7b. Assuming, as a first approximation, a constant disparity phase angle at threshold, the corresponding spatial disparities correlate reasonably well with the observed values of Fig. 9: 0.97, 0.83, 0.91, and 0.99 for observer S1 and 0.96, 0.64, 0.90, and 0.99 for observer S2 for the 0° , 30° , 45° , and 90° targets [all p ’s < 0.05 under F distribution (1 and 9 df)]. Reflected in these values is the fact that plaid disparity direction generally has a smaller effect on component phase thresholds (Fig. 8) than relative component orientation has on pattern spatial thresholds (Fig. 9).

3.2.2. Threshold phase angle increases as disparity direction approaches vertical

At threshold, the disparity phase-shift generally increases as the orientation of the mask grating—and the direction of the plaid’s disparity—approaches vertical (Fig. 8). This effect varies systematically across the non-horizontal target orientations, being greatest for 30° and least for 90° . Points plotting as approximately parallel vertical lines on the polar coordinates yield a $1/\cos \theta$ threshold elevation function, where θ represents mask orientation. The effect is noticeable in the plaids’ threshold spatial disparities (Fig. 9) as a modest increment peaking around 90° in the polar plots, dwarfed by the effect of the relative orientations of the components.

What causes the vertical elongation of the functions of Fig. 8? This anisotropy is a deviation from phase constancy. It is not easily attributable to variations in component disparities, for there are none: Rotating the mask grating to a more vertical orientation produces no change in either the direction or the magnitude of component disparities, since the disparity of the mask is zero at all orientations. It is also difficult to attribute the shape of these functions to mask–target interactions outside the disparity domain—contrast masking or orientation attraction or repulsion, for example. This can be appreciated by referring again to the symmetry of the functions of Fig. 8. Target thresholds increase as the mask orientation approaches vertical, with the increase proceeding at approximately the same rate regardless of whether the change in mask orientation widens or narrows the difference in orientation between the components.

To verify that the vertical elongation of the disparity threshold functions is due to the plaid’s disparity direction, rather than to a property of the components, the experiment was modified, as sketched in Fig. 10, to decouple the plaid’s disparity direction from the absolute value of the gratings’ orientations. The component gratings were oriented symmetrically about the plaids’ disparity direction and for each disparity direction between horizontal and vertical, the relative orientations of the components took on values over nearly the entire permissible range, between $\pm 15^\circ$ and $\pm 75^\circ$ of the disparity direction. Plaids containing horizontal gratings were deleted from the series. The gratings’ spatial frequencies were equal; their phase disparity were equal in

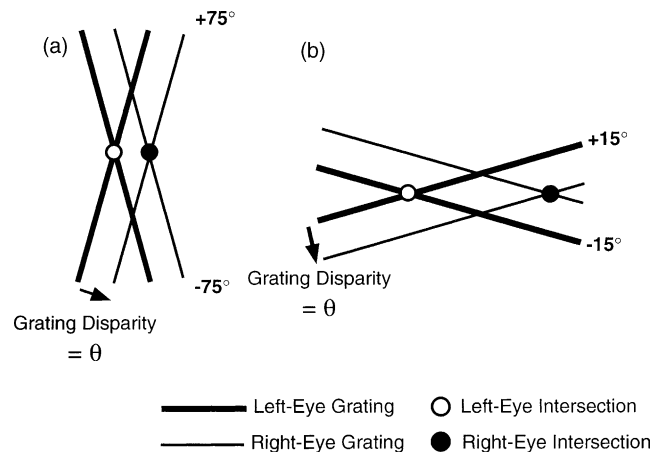


Fig. 10. Schematic rendering of symmetrical plaids. The component gratings are represented by pairs of lines with orientations symmetrically arrayed at angles of $\pm 15^\circ$ (a) and $\pm 75^\circ$ (b) of the plaid’s disparity direction, which here is horizontal. The components’ disparities (indicated by arrows) are equal in magnitude for the $\pm 15^\circ$ and $\pm 75^\circ$ cases, but the plaid’s disparity magnitudes, represented by the distance between left- and right-eye grating intersections (disks) differ by a factor of 3.7, given by the ratio of the sines of the angular differences between grating orientations and plaid disparity directions.

