

Orientation-Specific Computation in Stereoscopic Vision

Bart Farell

Institute for Sensory Research, Syracuse University, Syracuse, New York 13244-5290

The left and right eyes receive subtly different images from a visual scene. Binocular disparities of retinal image locations are correlated with variation in the depth of objects in the scene and make stereoscopic depth perception possible. Disparity stereoscopically specifies a stimulus; changing the stimulus in a way that conserves its disparity leaves the stimulus stereoscopically unchanged. Therefore, a person's ability to use stereo to see the depth separating any two objects should depend only on the disparities of the objects, which in turn depend on where the objects are, not what they are. However, I find that the disparity difference between two stimuli by itself predicts neither stereoacuity nor perceived depth. Human stereo vision is shown here to be most sensitive at detecting the relative depth of two gratings when they are parallel. Rotating one grating by as little as 10° lowers sensitivity. The rotation can make a perceptible depth separation invisible, although it changes neither the relative nor absolute disparities of the gratings, only their relative orientations. The effect of relative orientation is not confined to stimuli that, like gratings, vary along one dimension or to stimuli perceived to have a dominant orientation. Rather, it is the relative orientation of the one-dimensional components of stimuli, even broadband stimuli, that matters. This limit on stereoscopic depth perception appears to be intrinsic to the visual system's computation of disparity; by taking place within orientation bands, the computation renders the coding of disparity inseparable from the coding of orientation.

Key words: stereoscopic depth perception; relative disparity; orientation; stereoacuity; joint coding; binocular vision; psychophysics

Introduction

Stereo-enabled animals use binocular disparity for depth perception, but a disparity value does not in itself provide useful information about the depth of an object. This is because disparities change with the vergence angle of the eyes (the angle formed by the two lines of sight). In fact, humans are poor at detecting even large changes in disparity resulting from uniform full-field motion in depth, where the disparity of each stimulus changes by the same amount (Erkelens and Collewijn, 1985; Regan et al., 1986). What humans and other stereo-enabled animals are very sensitive to is relative disparity, the difference between the disparities of two stimuli that yields the perception of relative depth. For both humans (Westheimer, 1979) and monkeys (Prince et al., 2000), stereoacuity is approximately an order of magnitude better when the target stimulus appears in the presence of a nearby reference stimulus rather than appearing alone. Relative disparity is essentially immune to changes in eye position and is believed to be computed using a very simple algorithm, one that differences the absolute, or retinotopic, disparities of two stimuli (Westheimer, 1979).

Because absolute disparity is the only variable entering into this differencing algorithm, sensitivity to relative disparity should be invariant not only to changes in eye position but also to any change to the stimuli that conserves their absolute disparity signals. As Ogle (1962), referring to the contours within objects or at

their borders, has said, "the stimulus for stereoscopic vision is the disparity between the retinal images of those contours, and it is not concerned with the similarity or identification *per se* of the objects themselves." Invariance with respect to similarity or identification is an optimal strategy for computing relative disparity, for it allows a person to judge the depth separating any pair of objects, regardless of how the objects compare with respect to properties other than depth. I tested the invariance of human stereoacuity for stimuli differing in orientation. In the experiments reported here, observers viewed two stimuli, and each stimulus was seen with the same orientation by the right and left eyes. However, the two stimuli might differ in orientation from each other. The data show that even a small difference in the orientations of the two stimuli degrades stereoacuity. Relative orientation also affects perceived relative depth at disparities far beyond threshold. These effects are not caused by changes in physical disparity or stimulus visibility that might accompany changes in orientation. Therefore, whether the depth separating two objects can be resolved stereoscopically depends on the similarity of the objects in orientation and cannot be predicted from retinal disparities, or differences between retinal disparities, alone. The dependence of stereoscopic depth perception on relative orientation appears to be intrinsic to the visual system's computation of disparity.

Materials and Methods

Stimuli. Two basic kinds of stimuli were used, grating patches and noise patches. The center-surround grating stimuli consisted of a two-dimensional Gabor patch surrounded by a radial Gabor annulus. The Gaussian envelope of the central patch had horizontal and vertical SDs of 0.53° of visual angle. The windowing of the surround was Gaussian along the radial direction (i.e., away from the center of the display) with a SD of 0.34° (except for the data for S3 in Fig. 1B, where the Gaussian profile was

Received March 14, 2006; revised July 4, 2006; accepted July 27, 2006.

This work was supported by National Institutes of Health Grant EY012286. I thank Julian Fernandez, John Krauskopf, Suzanne McKee, and Denis Pelli for helpful discussions. I also thank two anonymous reviewers for thoughtful comments and valuable suggestions, including the check on the generality of the results to nonperiodic stimuli.

Correspondence should be addressed to Bart Farell at the above address. E-mail: bfarell@syr.edu.

DOI:10.1523/JNEUROSCI.1100-06.2006

Copyright © 2006 Society for Neuroscience 0270-6474/06/269098-09\$15.00/0

asymmetrical). Gaussians were truncated at $\pm 2\sqrt{2}\sigma$. The peaks of the center and surround Gaussian envelopes were separated by a distance of 2° (but 1.5° for observer S3 and 2.3° for S5) (see Fig. 1*B*); increasing the separation to 3° produced similar data for some observers and raised thresholds substantially for other observers. (Whether this increase was attributable to the greater eccentricity of the surround or to the separation per se has not yet been determined.) At the contrasts used, the center grating and the surround grating did not visibly overlap at any separation. Spatial frequencies of the sinusoidal carriers were three cycles/degree (c/d) in most experiments; 1 and 2 c/d patterns were also examined in several experiments and produced similar threshold data to those reported here.

Filtered noise patches (see Fig. 2) also appeared in center-surround arrangement. The noise had the same windowing as the gratings, with peak contrasts of the center and surround separated by 2° . Independent noise samples (prefiltered Gaussian luminance distribution, two-pixel checks) were used for center and surround and for each trial. Filtering confined the noise spatial frequency content to four octaves (0.5–8 c/d) and the orientation content to either ± 15 or $\pm 30^\circ$ of the mean orientation, which was either 60 or 120° . Filter gain was constant across the passband. Center and surround noise had the same orientation bandwidth, and stimuli with broad and narrow orientation bandwidths had the same expected noise rms contrast. Noise disparity was controlled to subpixel accuracy by manipulating the phase spectrum such that all component phase disparities corresponded to the same spatial disparity; this disparity, that of the center noise carrier as a whole, was horizontal in direction.

Stimuli were presented on luminance-calibrated monitors driven by a Macintosh G5 computer via attenuators that provided monochromatic resolution of ~ 12 bits (Pelli and Zhang, 1991). Michelson contrast of all gratings was 0.1. Increasing the contrast of the surround by a factor of two had no significant effect on stereo performance for any of the observers tested in this way, thereby allaying concern about surround-inducing changes (which could be relative-orientation contingent) in the perceived contrast of the center grating. (Direct measures of contrast sensitivity to the center as a function of surround orientation turned up only modest effects that were inconsistent across observers, with most observers showing no effect.) Presentation duration was 150 ms, with abrupt onsets and offsets for both grating and noise stimuli. Presenting the gratings in a 500 ms cosine temporal window lowered stereo thresholds slightly but preserved the pattern of results, as did drifting the gratings at 1 Hz for 500 ms, with the horizontal component of motion having opposite directions in the center and the surround. Stimuli generation and presentation were controlled by a MatLab (MathWorks, Natick, MA) program incorporating elements of Psychophysical Toolbox software (Brainard, 1997; Pelli, 1997).

