2014

Perceived depth in non-transitive stereo displays.

Bart Farell
Syracuse Univeristy, bfarell@syr.edu

Cherlyn J. Ng
Syracuse University, cng03@syr.edu

Recommended Citation

Farell, B. & Ng, C. J. (2014). Perceived depth in non-transitive stereo displays. Vision Research, 105, 137–150.

Vision Research 105 (2014) 137–150

Contents lists available at ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

Perceived depth in non-transitive stereo displays

Bart Farell *, Cherlyn Ng

Institute for Sensory Research, Syracuse University, Syracuse, NY, USA

ARTICLE INFO

Article history:
Received 27 June 2014
Received in revised form 2 October 2014
Available online 23 October 2014

Keywords:
Stereoscopic depth
Disparity
Attention
Binocular vision

ABSTRACT

The separation between the eyes shapes the distribution of binocular disparities and gives a special role to horizontal disparities. However, for one-dimensional stimuli, disparity direction, like motion direction, is linked to stimulus orientation. This makes the perceived depth of one-dimensional stimuli orientation dependent and generally non-veridical. It also allows perceived depth to violate transitivity. Three stimuli, *A*, *B*, and *C*, can be arranged such that $A > B$ (stimulus *A* is seen as farther than stimulus *B* when they are presented together) and $B > C$, yet $A \leqslant C$. This study examines how the visual system handles the depth of *A*, *B*, and *C* when they are presented together, forming a pairwise inconsistent stereo display. Observers' depth judgments of displays containing a grating and two plaids resolved transitivity violations among the component stimulus pairs. However, these judgments were inconsistent with judgments of the same stimuli within depth-consistent displays containing no transitivity violations. To understand the contribution of individual disparity signals, observers were instructed in subsequent experiments to judge the depth of a subset of display stimuli. This attentional instruction was ineffective; relevant and irrelevant stimuli contributed equally to depth judgments. Thus, the perceived depth separating a pair of stimuli depended on the disparities of the other stimuli presented concurrently. This context dependence of stereo depth can be approximated by an obligatory pooling and comparison of the disparities of one- and two-dimensional stimuli along an axis defined locally by the stimuli.

1. Introduction

Binocularly viewed one-dimensional (1-D) patterns such as gratings, lines, and edges are subject to the stereo 'aperture problem', which makes their disparity directions and magnitudes ambiguous (Farell, 1998; Morgan & Castet, 1997). The result is that stereoacuity and perceived depth for 1-D patterns vary with stimulus orientation, a fact known for many years but open to diverse interpretations (Blake, Camisa, & Antoinetti, 1976; Ebenholtz & Walchli, 1965; Farell & Ahuja, 1996; Friedman, Kaye, & Richards, 1978; Morgan & Castet, 1997; Ogle, 1955; see Howard & Rogers, 2002). In general, the psychophysical effects of 1-D stimulus orientation are consistent with an effective disparity that has a direction perpendicular to the orientation (Chai & Farell, 2009; Farell, 1998, 2006; Morgan & Castet, 1997; Patel, Bedell, & Sampat, 2006; Patel et al., 2003; Quaia et al., 2013), though the physiological evidence is mixed (e.g., Cumming, 2002; Durand, Celebrini, & Trotter, 2007; Maske, Yamane, & Bishop, 1986).

The perceived depth between a 1-D stimulus and a 2-D stimulus is a case in which horizontal disparities do not predict stereo depth perception (Chai & Farell, 2009; Farell, Chai, & Fernandez, 2009). The depth between a grating and a plaid, for example, is predicted instead by a version of the intersection-of-constraints rule (Fennema & Thompson, 1979; Adelson & Movshon, 1982) applied to the two-dimensional disparity vectors. This calculation orthogonally projects the plaid's disparity vector onto the grating's disparity axis; examples are shown in Fig. 1. The perceived depth separating the stimuli varies with the relative magnitude of disparity components along this axis (Farell, Chai, & Fernandez, 2009). Equivalently, the disparities can be compared in the direction of the plaid's disparity; the grating's disparity in this case is given by the intersection of its constraint line with the plaid's disparity axis. We call the results of either version of this calculation the *projected disparity* value. Because relative but not absolute disparity directions enter into the computation, it is possible for two simultaneously presented stimuli, one 1-D and the other 2-D, to appear at the same depth even though the horizontal disparity of one is negative and that of the other is positive (Farell, Chai, & Fernandez, 2009). This allows us to create sets of stimuli that have contradictory depth relations. We study the perception of such stimuli here.