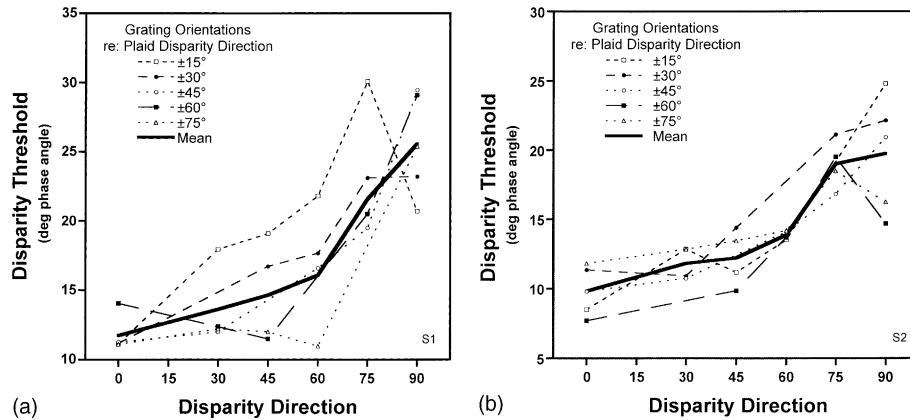


Fig. 11. Phase disparity thresholds for symmetrical plaids as a function of disparity direction. Data are shown for plaids with various component orientations and for their means (thick line). Both observers show a factor-of-two threshold difference between horizontal and vertical disparities. They show differing interactions between component orientation and disparity direction; for observer S1 (a), threshold elevation extend over a greater range of disparity directions the more similar the gratings orientations are to this direction, whereas observer S2 (b) shows no such relation. No effect of component orientation is seen when the disparity direction is horizontal (0° on abscissa).

magnitude and opposite in polarity. The result is a plaid disparity direction that bisects the angle between the gratings. The disparity direction and the gratings' orientations were fixed during a run in which two observers discriminated zero- and non-zero-disparity intervals. Other experimental details followed the pattern of the previous experiments.

Fig. 11 plots threshold disparity phase angles, which were equal for the two gratings, as a function of the plaid's disparity direction. Mean thresholds rise as the plaid's disparity direction veers away from horizontal, much as in Fig. 8, doubling upon reaching the vertical. Fluctuating about the means are data for the various grating orientation differences, which ranged from 30° to 150° ; this is a unsystematic fluctuation for one observer, but for the other observer (S1) thresholds begin to rise at disparity directions closer to horizontal the nearer the component orientations are to the disparity direction. Despite these differences between them, neither observer shows a sensitivity to the plaid's spatial disparity. For a given component disparity, the spatial disparity of a plaid with component orientations separated by 30° will exceed by a factor of 3.7 that of a plaid with component orientations separated by 150° . Yet the 30° separation did not give rise to lower thresholds, or the 150° separation to higher thresholds, than other conditions. Plaid disparity direction did have an effect that was independent of component orientation; the threshold increase as disparities approached the vertical was comparable in size to that observed in Experiment 2 (Fig. 8), where only one of the two components had a non-zero disparity.

It is well known that stereo thresholds for a variety of stimulus types rise when the disparity direction approaches vertical (Friedman, Kaye, & Richards, 1978; Morgan & Castet, 1997; Nielsen & Poggio, 1984; Ste-

venson & Schor, 1997; Westheimer, 1984). Figs. 8 and 11 show such an effect for plaids. Together these figures also show that the threshold elevation for plaids with vertical disparity is independent of the orientations of the components. Thus, the disparity thresholds for plaids are determined in part by a property of the plaid's 2-D structure—its disparity direction—that is not reducible to the properties of individual 1-D components. We turn now once more to the question of the other disparity parameter that determines threshold, the disparity magnitude. We have seen that the threshold limiting magnitude is a disparity phase angle, but we have not determined whether it is the 1-D component's phase angle or the 2-D pattern's phase angle.

3.3. Disparity phase of components vs. disparity phase of 2-D patterns

We can reject the spatial magnitude of a 2-D pattern's disparity as the factor limiting disparity thresholds. The data show instead that disparity detection is phase limited, with the threshold phase angle varying according to the pattern's direction of disparity. However, these results leave a central uncertainty unresolved. If we measure the extent of a plaid's periodic elements along any direction across its surface, we find that it is the same as the period of one of the plaid's components in this direction. Therefore, a disparity that is a certain proportion of the period of the 2-D features will be an equal proportion of the period of the component—an equality of phase angles. Returning to an example from Fig. 11, the threshold phase offset for horizontal disparities (0° on the abscissa) is approximately the same whether the components are 30° apart in orientation or 150° apart. But while the plaid in the 30° case has a threshold spatial disparity that is about 3.7 times larger than in the 150°

case (Fig. 10), both disparities are the same proportion of the plaids' 'blobs', which also differ by this same factor of 3.7. Thus, results showing phase-disparity constancy tell us only that threshold disparity is limited by phase. They do not tell us whether the limiting interocular phase shift is that of the plaid's 1-D components or that of the plaid's 2-D features.