Stimuli were centered on two 21 inch monitors, one for each eye, except in the case of the data for S3 in Figure 1*B*, for which the stimuli were presented on the left and right sides of a single 19 inch monitor. Monitor refresh rate was 75 Hz. Stimuli were viewed through a front-silvered mirror stereoscope. The optical distance was 1.26 m (0.93 m for observer S3; see Fig. 1*B*). Observers initiated each trial with a click of the computer mouse after fixating between aligned nonius lines. The nonius lines vanished 0.125 s before stimulus onset. No fixation stimulus was provided apart from the center-surround Gabor patterns, which were presented on an otherwise blank screen. The contours nearest these patterns were the edges of the monitor screens, which were located $\sim 3.8^\circ$ vertically and 5.8° horizontally from the outer visible limit of the annulus (and somewhat less than this for the data of S3; see Fig. 1*B*). The background luminance was 21 cd/m², which was also the mean luminance of the patterns. Observers used a chin rest and viewed the displays with natural pupils in a moderately lit room.

For threshold measurements, the surrounding grating or noise patch had zero disparity relative to nonius alignment. The carrier of the central Gabor patch had a disparity that varied from trial to trial according to an adaptive staircase method. No difference in threshold was noted if instead it was the carrier of the surrounding patch that varied in disparity; indeed, at near-threshold disparities it is not possible to tell whether it was the center or the surround whose disparity varied from trial to trial.

To obtain psychometric functions for perceived depth (see Fig. 6), the surrounding stimulus was a “plaid” annulus composed of two superimposed sinusoidal components. One component had a fixed disparity of either 1.67 or 2.22 arcmin (a phase angle of 30 or 40°). The other component had zero disparity. This gave the plaid carrier a disparity of $+45^\circ$ (when the component oriented at 135° had positive disparity) or -45° (when the component oriented at 45° had positive disparity). The disparity of the central grating (oriented at either 45° or 135°) varied over a large range, taking on both positive and negative values across trials under the control of a constant-stimulus procedure.

In all cases, the disparity of the Gaussian contrast envelopes was fixed at zero; the only nonzero disparities were interocular phase shifts of the grating or noise carrier. The absolute phases of the gratings were independently randomized (identically in the left and right eyes) on each trial; this translates the gratings within their envelopes unpredictably, disrupting potential local position cues without affecting disparity. Noise samples were independent both across trials and between the center and the surround.

Procedure. The observer’s task was to judge the center grating or noise as being near or far (i.e., closer to the observer or more distant) relative to the surrounding grating, plaid, or noise. After the observer initiated a trial with a click of the mouse, a brief “bing” sound introduced the stimulus presentation; observers signaled their decisions by clicking “Near” and “Far” buttons that appeared on the screen after termination of the stimulus. A correct response during threshold measurements was answered with a rewarding “clink”; no feedback was provided during perceived-depth measurements, where the correctness of responses is undefined.

Observers were given no information about which of the two stimuli had nonzero disparity. This applied to the center and surround stimuli and to the individual components of plaids. For gratings and noise near the threshold for depth discrimination and for plaid components at all disparities, this information is not assessable perceptually.

Before data collection, observers were given sufficient practice to stabilize performance. Data were gathered in runs of 40–48 trials for threshold measurements and 60 trials for perceived-depth measurements, each preceded by six to eight warm-up trials. At least four runs contributed to each data point. Thresholds were measured using the QUEST algorithm (Watson and Pelli, 1983), which guided trial-to-trial variation in disparity toward the disparity value required for the 84% correct discrimination level ($d' = 1$). Psychometric values for perceived depth were fit by Weibull functions.

Observers. Data were recorded from four observers, except in the grating threshold measures of Figure 1*B* (six observers; data for three shown) and in the noise threshold (see Fig. 2) and the perceived depth studies (see Fig. 6), for which two observers participated. Two observers, including the author, were highly experienced in psychophysical tasks, and four observers had exposure only to these and related experiments. The inexperienced observers were naive as to the purposes of the experiments. All of the observers had normal acuity (in some cases with corrective optics) and normal stereo vision. All observers gave their informed written consent before participating in the experiments; the protocol for the experiments was reviewed and approved by the Institutional Review Board of Syracuse University.

Results

The stimuli in the first experiment were patches of sinusoidal gratings arranged in a center-surround configuration (Fig. 1*A*). Threshold was defined as the disparity that enabled observers to discriminate, with an accuracy of 84%, between central gratings that were near (negative disparity) and those that were far (positive disparity) with respect to the zero-disparity surround. Threshold disparity of the center grating was expressed as a visual angle (that is, the subtense of the interocular spatial shift) and a phase angle (that is, the interocular carrier phase shift). I used various spatial separations between the center and the surround, finding little effect of separations below 2.3 – 3° of visual angle, provided the gratings did not visibly overlap (which they did not

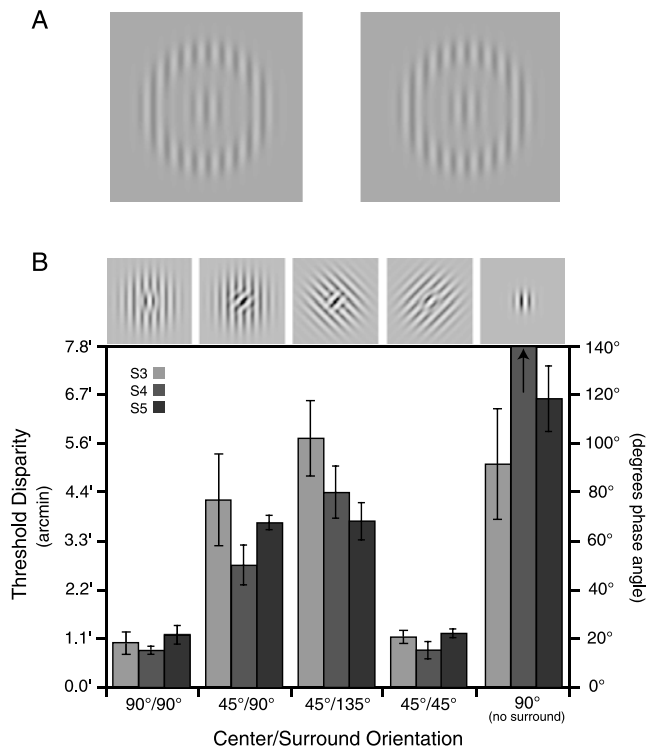


Figure 1. Stereoacuity for center and surround gratings as a function of relative grating orientation. *A*, Binocular center-surround Gabor stimuli. Fusing the left and right patterns by converging or diverging the eyes brings the gratings into stereoscopic view. The absolute disparity of the surrounding grating patch was set nominally to zero in the experiment; that of the center grating carrier was either positive or negative on any experimental trial. The observer's task was to judge the center as near or far with respect to the surround. The spacing between center and surround contrast envelopes had no major effect on stereoacuity as long as it was moderate in size and avoided overlapping the visible portions of the carriers; both of these conditions applied in experiments of this study. Grating contrast is higher here than in experiments. *B*, Effect of large orientation center-surround differences on stereoacuity for three observers. Thresholds for discriminating the polarity of the depth difference between center and surround are shown in spatial (arcmin) and phase (degrees) disparity units. Orientation differences of 45 or 90° elevated threshold substantially compared with same-orientation conditions. Stereo thresholds for the no-surround condition were not measurable for observer S3. The separation between center and surround stimuli differed between observers, with observer S3 viewing the tightly nestled stimuli shown at the top of the graph and observers S4 and S5 viewing stimuli similar to those shown in *A*. This difference had no obvious effect on stereoacuity. Stimulus displays at the top of the graph appear at higher contrast than those used in the experiment. Contrast thresholds for detecting the center grating as a function of the orientation of the surround (data not shown) were flat for most observers and correlated poorly with relative disparity thresholds. Error bars in this and subsequent figures show ± 1 SEM of threshold estimates.

do in these experiments). Figure 1*B* shows the basic result. When both gratings were oriented vertically, threshold disparity, measured perpendicular to the grating orientation, averaged just over 1 arcmin for these observers, equivalent to a phase disparity of $\sim 19^\circ$ or $\pi/9$ radians at the grating frequency of 3 c/d. Introducing a 45° orientation difference between center and surround by making the surrounding grating oblique raised the average relative disparity threshold by a factor of 3.5, to 3.7 arcmin, a phase shift of 66°. Doubling the orientation difference to 90°, by making both center and surround gratings oblique but with opposite directions from vertical, added a small additional increment to threshold.