1.1. Violations of transitivity

Transitivity asserts that if *A* is further than *B*, and *B* is further than *C*, then *A* should be further than *C*. A transitive series has a

---

* Corresponding author.

http://dx.doi.org/10.1016/j.visres.2014.10.012
0042-6989/

consistent ordering, so its consistency is quantitative, not merely qualitative. Considering the discussion above, however, it would not be surprising to find violations of transitivity in depth when *A*, *B*, and *C* include both 1-D and 2-D stimuli. For example, given the proper choice of stimulus dimensionality and disparity, stimuli *A* and *B* are seen at the same depth when they are viewed together; *B* and *C* are seen at the same depth when they are viewed together; but *A* and *C* are seen at different depths when they are viewed together (see Figs. 1 and 2). Bringing all three stimuli together into a single display would show whether these pairwise depth relations determine the depth structure of the display as a whole. We ask here whether humans can see stable depth relations among *A*, *B*, and *C* when they are presented all at once, creating a display with internal pairwise inconsistency. How does stereo processing of such displays differ from those in which *A*, *B* and *C* have consistent pairwise disparities? Are there alternatives to pairwise depth comparisons that can resolve the incompatible disparities? Or are the incompatibilities not resolved but seen?

Our interest here is in characterizing how the depth seen in displays made up of pairwise inconsistent stimuli differs from the depth seen in displays whose stimulus pairs have consistent relative depths. We describe three experiments, with a grating and two plaids playing the roles of stimuli *A*, *B*, and *C*. The first experiment assessed the perceived depth order of the three stimuli directly. The second and third experiments examined depth-order judgments to a relevant subset of stimuli within the displays. The data show a stimulus-dependent recalibration of the effective disparity direction. The disparities of all the stimuli in the display, whether relevant to the task or not, contribute to the resulting depth judgments. This resolves the perceptual inconsistencies between the stimuli within the display and reveals a global disparity computation of depth judgments of 1-D stimuli.

2. General methods

The displays contained three stimuli in Experiment 1 and five (two of which were redundant) in Experiments 2 and 3. One stimulus was a sinusoidal grating patch and the others were plaid patches formed by summing two orthogonal gratings. Gaussian contrast envelopes (with zero disparity) defined the location of these stimuli. Individual stimuli were characterized by three parameters: Dimensionality (1-D or 2-D), disparity magnitude (fixed for plaids, varying in magnitude across trials for gratings), and disparity direction (between +45° and −45°, plus one case of 135° and −135°, where the positive and negative horizontal directions are 0° and ±180°, respectively). The orientation of the grating was either 45° or 135° in all three experiments. Because a grating's disparity direction can be regarded as normal to its orientation, a grating with a 45° orientation has an associated disparity axis running along the +135°/−45° direction, and a grating with a 135° orientation has one running along the +45°/−135° direction.