The coupling between component and pattern phase offset rules out most options for experimentally dissociating them. However, two strategies allow some progress to be made on this problem. The first strategy is to enlist the aid of probability summation between the component gratings. Threshold facilitation through probability summation depends on separate opportunities for detection, as might occur if detection were component-based. Little probability summation would be expected in Experiment 1 because when the more-vertical of the two components is near its disparity threshold the other component is below threshold. This is a consequence of the orientation difference between the components of horizontally offset plaids: The more vertical the orientation, the greater the binocular phase shift. Likewise, Experiment 2 provides no opportunity for probability summation because only one component has a non-zero disparity. But if both components have identical phase disparities, they will approach threshold more or less synchronously. If detection is component-based, then threshold should drop because of probability summation between independent component detection processes. By contrast, the signal for pattern-based disparity detection is an increment in the plaid's disparity, not an increment in the components' disparities or the number of components with disparity increments. Thus, pattern-based detection, without access to the component disparities, should show no such benefit from probability summation.

We already have the data with which to put this issue to the test. One of our measures of the effect of disparity direction, that shown in Fig. 11, used plaids whose components had the same phase disparity. Threshold disparity tended to increase as the plaid's disparity direction approached vertical. And in Experiment 2, where only one component had non-zero disparity, thresholds also increased as the disparity direction neared vertical (Fig. 8). Fig. 12 compares these two sets of thresholds, normalized by single-grating data. The mean rate of threshold elevation as a function of disparity direction is comparable for the two conditions, as shown by the approximately parallel fits through the data points. However, when both components have the same phase disparity thresholds are lower than when one of the two has a fixed, and uninformative, disparity. The reduction in thresholds, about 1/6 log unit, is close to the value expected from probability summation based on the psychometric function slope of about 2.2 for single-grating disparity detection. If detection were not

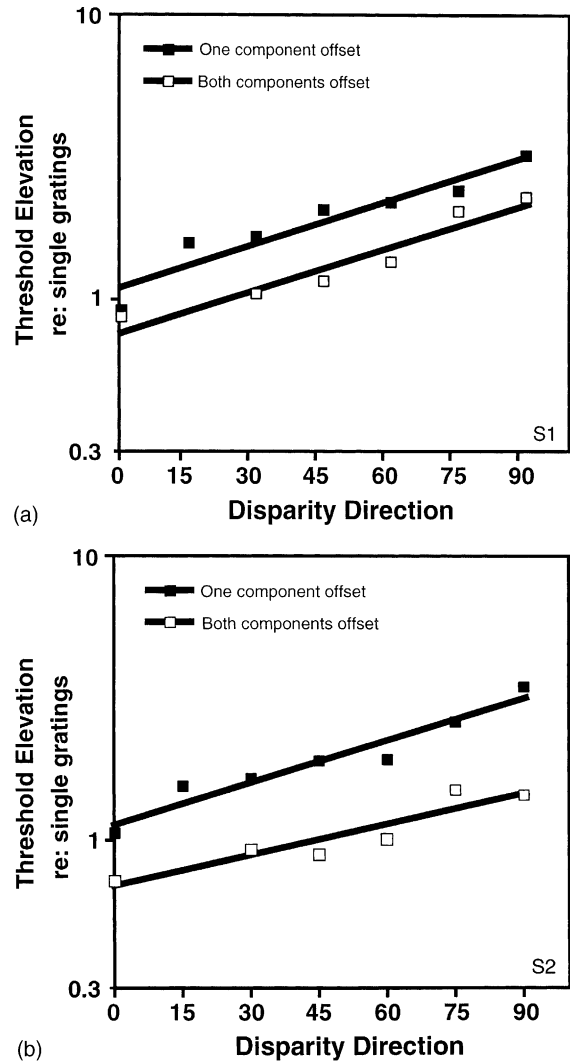


Fig. 12. Disparity threshold elevation for plaids with respect to single gratings as a function of disparity direction. The disparity of the plaid was carried by either one component (data from Fig. 8) or both components (data from Fig. 11): If by one, the other component had zero disparity; if by two, both components had equal disparities. (a) Observer S1; (b) observer S2.

based on component disparities, but rather on the composite plaid's disparity, no threshold reduction would be expected.

A second strategy for deciding between component- and pattern-disparity limits on threshold is to exploit threshold differences between single gratings. Horizontal gratings have higher detection thresholds than other gratings (Fig. 1). Plaids having one non-zero-disparity component grating can provide a test of whether the horizontal vs. non-horizontal orientation of this grating imposes different thresholds on plaids. Thresholds limited by the plaid's disparity should vary only with plaid disparity direction, not with target grating orientation. However, thresholds limited by the component's disparity should be elevated when the target grating is horizontal.