The literature contains suggestions that stereoscopic vision is specialized to extract the disparity of vertical or near-vertical contours (Marr and Poggio, 1979; Howard, 1982; DeAngelis et al.,

1991). Also, it is known that near-vertically oriented patterns are usually more readily detected and discriminated than obliquely oriented patterns [the "oblique effect" (Taylor, 1963)]. Hence, the high thresholds for the 90/45° and 45/135° pairs in Figure 1*B* might be attributable to the oblique orientation of one or both of the gratings. Yet when both gratings were given the same oblique orientation (the 45/45° case in Fig. 1*B*), disparity thresholds were as low as when both gratings were vertical. Thus, the visual system can extract disparity information from these obliquely oriented stimuli as precisely as it can from vertically oriented stimuli; it is relative, not absolute, stimulus orientation that determines whether the relative disparity threshold is high or low.

As Figure 1*B* makes clear, surrounding gratings that differ in orientation from the center by 45° or more contribute little or nothing to threshold for observer S3. Threshold in these cases is as high as threshold for an isolated grating, with no surround (Fig. 1*B*, right). This result is understandable if these thresholds are measures of the observer's sensitivity either to the absolute disparity of the grating or to the relative disparity between the carrier of the grating patch and its Gaussian-windowed contrast envelope (Hess and Wilcox, 1994; Schor et al., 1998; Langley et al., 1999; McKee et al., 2004), the disparity of which was fixed at zero (or perhaps between the carrier and other, more distant reference stimuli, such as the edges of the monitor). If the observer was more sensitive to any of these disparities than to the relative disparity between the two carriers, threshold would depend little on the grating orientation difference. Indeed, threshold should be independent of the presence of the second grating. And thus it is for different-orientation gratings: for observer S3, the presence of the surround is essentially irrelevant if it differs in orientation from the center by $\sim 45^\circ$ or more. Two other observers (not shown) displayed the same pattern of results, whereas observers S4 and S5 displayed very high thresholds for an isolated Gabor patch (immeasurably high for S4) (Fig. 1*B*). Evidently, there is considerable variability in observers' sensitivity to reference stimuli against which to judge the carrier disparity of this stimulus.

Thus, relative orientation affects stereoacuity for gratings. But sinusoidal gratings (one-dimensional, periodic, narrowband) are atypical visual patterns. In fact, the effect of relative orientation is not confined to gratings. The effect occurs also for patterns made of filtered two-dimensional noise (Fig. 2). Gaussian noise was filtered through a 4-octave (0.5–8.0 c/d) spatial-frequency passband and either a 60 or 30° orientation passband. Independent noise samples appeared in center-surround configuration, with the mean orientation being the same (either 60 or 120°) in center and surround or different (60° in one, 120° in the other). Stereoacuity depended on the mean orientation difference and on the orientation bandwidth. These effects interacted. Centers and surrounds with broadly filtered noise and a 60° mean orientation difference had abutting orientation distributions; their stereo thresholds were moderately higher (1.85 times; range, 1.3–2.4) than thresholds for centers and surround with the same orientations (Fig. 2, left). For narrowly filtered noise, the same 60° mean orientation difference left a 30° gap between the orientation distributions of center and surround. Here, the threshold elevation was distinctly larger (2.8 times; range, 2–3.5) (Fig. 2, right). Viewed differently, narrowing the orientation bandwidth of differently oriented patterns elevated threshold more (two times; range, 1.9–2) than narrowing the orientation bandwidth of same-orientation patterns (1.3 times; range, 1.3–1.3). As found using grating patterns, stereoacuity for two-dimensional, nonpe-

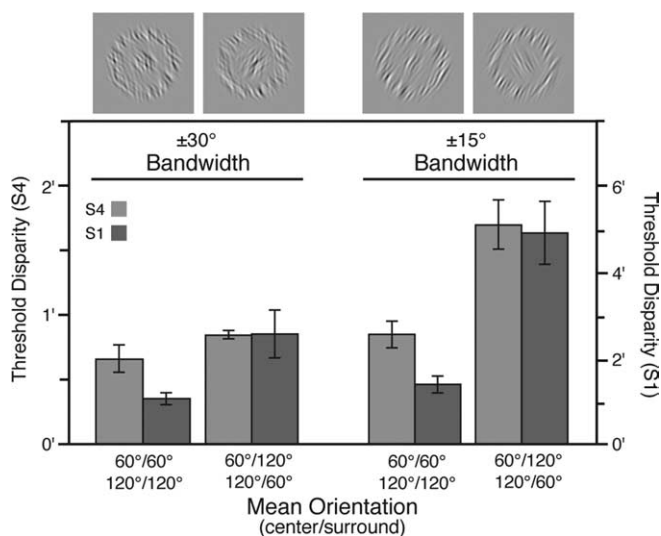


Figure 2. Stereoacuity for filtered noise as a function of mean center-surround orientation difference and orientation bandwidth. Mean orientation was either 60 or 120°. A 60° difference in mean orientation raised thresholds compared with a 0° difference, and this effect was larger for stimuli with narrow (30° range; right side) compared with broad (60° range; left side) orientation bandwidth. Spatial frequencies were confined to four octaves (0.5–8.0 c/d at the 126 cm viewing distance). The monocular examples appearing at the top are based on 3-octave (0.5–4.0 c/d) stimuli, for clarity at the printed scale.

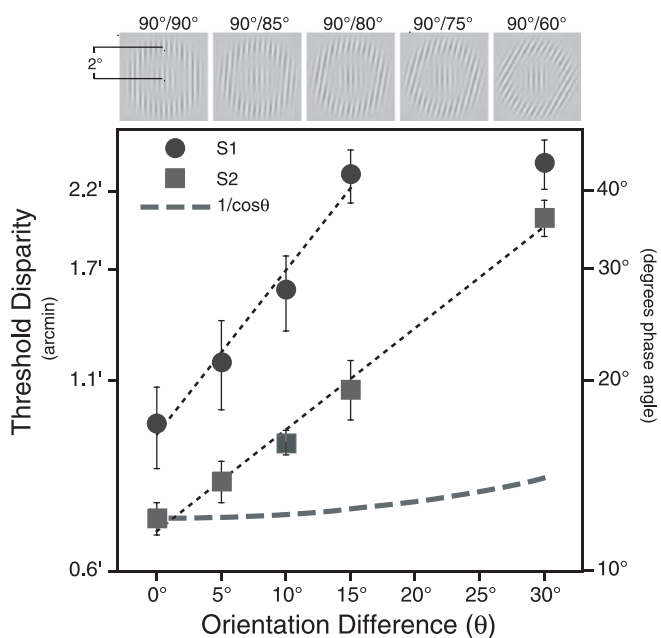


Figure 3. Stereoacuity for center and surround gratings for two observers: effect of small orientation differences. Monocular stimuli at the top of the figure show orientation-difference conditions. Exponential fits (up to saturating values) appear as dotted lines. The dashed curve showing $1/\cos\theta$ values gives an expected result for observer S2 if threshold was determined by the spatial disparity of one grating in the direction perpendicular to the other grating. For observer S1, this curve would be shifted vertically.

riodic patterns with broad spatial-frequency bandwidth depends on the relative orientations of the patterns.