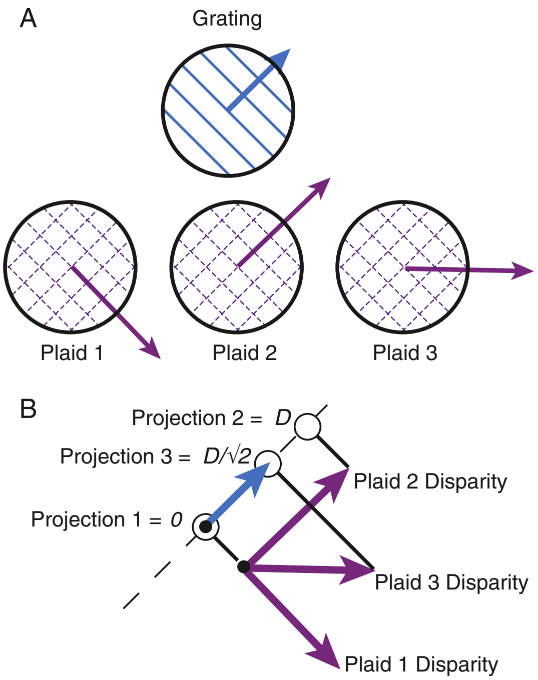

Fig. 1. Perceived depth predicted from projected disparities. (A) Arrows showing disparity vectors of sample grating (top) and three plaids (with disparity magnitudes exaggerated relative to the pattern wavelength). Disparity directions are 0° (horizontal) and ±45°. (B) Plaid disparities projected onto the grating's disparity axis. This axis is indicated by the dashed line. For clarity, the origins of the plaid disparity vectors are displaced from the origin of the grating disparity vector. The solid oblique lines intersect the grating's disparity axis perpendicularly, giving the projections of the plaids' disparities. The projected values assume a disparity magnitude is *D* for all three plaids. The relative sizes of the components along the grating's disparity predict that a grating with the disparity depicted here will appear farther in depth than one plaid, nearer than another, and at the same depth as the third, despite two of the plaids having equal horizontal disparities and therefore appearing in the same depth plane.

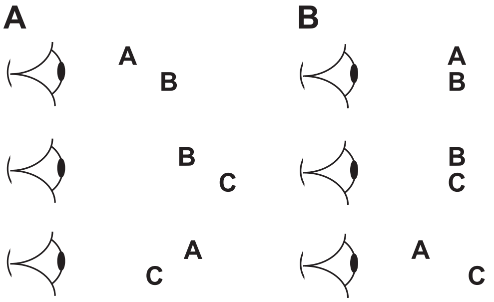

Fig. 2. Two examples of non-transitive depth, with stimuli *A*, *B*, and *C*. (A) The standard example of non-transitivity: *B* > *A*, *C* > *B*, A > *C*. (B) Alternative arrangement: *A* = *B*, *B* = *C*, *A* ≠ *C*.

The plaids had two sinusoidal components, one oriented at 45° and the other at 135°. In the case of a plaid with a disparity in the +45° direction, the right retinal image differed from the left solely by a phase shift of the 1-D component with the 135° orientation. The component oriented at 45° had the same phase in the two retinal images, a disparity of zero. When superimposed, these sinusoidal components perceptually cohere in depth, resulting in a unified 2-D stimulus seen in a single depth plane—a plaid rather than two distinct gratings (Adelson & Movshon, 1984; Calabro & Vaina, 2006; Delicato & Qian, 2005; Farell, 1998; Farell & Li, 2004). With the component disparities just described, the 2-D pattern features (for example, the 'blobs' formed at the intersections of the component gratings) have a disparity that is oblique, in the +45° direction. The horizontal component of this disparity is positive, corresponding to the 'far' depth at which the plaid is seen relative to a stimulus with zero disparity.[1]

All procedures carried out in the studies reported here followed the tenets of the World Medical Association Declaration of Helsinki and were approved by the Institutional Review Board of Syracuse University. All participants in the experiments participated with their informed consent.

---

[1] Direction is defined here by a vector from a location of a 2-D feature in the retinal image of the left eye to the nearest identical feature in the retinal image of the right eye after the two retinas have been overlaid in anatomical correspondence (where 'identical' discounts differences in contrast due to the Gaussian envelope).

3. Experiment 1. Perceived depth order

Previous work has shown that a grating and a plaid appear at the same depth when their disparity magnitudes are approximately the same, provided their disparity directions are parallel. If their disparity directions are perpendicular, however, they appear at the same depth when the grating's disparity magnitude is the smaller of the two and approaching zero, regardless of the magnitude of the plaid's disparity within the range tested (Chai & Farell, 2009; see Fig. 1). In the case of a pair of 2-D stimuli, a different rule applies. A depth match between two plaids occurs when the horizontal components of the stimulus disparities are equal, a result holding over a considerable range of vertical carrier disparity differences (Farell, Chai, & Fernandez, 2010). In all these cases, the disparity in question was that of the grating or plaid carrier. The contrast envelope had a disparity of zero.