Again, we already have the data to make this comparison. Disparity detection thresholds for plaids in which only one component had non-zero disparity appeared in Fig. 8. When this target grating was horizontal, threshold elevations overlapped the values observed when target gratings were non-horizontal (30° , 45° , and 90°). However, the single-grating thresholds by which these data are normalized are higher for horizontal gratings than for the others, by roughly a factor of 2 for all observers (see Footnote 3). Fig. 13 shows unnormalized thresholds. Plotted here are data for the subset of plaids from Experiment 2 that had both a horizontal and a non-horizontal component—plaids with orientations 0° and 30° , 0° and 45° , and 0° and 90° . Thresholds measured when the 0° component functioned as the target were compared to those measured when it functioned as the mask, a comparison that controls for effects of relative component orientations and resulting changes in the plaids' spatial structure.

Fig. 13 shows that thresholds are about twice as high, or higher, when the horizontal component was the target than when the non-horizontal component was the target. The effect of the plaid's disparity direction can be seen in the thresholds for the 'horizontal target' condition, which of course are higher for vertical disparities (90° on abscissa) than for other disparity directions (30° and 45°). The disparity direction is horizontal for all the points in the 'non-horizontal target' condition, but this does not account for the 2:1 threshold ratio between the two conditions, because there is little threshold elevation for plaid disparity directions within $\pm 30^\circ$ or more of horizontal (Fig. 8). The 2:1 ratio seen within this range here (best measured at 30° on the abscissa) is thus the result of differences due to horizontal vs. non-horizontal target gratings. Thresholds for single horizontal and non-horizontal gratings, also appearing in the figure, show the 2:1 ratio and are similar in value to the thresholds for plaids. From this result, and the effect of probability summation (Fig. 12), we can locate the source of the threshold-limiting disparity of plaids in their 1-D components.

3.4. Effect of disparity direction: grating vs. plaid thresholds

We find that the threshold for detecting near-horizontal disparities of 2-D patterns is limited by disparity phase. The limiting phase angle for plaids is the same as that found for single gratings, suggesting a component limitation on pattern disparity thresholds. Additional evidence for a component analysis of 2-D pattern disparity comes from component probability summation and the effect of horizontal-vs.-non-horizontal component orientations.

Thresholds for plaids, however, cannot be reduced to thresholds for components. 2-D patterns with near-

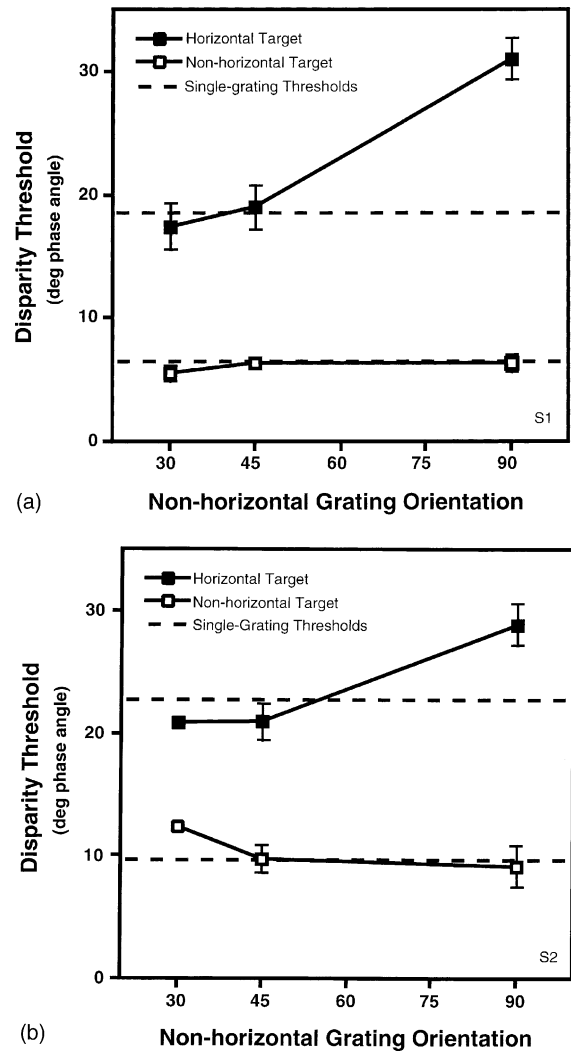


Fig. 13. Disparity thresholds for plaids containing a horizontal component. Threshold was measured either on the horizontal target grating (filled symbols) or the non-horizontal target grating (open symbols), whose orientation is given on the abscissa; the other component had zero disparity. Dashed lines give thresholds for single gratings, either horizontal (with the higher threshold) or non-horizontal, averaged over the three orientations. Data are taken from Fig. 8, where they appear as normalized thresholds. Disparity directions are horizontal for the non-horizontal target condition and parallel to the non-target orientation for the horizontal target condition. (a) Observer S1; (b) observer S2.