Small orientation differences ($\leq 30^\circ$) provide the most revealing view of the effect. We see that relative disparity thresholds increase exponentially with the magnitude of orientation difference between center and surround gratings (Fig. 3). An orientation difference of 10° elevated threshold significantly, and an

orientation difference of $15\text{--}20^\circ$ doubled threshold from its minimal value. This threshold elevation greatly exceeds what one would expect from purely geometric considerations. For example, a projection model predicts thresholds proportional to the magnitude of the horizontal disparity of the center grating as projected onto the perpendicular disparity axis of the surrounding grating. Thresholds in this case would increase with the orientation difference, θ , only as $1/\cos\theta$. This is much smaller than the observed increase (Fig. 3). The orientation-difference bandwidth for relative disparity shown in Figure 3 approximates the orientation bandwidth for luminance contrast found for cortical neurons in the monkey (Schiller et al., 1976; De Valois et al., 1982). Many of these contrast-sensitive neurons are also sensitive to absolute disparity. Human psychophysical data reveal a similar orientation channel bandwidth (Campbell and Kulikowski, 1963; Mostafavi and Sakrison, 1976; Thomas et al., 1982; Phillips and Wilson, 1984). Thus, the data in Figure 3 suggest that relative disparity is computed within, rather than between, orientation-selective channels.

Horizontal versus perpendicular matching directions

I also tested center and annular surround gratings with orientations that differed symmetrically about the vertical (e.g., center at 80° and surround at 100°). The effect of relative orientation was indistinguishable from that of Figure 3, where the center had a fixed vertical orientation. Considerations of the stereo matching direction make this important. Matching between left- and right-eye stimulus images is conventionally measured in the horizontal or epipolar direction. A fixed spatial disparity measured in this direction corresponds to a larger phase disparity for vertical than for nonvertical stimuli, with implications that depend on whether threshold reflects sensitivity to interocular spatial offset or interocular phase offset. However, the stereo matching direction for one-dimensional stimuli, such as gratings, may well be perpendicular to stimulus orientation (Zhu and Qian, 1996; Anzai et al., 1997; Morgan and Castet, 1997; Farell, 1998, 2003a; van Ee and Schor, 2000; Patel et al., 2003), although this issue is still controversial (Cumming, 2002). Stimuli with mirror-imaged orientations allow us to side-step the issue entirely. We can safely assume that the visual system estimates the absolute disparity of a grating with orientation $+\theta$ from vertical with approximately the same precision as a grating with orientation $-\theta$ from vertical (at least for stimuli centered on the fovea, as is the case here). Therefore, if the relative-disparity computation was indifferent to relative orientation, then a pair of gratings with the mirror-image orientations $+\theta$ and $-\theta$ would be computationally equivalent to a pair with the same orientation, either $+\theta$ or $-\theta$. Because of the symmetry of the stimuli, this holds whether the visual system carries out stereo matches in the horizontal or the perpendicular direction. Yet, center and surround gratings with symmetrical orientations support the same sensitivity to relative disparity as grating pairs in which the center had a fixed vertical orientation (supplemental Fig. 1, available at www.jneurosci.org as supplemental material). Hence, whatever the stereo-matching direction, it does not interact with the effect of relative orientation on sensitivity to relative disparity observed here.

Orientation effect not tied to spatial arrangement

The center/surround organization used here could have placed parts of the stimuli into a spatial arrangement that expressly benefits same-orientation depth discrimination. Analogous effects occur in contrast detection of oriented patterns, as in the well-known enhancement of sensitivity and figure-ground segrega-

tion for collinear Gabor patches (Field et al., 1993; Polat and Sagi, 1993). [On the other hand, the apparent contrast of a central grating is lowered by iso-orientation surrounds (Cannon and Fullencamp, 1991; Yu et al., 2001) and can increase with cross-orientation surrounds (Yu et al., 2001), and spike rates of individual cortical cells are analogously affected (Sillito et al., 1995; Marcus and van Essen, 2002).]

In fact, the effect of relative orientation on stereoacuity depends neither on a specific spatial arrangement of the two stimuli, nor on specific absolute orientations. To show this, I reduced the annular surround to two soft-edged flankers appearing on opposite sides of the central Gabor patch. Relative to the center, the flankers appeared above and below, to the left and right, or on either of two pairs of oblique positions (Fig. 4A). Most observers found the depth discrimination task difficult with these smaller reference stimuli. Observers also showed marked differences in absolute thresholds and considerable idiosyncratic variation across the four flanking arrangements. Still, increasing the center-flanker orientation difference from 0 to 20° produced on average the same factor-of-two drop in stereoacuity as seen with the full annular surround (Fig. 4B). The “20° difference” data shown here for observer S1 are average thresholds for 70 and 110° flankers (the center orientation was 90°); thresholds for 70 and 110° flankers were similar, as was typically the case. However, interactions between the position and orientation of the flanker appeared in the data of two of the observers. Figure 4C shows an example of this by plotting the different-orientation data of observer S4 for a single flanker orientation. The 20° difference between the 90° center and 110° right oblique flankers raised threshold 3.9-fold over the 0°-difference threshold for this observer. Yet the same 20° difference between center and left oblique flankers lowered her thresholds marginally, bringing the average oblique threshold elevation to 1.9. Other interactions of this kind also tended to be balanced around a mean threshold elevation of approximately two. The overall factor-of-two effect shown here was also observed for center orientations of 70 and 110° as well (again with flankers at a relative orientation of +20, 0, or –20°). Overall, the data show that sensitivity to the disparity of a central grating depends as much on the relative orientation of reference gratings that are aligned perpendicularly or obliquely to the central orientation as it does for collinear reference gratings. Whatever effect a particular stimulus spatial arrangement has on stereoacuity, it appears to be independent of the effect of relative orientation.

Effect of the orientation of components of two-dimensional patterns

Some natural-image textures resemble gratings in being strongly oriented. Others have a broad range of orientation components and lack a perceptually dominant orientation. How do the individual components of an orientationally broadband pattern contribute to stereo vision? Let us examine the influence of component orientation first on stereoacuity and then, in the next section, on perceived depth. As shown previously (Farell, 1998; Farell and Li, 2004; Delicato and Qian, 2005), a plaid formed by superimposing two static sinusoids allows observers no access to the disparities of the individual grating components. The gratings are seen to cohere in depth, forming a plaid localized in a single depth plane, even when the disparities of the two component gratings differ greatly. Therefore, neither of the superimposed components can act as a reference stimulus for judging the depth of the other. Yet, I found that superimposed gratings, oriented $\pm\theta$ degrees of vertical, act as an effective reference stimulus for