In Experiment 1, the grating and pair of plaids making up a display had disparity directions drawn from the set {+45°, −45°}. Thus, any two stimuli selected from a display had disparities that had either parallel or orthogonal directions. Suppose the plaids had disparities with equal magnitudes and orthogonal directions, +45° and −45°. They would then have equal horizontal disparities and should appear at the same depth (Farell, Chai, & Fernandez, 2010). But if the grating (oriented at, say, 135°) had a disparity that gave it the same apparent depth as one of the plaids, it should appear at different a depth from the other plaid. This is because one grating-plaid pair has parallel disparity directions and the other pair has orthogonal disparity directions. Thus, when the plaids have orthogonal disparity directions, the perceived depths of the three component stimulus pairs—two plaids; the grating and one plaid; and the grating and the other plaid—are non-transitive. Note that this non-transitivity is based on a *pairwise* analysis of the perceived depths of component stimuli. It does not consider the influence on perceived depth arising from a more global calculation. An example of this display type appears in Fig. 3.

Experiment 1 used a direct method to answer the questions: Does the visual system resolve the contradictory depth information contained in non-transitive stimulus pairs when the three stimuli making up the pairs are presented at once? And if so, how does it do it? We compare these Non-Transitive displays with Transitive displays whose plaids had parallel disparity directions and no contradictory component depth signals.

3.1. Method

3.1.1. Stimuli

A grating and two plaids were arranged to form an isosceles triangle, with the grating occupying the top vertex (Fig. 3). All three stimuli were equidistant (2.25°) from the nominal point of fixation located at the middle of the lower edge of the triangle defined by the stimulus mid-points. The mid-point of the grating was separated from the mid-point of each plaid by a visual angle of 3.18°. The grating, like the components of the plaid, was sinusoidal with a spatial frequency of 2.0 c/d. The contrast of the grating was 0.2; the contrast of the each of the plaids' sinusoidal components was 0.1. The standard deviation of the Gaussian envelopes of both the gratings and the plaids was 0.53° of visual angle in both the horizontal and vertical directions.

Within a block of trials the plaid's disparities were fixed, while the magnitude of the grating's disparity varied from trial to trial. The plaids had a disparity magnitude of 1.67 arcmin, which is equivalent to a 20° phase disparity for the grating and a 1-D component of the plaid. There were four possible disparity direction combinations for the plaids on the left/right sides of the display: +45°/−45°, −45°/+45°, +45°/+45°, and −45°/−45°. The first two pairs, with orthogonal disparity directions, appeared in Non-Transitive displays. The second two, with parallel disparity directions, appeared in Transitive displays. The horizontal disparities of all the plaids were the same and were positive, which would place the plaids on the far side of a zero-disparity reference stimulus. The disparity direction of the grating was either parallel or orthogonal to the disparity directions of the two plaids in Transitive displays; it was parallel to the disparity direction of one plaid and orthogonal to that of the other in Non-Transitive displays. These relative disparity directions define the three conditions of the experiment: Parallel Transitive (shown schematically in Fig. 4, left), Orthogonal Transitive (Fig. 4, center), and Non-Transitive (Fig. 4, right).

Each plaid disparity direction (+45° and −45°) appeared equally often in the left and right plaid positions across blocks of trials. The duration of stimulus presentations was 174 ms (13 frames), with abrupt onsets and offsets. These presentations were observer-initiated. The stimuli were centered on CRT monitors with visible screen diagonals measuring 49 cm, one monitor for each eye. Viewing was at an optical distance of 1.25 m through a front-silvered mirror stereoscope. Observers' eyes were on or very nearly on the same horizontal plane as the centers of the monitors; their heads were perched on a chin-rest in upright posture. The apparatus gave observer's vergence angle the value appropriate for this viewing distance. Black nonius lines (191.3 × 1.8 arcmin) centered horizontally on each screen preceded stimulus presentation. The mean luminance of the monitors was 21 cd/m^2 during the trial except during presentation of the response screen (described below), when 21 cd/m^2 was the background luminance. A look-up table linearized luminance of the monitors, which were driven through their green guns after the R, G, and B signals were combined via attenuators to increase luminance resolution (Pelli & Zhang, 1991). The testing room was illuminated indirectly with an incandescent bulb and had an average luminance of approximately 6 cd/m^2.