vertical disparities, whether plaids (Figs. 8 and 12) or other stimulus types (Friedman et al., 1978; Morgan & Castet, 1997; Nielsen & Poggio, 1984; Stevenson & Schor, 1997; Westheimer, 1984), have elevated thresholds. This is not due to elevated thresholds for the components of these patterns; the threshold elevation for plaids occurs whatever the orientation of the components (Fig. 11). Yet horizontal gratings, whose perpendicular disparities are vertical, also display elevated thresholds (Fig. 1). Therefore, disparity direction might exert an equivalent effect on grating and plaid thresholds

and might do so through a common mechanism. The data, however, suggest otherwise. Elevated threshold are found for plaid disparity directions far from vertical (Fig. 8), including directions that yield no elevation from gratings whose orientations (for example, 45°) are orthogonal to these directions (Fig. 5). So, plaids and gratings differ with respect to the range of disparity directions that give rise to elevations in threshold, with gratings showing a rather abrupt effect only within 15° or so of vertical and plaids showing a more gradual transition, with thresholds rising in some cases within 30° to 45° of horizontal.

These differences between results for plaids and gratings, including differences between observers in their plaid data that have no evident analog in their grating data, will frustrate attempts to reduce performance for 2-D patterns to processing of 1-D components. Pattern disparity directly influences threshold. Yet our data also show that pattern disparity thresholds are limited by the components. What this might mean for disparity processing models will be taken up in the Section 4.

4. Discussion

Grating disparity thresholds fall into two orientation-specific regimes. As seen in Fig. 1, thresholds for non-horizontal gratings are a constant disparity phase angle, at least for spatial frequencies in the low-to-moderate range (Schor et al., 1984). For horizontal gratings thresholds are considerably higher and, for some observers, display a near-spatial-constancy at disparity threshold, rather than a phase constancy. If we take the disparity direction for gratings to be perpendicular to their orientations, then these disparity thresholds for gratings are similar to disparity thresholds for plaids, but only superficially. When their disparities are near horizontal, plaids, like gratings, display a constant disparity phase angle at detection threshold (Figs. 4 and 5). When their disparities are non-horizontal—in particular, near-vertical—plaids have elevated thresholds (Fig. 8), again rather similar to the grating data. Parsimony would argue that these commonalities derive from the same mechanisms and, by extension, that the threshold criterion for stereo disparity is the same whether the stimulus is 1-D or 2-D. In fact, the plaid data argue otherwise and support the existence of separate thresholds for 1-D and 2-D stimulus elements. The proposed relation between 1-D and 2-D thresholds is this: Observers do not have direct access to the disparity of 1-D components, only to the depth of 2-D patterns; however, observers' thresholds are limited by the lowest of the component thresholds, because pattern depth is computed from component disparities.

Plaid threshold disparity phase angles were smallest when disparities were horizontal. This lower limit on

phase disparity thresholds for plaids equaled the phase disparity threshold for gratings. Thus differences in the spatial disparity gain between plaids—the change of horizontal spatial disparity for a given change in component phase disparity—had no effect on the threshold phase angle. No evidence was found in any of the experiments that it is possible to shape the plaid's 2-D spatial structure in a way that increases the observer's disparity sensitivity beyond the level set by the plaid's 1-D components; thresholds could not be driven below the just-detectable disparity phase angle measured on isolated gratings. van Ee et al. (2001) found related results using bar and letter stimuli.

These data argue for component-based detection, but they would be trivial if the visual system were acting linearly, independently processing the 1-D components, each with its own disparity and distinctly perceived depth. Superimposed squarewave gratings are generally perceived this way, segregating into transparent surfaces (Ahuja & Farell, 1997; Farell & Li, 2002), but superimposed sinewave gratings are not. Indeed, observers reported that the stimuli used here consistently appeared as depth-coherent plaids, never as transparent gratings. That the gratings were at some point processed as 2-D patterns, with disparities distinct from those of the component gratings, can be seen from the phenomenon of depth reversal (Fig. 3) and the effect of the plaids' disparity direction (e.g., Figs. 9 and 11).

The effect of disparity direction, in which threshold increases as the direction approaches vertical, is consistent with results of several recent experiments using 2-D stimuli. Morgan and Castet (1997) found that disparity thresholds for Gaussian blobs increased steadily as the disparity direction approached vertical. Stevenson and Schor (1997) found similar results for random-dot stereograms, though others using similar stimuli have found the tolerance to vertical disparity to be quite small (Nielsen & Poggio, 1984; Prazdny, 1985). Using dots, Westheimer (1984) measured thresholds that were far higher (~20×) for vertical disparity than for horizontal disparity.