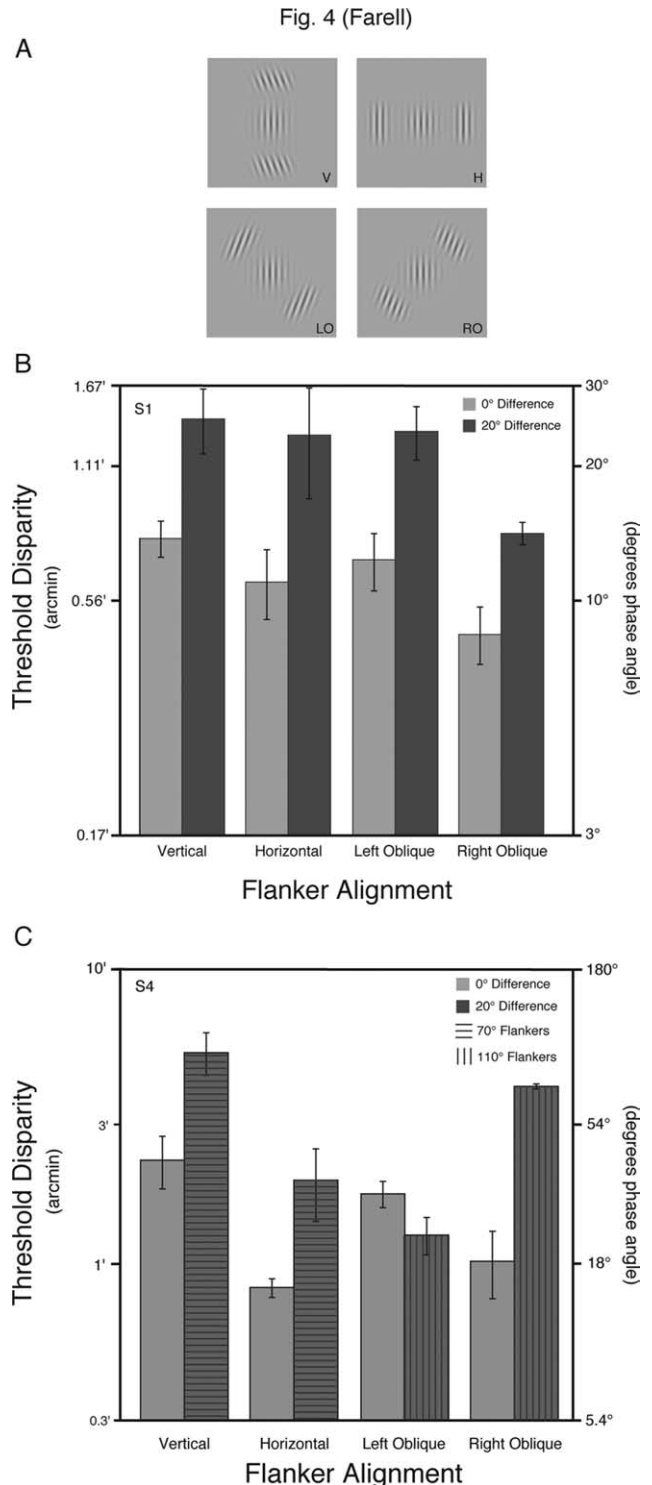


Figure 4. Stereoacuity for four center-flanker grating alignments. **A**, The flankers were formed by reducing the annular surround to a pair of wedges on opposite sides of the center grating. In separate conditions, these flankers were aligned with the center vertically (v), horizontally (h), or obliquely (lo and ro). Each Gaussian-edged flanker envelope subtended a polar angle of 90°, a quarter of the full surround. The orientation of the center Gabor was 70, 90, or 110°, and the orientation of the flanking grating differed from that of the center by 0, +20, or –20° in conditions run separately. **B**, **C**, Threshold as function of center-surround grating orientation difference for the four stimulus alignments. Compared with a 0° difference, a 20° orientation difference raised threshold by a factor of 2, on average. For these data, the center orientation was 90°; centers oriented at 70 and 110° produced similar results. Different-orientation thresholds for observer S1 (**B**) are averages of 70° and 110° flanker data; for S4 (**C**) thresholds are for 70° (vertical and horizontal flankers) and for 110° (oblique flankers), showing an example of interaction between absolute orientation and flanker position.

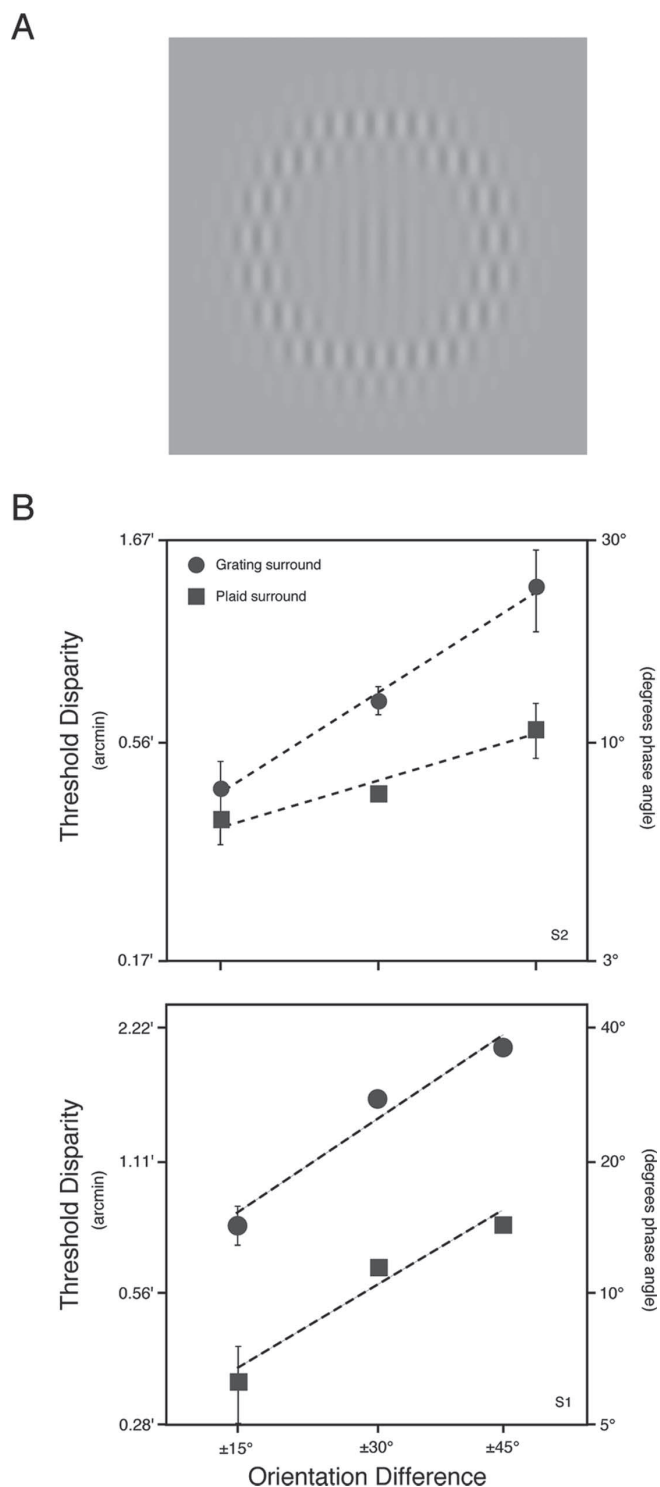


Figure 5. Stereoacuity with one- and two-dimensional patterns as stimuli. **A**, Center-surround Gabor stimulus. The surrounding plaid is composed of a pair of superimposed sinusoidal gratings, each with zero disparity, symmetrically oriented about the vertical. **B**, Stereoacuity as a function of the difference in orientation between the central grating patch and the surrounding plaid components.

judging the depth of the spatially separate vertical grating (Fig. 5). The components interacted; the plaid surround improved observers' stereoacuity compared with surrounds made up of either component grating alone. The lowering of threshold (by factors of 2.4 and 2.2 for the two observers shown, and similarly for the other two observers) is greater than can be accounted for by

independent noise sources from the two plaid components. Still, threshold increased with the orientation difference between the components of the plaid and the grating, as shown on the log-linear plots of Figure 5B, although overall the increase was less steep than when gratings formed both center and surround. Thus, stereoacuity does show an influence of the relative orientation of the central grating and the components of the surrounding plaid: a spatially two-dimensional pattern carries an imprint of the orientation of its one-dimensional components onto measures of stereo sensitivity.