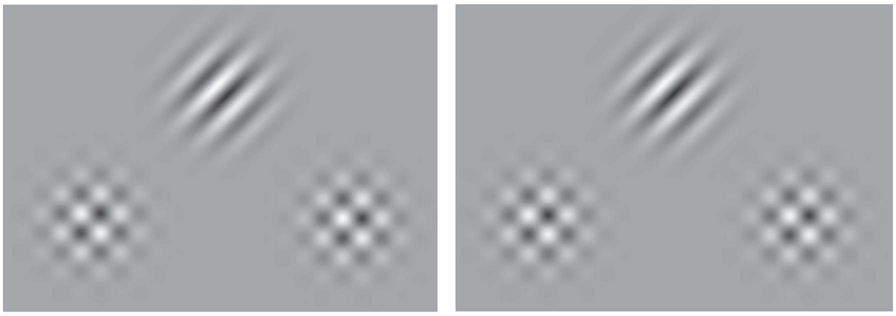

Fig. 3. A binocular Non-Transitive display. Under the conditions of the experiment, the stimulus contrasts would be lower than in this figure and the gray background would extend well beyond the boundaries shown here. The display would be preceded by nonius lines and followed by a response screen showing a display of three circles, each one coinciding with the visible borders of one of the three test stimuli. The images are meant to be fused convergently.

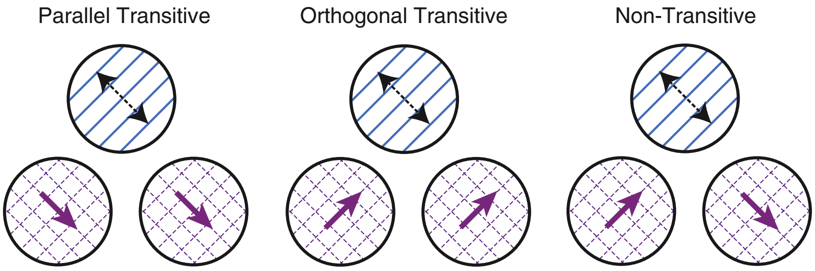

Fig. 4. Sketches of the disparity directions found in the displays of Experiment 1. The maroon arrows show the direction of plaid disparity. All are oblique (+45° or −45°) and have the same magnitude. The dashed black arrows designate variable-magnitude grating disparities (here for gratings oriented at 45°), which were either parallel or orthogonal to the plaid disparities.

Our intention was to make observers' judgments depend directly on the relative disparities of the grating and plaids that were the focus of the study. We attempted to avoid judgments mediated by the relative disparity of each of these stimuli with other, task-irrelevant, stimuli with disparity parameters of their own. Therefore, no fixation stimulus appeared along with the grating and plaids, the stimuli had soft-edged contrast envelopes, and contours and terminators that might function as reference stimuli, including the monitors' vertical edges and the ends of their horizontal edges, were excluded from binocular viewing by occluders (Chai & Farell, 2009).

3.1.2. Procedure

The task was to specify the depth order of the three stimuli. Observers did this by clicking with a mouse twice on a response screen. The first click indicated the stimulus that was seen as nearest and the second indicated the stimulus that was seen as farthest. These clicks were made on a screen that showed three circles that were approximately coincident with the visible boundaries of the three test stimuli, forming a triangle. This response screen appeared 250 ms after the offset of the test display and remained in view until the observer concluded her second response, ending the trial. The nonius lines then reappeared and remained until a subsequent click initiated the presentation of the next test display.