Despite this consistency with earlier data, the finding that thresholds vary with the plaid's disparity direction presents interpretive challenges, for the data also show that threshold is not limited by the *magnitude* of the plaid's disparity in this direction, whether measured by a spatial or a phase metric. Rather, threshold is limited by the magnitude of the disparity of the plaid's 1-D components. We must ask, therefore, how component-limited threshold disparity is modulated by pattern-disparity direction. This surprising result, a combination of component and pattern constraints on threshold reminiscent of 1- and 2-D differences in perceived speed (Farell, 1999), can be accommodated by making three assumptions: pattern disparities are computed from component disparities, separate threshold nonlinearities

exist for pattern disparities and for component disparities, and perceptual decisions have access only to pattern disparities. These assumptions imply that single-grating disparities are detected through the same two-stage process as plaid disparities and their minimal threshold disparities should be the same, as observed in the data. An intriguing parallel is found in recognition of familiar objects—words—for which efficiency is limited by the independent detection of their component features, with no benefit coming from the processing of higher-order structure (Pelli, Farell, & Moore, 2003).

A 2-D pattern disparity computed on the output of stereo matches of 1-D components will depend jointly on the orientations, spatial frequencies, and disparities of the components. The results of combining these component parameters, not the individual component parameters themselves, are presumed to be accessible to the observer's decision process. Indeed, information about individual component disparities tells us little about the depth of the 2-D pattern (the depth-consistent and depth-reversed plaids of Figs. 2 and 3 serve as examples); the option may also exist to see no coherent 2-D pattern, but rather to see transparent 1-D patterns instead. As previously noted, changes in orientation and spatial frequency over fairly large ranges have only small effects on grating phase disparity thresholds. Yet because these changes can alter the disparity direction of the plaid (among other plaid parameters), they can have sizable effects on plaid disparity thresholds. Suppose the component disparities required for detecting the disparity of the plaid is specified by a criterion (i.e., a threshold nonlinearity) that varies with the plaid's disparity direction. For a particular direction θ the criterion may be low or high relative to the criterion for detecting the disparities of the individual components. If matching is carried out on components, and the disparities of the 2-D features are computed on the output of these matches, then a high criterion for the plaid disparity will result in a high threshold. However, a low criterion will be without effect; the threshold disparity of the components will have to be reached before the pattern's disparity can be determined. Therefore, plaids will not yield thresholds lower than the thresholds for the individual components (see Fig. 8). The threshold elevation seen as the plaid's disparity direction veers away from horizontal (Fig. 8) can be understood as a change in the threshold for combining component disparities into specific 2-D pattern disparity directions, independent of the thresholds of the individual components themselves. This conclusion takes into account the finding that threshold is not determined by the magnitude of this disparity. Following a priori environmental regularities, 1-D components are more readily combined into 2-D patterns with near-horizontal disparities, regardless of their own disparity direction (Farell & Li, 2002), with other directions leading to higher thresholds.

The observation that the threshold disparity for plaids varies with the plaids' disparity direction, independent of the components' properties, and yet is no lower than the components' threshold disparity, thus reflects decisions based on the disparity of the plaid's 2-D features. The component disparities from which plaid disparity is computed provide the threshold limiting noise to the decision process but they do not have access to the decision process. Nor, because of depth coherence, do they have access to perception. Decision and perception operate, perhaps inseparably, on the plaid.

In agreement with data presented here, van Ee et al. (2001) observed no threshold facilitation due to the disparities of occlusion boundaries, even when these disparities were large. Neither, however, did they observe threshold elevation, even when the direction of these disparities was near-vertical, in disagreement with the present data. This difference in results is understandable if, as proposed here, the higher thresholds for near-vertical plaid disparities reflects a second-stage threshold for combining component disparities into a 2-D pattern disparities. van Ee et al. used bars and letters, which are, like squarewave gratings and other spectrally broadband patterns (Farell, 1997, 1998; Farell & Li, 2002), quite resistant to depth coherence. So, their overlapping and non-overlapping conditions would not have differed in the need to combine component disparities; stimuli in both conditions may have been spatially separated, either laterally or in depth. With no difference in the component-combination stage, there would be no difference in threshold.