Perceived depth of two-dimensional patterns

Seeking to generalize further the effect of relative orientation, one might ask whether the imprint of component orientations extends to suprathreshold disparities, in the realm of perceived depth. Again, I used a central grating and a surrounding plaid, this time giving the surround a nonzero disparity. As noted above, superimposed sinusoids of different orientations (but similar spatial frequencies) are seen to cohere in depth, overcoming their disparity difference to form a plaid at a single apparent depth (Farrell, 1998; Farrell and Li, 2004; Delicato and Qian, 2005). This coherence holds even if one grating has zero disparity or if one has positive disparity and the other negative disparity. Whether the plaid is seen as near or far relative to an isotropic reference stimulus depends on the sign of the horizontal component of the plaid disparity (Farrell, 1998). To determine the perceived depth of a plaid relative to an oriented reference stimulus, I presented an annular plaid with fixed disparity and probed its perceived depth using a central Gabor patch with variable disparity. The plaid was made from superimposed sinusoids oriented at 45 and 135°. The components of the plaid had fixed disparities, one at zero and the other at 1.67 or 2.22 arcmin (phase angles of 30 or 40°) for different observers. This gave the interocular disparity of the plaid as a whole either a direction of +45° or a direction of -45°, depending on which of the two components had nonzero disparity (Fig. 6A); these alternative plaids had disparities that were equal in magnitude and that differed only in the sign of their vertical components. The central grating, oriented at either 45 or 135°, had a disparity that was selected randomly on a trial-by-trial basis from among a wide range of values, both positive and negative.

The horizontal disparity of the plaids was fixed across conditions; their vertical disparity varied only in sign. Therefore, if the disparity of a plaid determines its perceived depth, we would expect the depth of the plaid relative to the central grating to depend only on the disparity of the grating. However, the perceived relative depth depended on the orientation of the grating as well as its disparity. In particular, relative depth depended on whether the grating matched the plaid's zero-disparity component or its nonzero-disparity component in orientation.

Data take the form of psychometric functions that give the probability of judging the central grating as "far" for each of the disparity values of the grating (Fig. 6B); the center appeared at the same depth as the surround when the disparity of the center gave a "far" probability of 0.5. For the plaid with a disparity direction of -45°, the component oriented at 45° had a disparity of 1.67 (observer S1) or 2.22 (observer S2) arcmin. Observers matched the depth of this plaid with a central 45° grating having a disparity that was a large fraction of 1.67 or 2.22 arcmin. However, observers matched the same plaid with a 135° central grating having a disparity of approximately zero. And zero, recall, is the disparity of the component of the plaid oriented at 135°. Conversely, a 45° central grating with a disparity of approximately

zero and a 135° central grating with a disparity that was a large fraction of 1.67 or 2.22 arcmin both matched the depth of the plaid that had a disparity of +45°. Thus, flipping the disparity direction of the plaid vertically, from +45 to -45°, changed its perceived depth relative to the grating; how the depth changed depended on the orientation of the grating. On average, the vertical shift in the disparity direction of the plaid changed the depth-matching disparity of the central grating by 73.9% ($\pm 6\%$) of the disparity magnitude of the plaid. This value corresponds to the proportional magnitude of the horizontal projection of the disparity of the plaid (cosine of 45° = 0.707), a relation whose generality needs to be tested with different plaid angles. The perceived depth of the plaid depended on the orientation of the reference grating over the entire disparity range. In summary, relative to an oriented reference stimulus, a two-dimensional pattern is not seen at a depth corresponding to the disparity of the pattern as a whole. Instead, its perceived depth is proportional to the relative disparity of similarly oriented components in the pattern and the reference stimulus.

Discussion

Signal differencing is a general sensory coding strategy for gaining invariance and sharpening contrast (Hartline and Ratliff, 1957). It is basic to the processing of contrast and color as well as depth. In color vision, for example, each of the three cone classes is individually color blind. Taking differences of responses across cone types yields opponent-color signals that deconfound wavelength and intensity information and allow fine chromatic discriminations. In a similar way, humans are poor at judging absolute disparities, which confound stimulus depth and ocular vergence. Differencing absolute disparities yields a relative disparity signal that eliminates the vergence component and permits fine depth discrimination.

A relative disparity computed on absolute disparities alone would also have the advantage of generality. It would be invariant to stimulus differences along other dimensions, provided only that variation on these dimensions conserves absolute disparity. Thus, any stimulus capable of evoking an absolute-disparity response should in principle be able to function as a reference stimulus for judging the relative stereoscopic depth of any other such stimulus. However, the present experiments show that this invariance does not hold; stereoscopic reference stimuli are orientation specific. Orientation specificity modulates how sensitive humans are to relative disparity (Fig. 3) and what relative depth they perceive (Fig. 6). On the surface, this specificity appears to cancel the generality that would be gained by simply differencing absolute disparities and ignoring other stimulus dimensions.

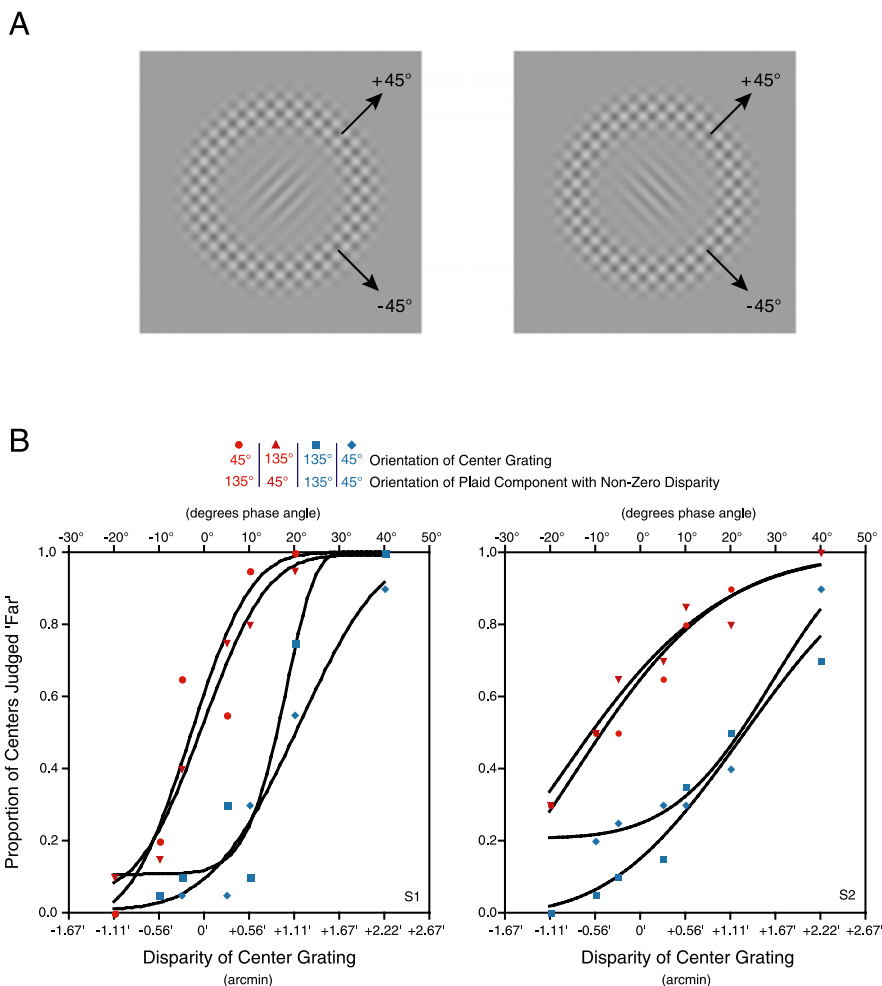


Figure 6. Perceived depth as a function of relative orientation. **A**, Annular plaid stimulus. Gratings oriented at 45 and 135° made up the plaid. When the 45° grating had positive disparity and the 135° grating had zero disparity, the plaid as a whole had a disparity direction of -45°. When the 135° grating had positive disparity and the 45° grating had zero disparity, the plaid as a whole has a disparity direction of +45°. The two plaids differ only in the sign of their vertical disparity component and are perceptually indistinguishable relative to an isotropic reference stimulus. Reference stimuli in this experiment were gratings oriented at 45° (left) or 135° (right). **B**, Perceived depth of a plaid surround relative to a central grating patch, for two observers. The plaid had a constant horizontal and vertical disparity magnitude; the sign of the vertical disparity could be positive or negative. For observer S1, the 3 c/d components of the plaids had disparities of 0 and 1.67 arcmin (a 30° interocular phase shift); for observer S2, the values were 0 and 2.22 arcmin (40°). The perceived depth of the plaid varied with its disparity direction and with the orientation of the central grating. Yet, for each observer, the horizontal disparity of the plaid was identical throughout the experiment.