Seventy test trials appeared in a block. Two parameters changed across these trials: The grating disparity magnitude and the absolute phases of the grating and the 1-D components of the plaids. The grating disparity on each trial was chosen from a set of seven values selected in light of pilot data to span the perceptual range from 'grating nearest' to 'grating farthest'. The selection of disparities across trials was random, subject to the constraint that the seven disparities were presented an equal number of times per block. Treatment combinations, which were defined by grating orientation, plaid disparity direction, and (for Non-Transitive displays) plaid position, appeared in separate blocks of trials. Their presentation order was randomized within the constraint that the observer complete $N-1$ trial blocks for each combination before encountering any combination for the Nth time. Observer T1 completed 24 blocks of Transitive displays, half of them Parallel and half Orthogonal, and 24 blocks of Non-Transitive displays; observer T2 completed one-third as many blocks in the same proportion. In each block, experimental trials were preceded by 5 warm-up trials.

3.1.3. Observers

Two female undergraduate students served as observers. Their previous experience in psychophysical experiments was moderate and restricted to stereo studies in this laboratory. Neither observer was informed about the purpose of the experiment until her participation in it had ended. Both had normal acuity with spectacle correction; observer T2 was strongly myopic without corrective lenses.

3.2. Results

The purpose of Experiment 1 was to find out whether observers could perceive depth reliably among the pairwise inconsistent elements of Non-Transitive displays and to characterize differences between the depth order seen in Transitive and Non-Transitive displays. Of primary interest was how the grating's perceived depth changed with its disparity and how this differs across the three display types depicted in Fig. 4. After a look at raw probabilities for the Non-Transitive displays, we will examine psychometric functions.

3.2.1. Non-Transitive response probabilities

Fig. 5 shows the probabilities of 'Nearest' and 'Farthest' responses for Non-Transitive displays and reveals a surprising outcome of Experiment 1. The plotting convention, shown in Fig. 5A, mirrors the triangular arrangement of stimuli in the experimental displays, with the grating represented at the top vertex and the plaids represented on the left and right vertices. Data for the two observers appear in separate columns of panels b and c of Fig. 5. (The response probabilities for the remaining [Transitive] displays appear in Fig. S1; see Supplementary Material.) The position of a particular point gives the probability of judging each of the three stimuli as 'Nearest' (red disks) and 'Farthest' (blue disks) when the grating disparity has a particular value, coded by color saturation. The probability values are given by the proximity of each data point to the vertices of the triangle. For example, consistently selecting the grating as 'Nearest' when it had a particular disparity value would result in a data point at the upper vertex. A lower probability would make the point fall below this vertex.

Our interest here is in the relative probabilities of judging the two plaids as "Nearest" and "Farthest" as a function of the disparity of the grating. These probabilities vary along the left–right direction in the plots of Figs. 5 and S1. (The psychometric functions considered in the next section examine variation along the up-down direction.) Data for Non-Transitive displays can be plotted in two ways. Fig. 5B plots response probabilities according to plaid position. The left and right vertices represent a probability of 1.0 for selecting the left- and right-side plaid, respectively. "Nearest" and "Farthest" probabilities in Fig. 5B do not fall in the center of the plot. They fall on opposite sides of the midline, indicating that the plaid on one side is more often than the other to be selected as 'Nearest' or as 'Farthest', even though the left- and right-side plaids differ only in position. The same spatial response bias appears also in the data for Transitive displays (Fig. S1). (A third observer, an author who judged Non-Transitive displays only, produced results [not shown] very similar to those in Fig. 5, but with a response bias for left- and right-side plaids in the opposite direction from those appearing in Fig. 5B.)

Fig. 5C plots the same probabilities, but now with respect to the plaids' disparity direction, parallel versus orthogonal relative to the grating's disparity direction. The left vertex represents a probability