There has been a long tradition of thinking about stereo correspondence as a 1-D matching problem with spatial limits. This view invites the assumption that the stimuli and receptive fields pertinent to disparity coding are unoriented (e.g., Bishop, 1970; Julesz, 1971) or vertical (e.g., Howard, 1982; Read & Eagle, 2000) and is neutral-to-weak on the assumption of a correlation between spatial frequency and disparity coding. By comparison, component phase-disparity coding makes both orientation and spatial frequency intrinsically significant parameters for stereopsis and links them. Phase coding correlates low spatial frequencies with large spatial disparities, and high frequencies with small disparities (Felton, Richards, & Smith, 1972; Marr & Poggio, 1979; Richards & Kaye, 1974; Smallman & MacLeod, 1994). Orientation modulates this correlation; matching non-vertical components can extend the useful range of horizontal spatial offsets beyond the half-wavelength limit imposed by the phase-disparity coding of vertical components. Similarly, vertical receptive fields can signal offsets that have a greater spatial extent in nearly vertical directions. This link between component orientation and pattern disparity direction may provide a role for interactions between target and mask orientations, as seen in Experiment 2 (Fig. 9), which are otherwise

puzzling. However, the vertically oriented stimulus contours and receptive fields sometimes emphasized by position-coding (Howard, 1982; Read & Eagle, 2000) and even phase-coding treatments of disparity (DeAngelis, Ohzawa, & Freeman, 1991) do have a special role in a component analysis of stereopsis, for they provide the highest phase gain for horizontal disparities, thereby setting the stereoacuity limit—i.e., determining horizontal *spatial* thresholds.

Both phase and spatial metrics are used in stereopsis. Phase disparity limits the threshold response of stereo matches, while spatial disparities map onto perceived depth. It is the disparities of 1-D components—phase-based and normal to the component—that are the likely elementary signals from which spatial disparities and depth are computed. The results suggest that the horizontal disparities of the 2-D patterns typical of naturalistic viewing conditions should become detectable when the most-nearly vertical of the patterns' 1-D components reaches its disparity threshold. Patterns with only near-horizontal components will have greater horizontal spatial thresholds than others (because a constant threshold phase disparity corresponds to a horizontal spatial disparity that is proportional to the reciprocal of the sine of the component's orientation). But for all patterns the disparity phase of the 1-D components is the common limiting source of noise.

Acknowledgements

I thank Dr. Suzanne McKee for many enlightening discussions and helpful suggestions, Dr. Denis Pelli for raising pointed questions, Ms. Simone Li for rejecting pat answers, and two anonymous reviewers for useful comments. Portions of Experiments 1 and 2 were presented earlier in abstract form (Farell, 2000). This research reported here was supported by NEI Grant EY 12286.