No traditional “fixation point” appeared in the present experiments, and this goes to the gist of the results: relative judgments depend as much on the reference stimulus (often the fixation point in vision studies) as on the target stimulus. A broadband, isotropic fixation point serves as an all-purpose reference stimulus, performing that function for most any target stimulus and thereby masking the specificity of mechanisms that compare the target with target-like components of the fixation stimulus. For stereopsis, these mechanisms are specific to orientation.

Relative orientation and relative disparity appear to be non-separable dimensions in human vision. This is not a general property of either orientation or disparity in combination with other dimensions but instead is specific to the combination of orientation and disparity. For example, humans discriminate spatial-frequency differences equally well whether the stimuli are parallel or orthogonally oriented (Burbeck and Regan, 1983; Bradley and Scottun, 1984; Olzak and Thomas, 1991), and hu-

man sensitivity to relative disparity as a function of spatial-frequency differences is distinct from that of orientation differences reported here (Farell, 2003b).

Correlated noise?

The data also speak to an alternative interpretation of this orientation specificity. Instead of resulting from a within-orientation-channel computation of relative disparity, the effect of relative orientation might result from correlated disparity noise. Unless the stimulus is intrinsically noisy or corrupted by adding noise to it, it is the observer's internal noise that limits threshold (Fechner, 1860/1966; Barlow, 1977). If internal noise is correlated across disparity-sensitive mechanisms tuned to similar orientations, then differencing the output of these mechanisms will subtract off a portion of the noise, resulting in a comparatively high signal-to-noise ratio. Conversely, differencing would add the uncorrelated disparity noise of channels tuned to very different orientations, resulting in a lower signal-to-noise ratio. Thus, threshold disparity for stimuli with similar orientations would be lower, even if the relative disparity computation was indifferent to relative stimulus orientation. However, relative orientation has a powerful effect not only when the disparity difference is near threshold and noise-dominated but also when it is large, and this argues against the correlated noise hypothesis. Perceived depth varies with disparity at values well beyond the disparity threshold, and even at these large disparities perceived depth depends crucially on the relative orientation of stimulus components, as seen in the psychometric functions of Figure 6*B*. The effect of relative orientation is to shift these perceived-depth functions over their entire disparity range. A multiplicative noise, kept proportional to the magnitude of the disparity signal, could maintain the signal-to-noise ratio as disparity varies, but then the change in perceived depth with changing disparity would be unexplained. Perceived depth escapes the noise and increases with disparity magnitude but still shows the effect of relative orientation.

Joint coding of disparity and orientation

The results reported here support a nonseparable coding of disparity and orientation. Its source can be found in the link between the stereo-matching direction and stimulus orientation. Matching stimulus images between the left and right eyes in a direction perpendicular to the stimulus orientation encodes disparity as a two-dimensional vector (Anzai et al., 1997; Farell, 1998, 2003a; Mikaelian and Qian, 2000; van Ee and Schor, 2000; Matthews et al., 2003; Patel et al., 2003). In this case, a magnitude-differencing algorithm will yield accurate relative disparities only to the extent that stimulus orientations are the same (or symmetrical about the horizontal axis). (Spatial-frequency similarity is necessary, too, if disparity is coded in terms of phase.) This vector code would be advantageous for such purposes as integrating disparity with frontoparallel motion signals, which are intrinsically two-dimensional; there is evidence of such integration occurring in cortical area MT (Fernandez et al., 2002; Pack et al., 2003). However, unless disparity magnitudes are normalized with respect to the horizontal direction before the calculation of relative disparity, stereoscopic depth will vary with relative orientation.

Normalization solves the problem that a vertical stimulus, for example, needs a larger phase disparity than an oblique stimulus of the same spatial-frequency content to signal a disparity of a particular horizontal spatial extent. The data presented here point to the absence of normalization. How, then, do humans perceive the relative depth of stimuli with very different orientations? Perhaps we are unable to do so in the absence of intermediate

orientations in the scene to serve as common reference stimuli. Alternatively, second-order stimulus features (Hess and Wilcox, 1994; Schor et al., 1998; Langley et al., 1999; McKee et al., 2004) might play this role in naturalistic scenes. Another possibility is that normalization does occur, but only among the spatially overlapping components of a single object, such as the components of the plaids used here. This could account for the reduced size of the relative-orientation effect on threshold disparity between a plaid and a grating (Fig. 5). But the effect persists even there, suggesting that the normalization is either imperfect or requires more than just two components to overcome the effect of orientation difference.

The inseparable coding of disparity and orientation in relative depth judgments has counterparts in the inseparability of spatial and temporal frequency in motion energy computations (Adelson and Bergen, 1985; Reid et al., 1991; Basole et al., 2003; Mante and Carandini, 2005) and the conjoint coding of motion and disparity that may underlie phenomena such as the Pulfrich effect (Qian and Andersen, 1997; Morgan and Fahle, 2000; Anzai et al., 2001), although this effect can be explained by other means (Read and Cumming, 2005a,b). In the studies reported here, disparity and orientation interact between stimuli rather than within a single stimulus, underscoring the generality and importance of nonseparable coding in vision, and showing the existence of a pathway supporting relational judgments between stimuli in which the linked dimensions are not orthogonalized at any preperceptual stage.

The orientation specificity of relative disparity places an intrinsic limit on stereo processing. It can limit stereo resolution without changing disparities: an oriented zero-disparity stimulus retains its disparity after a rotation in the X-Y plane, regardless of how one measures disparity. However, it does not retain its utility as a stereo reference for a second oriented stimulus. Of particular importance will be to determine how the effect of relative orientation plays out in scenes of three or more stimuli. In a scene containing two stimuli, the relative orientation of stimuli judged with respect to their relative depth is unambiguous. In scenes of three or more stimuli, neither relative orientation nor relative disparity needs to be calculated on stimulus pairs. The calculation might be more global. For example, observers might compute the relative disparity of each stimulus using the mean orientation and disparity of the ensemble as a common nonlocal reference value. Thus, in multiobject scenes, the distribution of texture orientations across object surfaces, and the spatial distribution of surfaces in the scene, could affect observer's perception of the relative depth of a particular pair of surfaces. In a simple test of this, I measured thresholds for detecting depth differences among six Gabor patches forming a circular array. The orientation difference between adjacent patches mattered (observers were more sensitive to relative depth when neighboring orientations were similar) even when the orientation distribution was fixed and only the spatial ordering of the orientations varied (supplemental Fig. 2, available at www.jneurosci.org as supplemental material). This argues for a relative-disparity computation that is local in space as well as in orientation.