References

- Adelson, E. H., & Movshon, J. A. (1982). Phenomenal coherence of moving visual patterns. *Nature*, *300*, 523–525.
- Adelson, E. H., & Movshon, J. A. (1984). Binocular disparity and the computation of two-dimensional motion. *Journal of the Optical Society of America A*, *1*, 1266.
- Ahuja, S., & Farell, B. (1997). Perceiving transparency in stereo plaids. *Investigative Ophthalmology and Visual Science*, *38*, 904.
- Anzai, A., Ohzawa, I., & Freeman, R. D. (1999). Neural mechanisms for encoding binocular disparity: receptive field position versus phase. *Journal of Neurophysiology*, *82*, 874–890.
- Badcock, D. R., & Schor, C. M. (1985). Depth-increment detection function for individual spatial channels. *Journal of the Optical Society of America A*, *2*, 1211–1215.
- Bishop, P. O. (1970). Beginning of form vision and binocular depth discrimination in cortex. In F. O. Schmitt (Ed.), *The neurosciences: second study program*. New York: Rockefeller University Press.
- Blake, R., Camisa, J., & Antoinetti, D. N. (1976). Binocular depth discrimination depends on orientation. *Perception and Psychophysics*, *20*, 113–118.
- Boothroyd, K., & Blake, R. (1984). Stereopsis from disparity of complex grating patterns. *Vision Research*, *24*, 1205–1222.
- Cumming, B. C. (2002). An unexpected specialization for horizontal disparity in primate visual cortex. *Nature*, *418*, 633–636.
- DeAngelis, 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.
- Ebenholtz, S., & Walchli, R. (1965). Stereoscopic thresholds as a function of head- and object-orientation. *Vision Research*, *5*, 455–461.
- Farell, B. (1997). Depth coherence and depth reversal in stereo plaids. *Investigative Ophthalmology and Visual Science*, *38*, 904.
- Farell, B. (1998). Two-dimensional matches from one-dimensional stimulus components in human stereopsis. *Nature*, *395*, 689–693.
- Farell, B. (1999). Color and luminance in the perception of 1- and 2-dimensional motion. *Vision Research*, *39*, 2633–2647.
- Farell, B. (2000). What determines disparity thresholds for 2-D patterns? *Investigative Ophthalmology and Visual Science*, *41*, 737.
- Farell, B., & Ahuja, S. (1996). Binocular disparity of 1- and 2-D contours. *Investigative Ophthalmology and Visual Science*, *37*, 284.
- Farell, B., & Li, S. (2002). Depth coherence and transparency. Submitted for publication.
- Farell, B., Li, S., & McKee, S. P. (2002). Coarse-to-fine interactions in disparity discriminations. Submitted for publication.
- Felton, T. B., Richards, W., & Smith, J. R. A. (1972). Disparity processing of spatial frequencies in man. *Journal of Physiology (London)*, *25*, 349–362.
- Friedman, R. B., Kaye, M. G., & Richards, W. (1978). Effect of vertical disparity upon stereoscopic depth. *Vision Research*, *18*, 351–352.
- Heckmann, T., & Schor, C. M. (1989). Is edge information for stereoacuity spatially channeled? *Vision Research*, *29*, 593–607.
- Howard, I. P. (1982). *Human visual orientation*. New York: Wiley and Sons.
- Julesz, B. (1971). *Foundations of cyclopean perception*. Chicago: University of Chicago Press.
- King-Smith, P. E., Grigsby, S. S., Vingrys, A. J., Benes, S. C., & Supowit, A. (1994). Efficiency and unbiased modifications of the QUEST threshold method: theory, simulations, experimental evaluation and practical implementation. *Vision Research*, *34*, 885–912.
- LeVay, S., & Voigt, T. (1988). Ocular dominance and disparity coding in cat visual cortex. *Visual Neuroscience*, *1*, 395–414.
- Levinson, E., & Blake, R. (1979). Stereopsis by harmonic analysis. *Vision Research*, *19*, 73–78.
- Marr, D., & Poggio, T. (1979). A computational theory of human stereo vision. *Proceedings of the Royal Society, London, B*, *204*, 301–328.
- Maske, R., Yamane, S., & Bishop, P. O. (1986). End-stopped cells and binocular depth discrimination in the striate cortex of cats. *Proceedings of the Royal Society, London, B*, *229*, 227–256.
- Mayhew, J. E. W., & Frisby, J. P. (1978). Contrast summation effects and stereopsis. *Perception*, *9*, 69–86.
- Morgan, M. J., & Castet, E. (1997). The aperture problem in stereopsis. *Vision Research*, *39*, 2737–2744.
- Nielsen, K. R. K., & Poggio, T. (1984). Vertical image registration in stereopsis. *Vision Research*, *24*, 1133–1140.
- Ogle, K. N. (1955). Stereopsis and vertical disparity. *Archives of Ophthalmology*, *53*, 495–504.
- Ohzawa, I., & Freeman, R. D. (1986). The binocular organization of simple cells in the cat's visual cortex. *Journal of Neurophysiology*, *56*, 221–242.

- Pelli, D. G., Farell, B., & Moore, D. C. (2003). The remarkable inefficiency of word recognition. *Nature*, in press.
- Prazdny, K. (1985). Vertical disparity tolerance in random-dot stereograms. *Bulletin of the Psychonomic Society*, 23, 413–414.
- Read, J. C. A., & Eagle, R. A. (2000). Reversed stereo depth and motion direction with anti-correlated stimuli. *Vision Research*, 40, 3345–3358.
- Richards, W., & Kaye, M. G. (1974). Local versus global stereopsis: two mechanisms? *Vision Research*, 14, 1345–1347.
- Rohaly, A. M., & Wilson, H. R. (1993). Nature of coarse-to-fine constraints on binocular fusion. *Journal of the Optical Society of America A*, 10, 2433–2441.
- 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. C., & Ogawa, J. (1984). Spatial tuning of static and dynamic local stereopsis. *Vision Research*, 24, 573–578.
- Smallman, H. S., & MacLeod, D. I. A. (1994). Size-disparity correlation in stereopsis at contrast threshold. *Journal of the Optical Society of America*, 11, 2169–2183.
- Smallman, H. S., & MacLeod, D. I. A. (1997). Spatial scale interactions in stereo sensitivity and the neural representation of binocular disparity. *Perception*, 6, 977–994.
- Stevenson, S. B., & Schor, C. M. (1997). Human stereo matching is not restricted to epipolar lines. *Vision Research*, 37, 2717–2723.
- van Ee, R., & Schor, C. M. (2000). Unconstrained stereoscopic matching of lines. *Vision Research*, 40, 151–162.
- van Ee, R., Anderson, B. L., & Farid, H. (2001). Occlusion junctions do not improve stereoacuity. *Spatial Vision*, 15, 45–59.
- Watson, A. B., & Pelli, D. G. (1983). QUEST: a Bayesian adaptive psychometric method. *Perception and Psychophysics*, 33, 113–120.
- Westheimer, G. (1984). Sensitivity for vertical retinal image differences. *Nature*, 307, 632–634.