References

- Adelson EH, Bergen JR (1985) Spatiotemporal energy models for the perception of motion. *J Opt Soc Am A* 2:284–299.
- Anzai A, Ohzawa I, Freeman RD (1997) Neural mechanisms underlying binocular fusion and stereopsis: position vs. phase. *Proc Natl Acad Sci USA* 94:5438–5443.
- Anzai A, Ohzawa I, Freeman RD (2001) Joint encoding of motion and depth by visual cortical neurons: neural basis of the Pulfrich effect. *Nat Neurosci* 4:513–518.
- Barlow HB (1977) Retinal and central factors in human vision limited by

- noise. In: Vertebrate photoreception (Barlow HB, Fatt P, eds), pp 337–358. New York: Academic.
- Basole A, White LE, Fitzpatrick D (2003) Mapping multiple features in the population response of visual cortex. *Nature* 423:986–990.
- Bradley A, Scottun BC (1984) The effects of large orientation and spatial frequency differences on spatial discriminations. *Vision Res* 24:1889–1896.
- Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10:433–436.
- Burbeck CA, Regan D (1983) Independence of orientation and size in spatial discriminations. *J Opt Soc Am* 73:1691–1694.
- Campbell FW, Kulikowski JJ (1963) Orientation selectivity of the human visual system. *J Physiol (Lond)* 186:437–445.
- Cannon MW, Fullencamp SC (1991) Spatial interactions in apparent contrast: inhibitory effects among grating patterns of different spatial frequencies, spatial positions and orientations. *Vision Res* 31:1985–1998.
- Cumming BC (2002) An unexpected specialization for horizontal disparity in primate visual cortex. *Nature* 418:633–636.
- DeAngelis GC, Ohzawa I, Freeman RD (1991) Depth is encoded in the visual cortex by a specialized receptive field structure. *Nature* 352:156–159.
- Delicato LS, Qian N (2005) Is depth perception of stereo plaids predicted by intersection of constraints, vector average or second-order feature? *Vision Res* 45:75–89.
- De Valois RL, Yund EW, Hepler N (1982) The orientation and direction selectivity of cells in macaque visual cortex. *Vision Res* 22:531–544.
- Erkelens CJ, Collewijn H (1985) Eye movements and stereopsis during dichoptic viewing of moving random-dot stereograms. *Vision Res* 25:1689–1700.
- Farrell B (1998) Two-dimensional matches from one-dimensional stimulus components in human stereopsis. *Nature* 395:689–693.
- Farrell B (2003a) Detecting disparity in two-dimensional patterns. *Vision Res* 43:1009–1026.
- Farrell B (2003b) Relative binocular disparity: relative to what? *J Vision* 3:462a.
- Farrell B, Li S (2004) Perceiving depth coherence and transparency. *J Vision* 4:209–223.
- Fechner GT (1860/1966) *Elemente der Psychophysik*. Leipzig: Brietkopf & Hartel. *Elements of psychophysics* (Adler HE, translator). New York: Holt, Rinehart and Winston.
- Fernandez JM, Watson B, Qian N (2002) Computing relief structure from motion with a distributed velocity and disparity representation. *Vision Res* 42:883–898.
- Field DJ, Hayes A, Hess RF (1993) Contour integration by the human visual system: evidence for a local association field. *Vision Res* 33:173–193.
- Hartline HK, Ratliff F (1957) Inhibitory interaction of receptor units in the eye of limulus. *J Gen Physiol* 40:357–376.
- Hess RF, Wilcox LM (1994) Linear and non-linear filtering in stereopsis. *Vision Res* 34:2431–2438.
- Howard IP (1982) Human visual orientation. New York: Wiley.
- Langley K, Fleet DJ, Hibbard PB (1999) Stereopsis from contrast envelopes. *Vision Res* 39:2313–2324.
- Mante V, Carandini M (2005) Mapping stimulus energy in primary visual cortex. *J Neurophysiol* 94:788–798.
- Marcus DS, van Essen DC (2002) Scene segmentation and attention in primate cortical areas V1 and V2. *J Neurophysiol* 88:2648–2658.
- Marr D, Poggio T (1979) A computational theory of human stereo vision. *Proc R Soc Lond B Biol Sci* 204:301–328.
- Matthews N, Meng X, Xu P, Qian N (2003) A physiological theory of depth perception from vertical disparity. *Vision Res* 43:85–99.
- McKee SP, Verghese P, Farrell B (2004) What is the depth of a sinusoidal grating? *J Vision* 4:524–538.
- Mikaelian S, Qian N (2000) A physiologically-based explanation of disparity attraction and repulsion. *Vision Res* 40:2999–3016.
- Morgan MJ, Castet E (1997) The aperture problem in stereopsis. *Vision Res* 39:2737–2744.
- Morgan MJ, Fahle M (2000) Motion-stereo mechanisms sensitive to interocular phase. *Vision Res* 40:1667–1675.
- Mostafavi H, Sakrison DJ (1976) Structure and properties of a single channel in the human visual system. *Vision Res* 16:957–968.
- Ogle KN (1962) Spatial localization through binocular vision. In: *The eye, Vol 4, Visual optics and the optical space sense* (Davson H, ed), pp 271–324. New York: Academic.
- Olzak LA, Thomas JP (1991) When orthogonal orientations are not processed independently. *Vision Res* 31:51–57.
- Pack CC, Born RT, Livingstone MS (2003) Two-dimensional substructure of stereo and motion interactions in macaque visual cortex. *Neuron* 37:525–535.
- Patel SS, Ukwade MT, Stevenson SB, Bedell HE, Sampath V, Ogmen H (2003) Stereoscopic depth perception from oblique phase disparities. *Vision Res* 43:2479–2792.
- Pelli DG (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis* 10:437–442.
- Pelli DG, Zhang L (1991) Accurate control of contrast on microcomputer displays. *Vision Res* 31:1337–1350.
- Phillips GC, Wilson HR (1984) Orientation bandwidths of spatial mechanisms measured by masking. *J Opt Soc Am A* 1:226–232.
- Polat U, Sagi D (1993) Lateral interactions between spatial channels: suppression and facilitation revealed by lateral masking experiments. *Vision Res* 33:993–1000.
- Prince SJD, Pointon AD, Cumming BG, Parker AJ (2000) The precision of single neuron responses in cortical area V1 during stereoscopic depth judgments. *J Neurosci* 20:3387–3400.
- Qian N, Andersen RA (1997) A physiological model for motion-stereo integration and a unified explanation of Pulfrich-like phenomena. *Vision Res* 37:1683–1698.
- Read JCA, Cumming BG (2005a) The stroboscopic Pulfrich effect is not evidence for the joint encoding of motion and depth. *J Vision* 5:417–434.
- Read JCA, Cumming BG (2005b) All Pulfrich-like illusions can be explained without joint encoding of motion and disparity. *J Vision* 5:901–927.
- Regan D, Erkelens JC, Collewijn H (1986) Necessary conditions for the perception of motion in depth. *Invest Ophthalmol Vis Sci* 27:584–597.
- Reid RC, Soodak RE, Shapley RM (1991) Directional selectivity and spatio-temporal structure of receptive fields of simple cells in cat striate cortex. *J Neurophysiol* 66:505–529.
- Schiller PH, Finlay BL, Volman SF (1976) Quantitative studies of single-cell properties of monkey striate cortex. II. Orientation specificity and ocular dominance. *J Neurophysiol* 39:1320–1333.
- Schor CM, Edwards M, Pope DR (1998) Spatial-frequency and contrast tuning of the transient-stereopsis system. *Vision Res* 38:3057–3068.
- Sillito AM, Grieve KL, Jones HE, Cudeiro J, Davis JW (1995) Visual cortical mechanisms detecting focal orientation discontinuities. *Nature* 378:492–496.
- Taylor MM (1963) Visual discrimination and orientation. *J Opt Soc Am* 53:763–765.
- Thomas JP, Gille J, Barker RA (1982) Bandwidths of orientation channels in human vision. *J Opt Soc Am* 72:1642–1651.
- van Ee R, Schor CM (2000) Unconstrained stereoscopic matching of lines. *Vision Res* 40:151–162.
- Watson AB, Pelli DG (1983) QUEST: a Bayesian adaptive psychometric method. *Percept Psychophys* 33:113–120.
- Westheimer G (1979) Cooperative neural processes involved in stereoscopic acuity. *Exp Brain Res* 36:585–597.
- Yu C, Klein SA, Levi DM (2001) Surround modulation of perceived contrast and the role of brightness induction. *J Vision* 1:18–31.
- Zhu Y-D, Qian N (1996) Binocular receptive field models, disparity tuning, and characteristic disparity. *Neural Comp* 8:1611–1641.