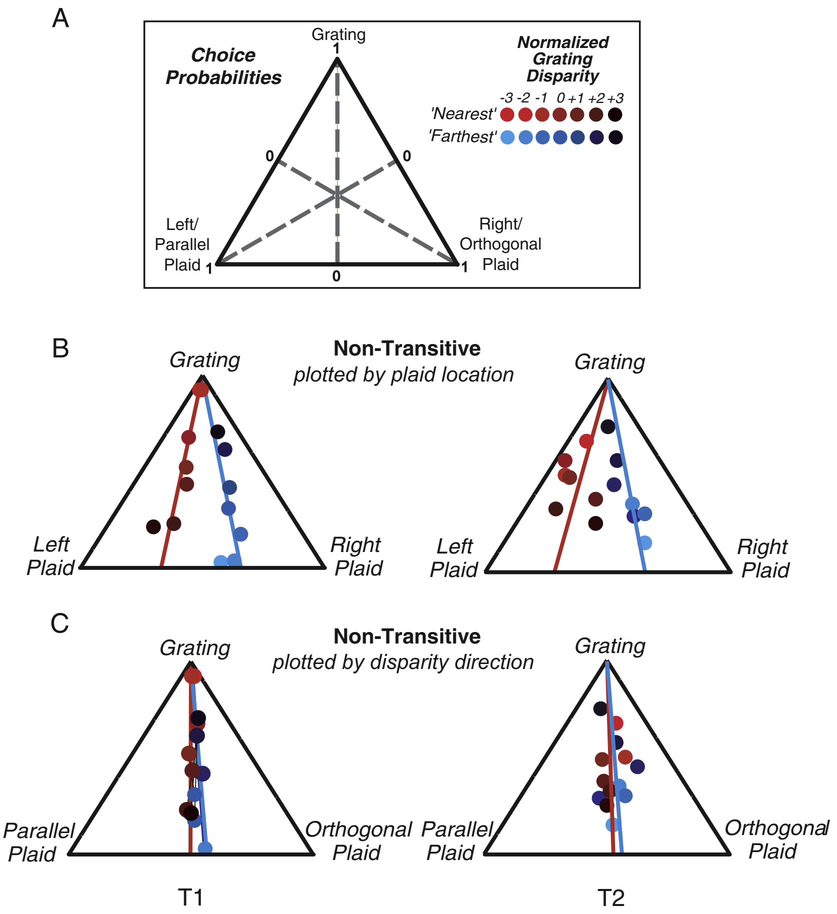

Fig. 5. 'Nearest' and 'Farthest' response probabilities for Non-Transitive displays in Experiment 1. (A) Plotting conventions. Probabilities of 'Nearest' and 'Farthest' judgments (red and blue disks, respectively) are given by the reciprocal distance from the triangle's vertices. Each vertex represents the grating or a plaid type. Disk saturation indicates grating disparity, rank ordered from −3 (largest negative disparity) to +3 (largest positive disparity). (B) 'Nearest' and 'Farthest' response probabilities for two observers, T1 and T2. Red and blue lines are best-fitting linear functions for 'Nearest' and 'Farthest' responses, respectively, constrained to pass through the upper vertex. The triangle's left and right vertices correspond here to plaid position. Plaids with disparities parallel and orthogonal to the grating's disparity appeared equally often in these two display positions. (C) The same probabilities plotted with triangle's left and right vertices representing plaid disparity direction, either parallel or orthogonal to the grating's disparity direction.

of 1.0 for 'Nearest' or 'Farthest' judgments of the plaid with a disparity direction parallel to that of the grating and the right vertex represents the same for the plaid with a disparity direction orthogonal to that of the grating. Here, the 'Nearest' and 'Farthest' functions fall very near the vertical mid-line of the plots. That is, the grating is equally likely to be seen as 'Nearest' or 'Farthest' relative to either plaid. There is scant difference between the response probabilities for parallel and orthogonal Non-Transitive plaids. The 'Nearest' and 'Farthest' probabilities are well described by linear functions (red and blue lines, respectively) constrained to pass through the upper vertex, keeping the relative probabilities between the two plaids constant. These functions do not differ significantly in slope ($t(10) = 1.95$, $p = 0.08$ and $t(10) = 1.20$, $p = 0.26$, for T1 and T2, respectively); both are very close to vertical and hence the same for parallel and orthogonal plaids. This surprising indifference of observers' judgments to the relative disparity directions of grating and plaid in Non-Transitive displays contrasts with the difference between Parallel and Orthogonal Transitive displays described in the next section and provides the impetus for Experiment 2.

3.2.2. Psychometric functions

Measures of horizontal disparity and measures of projected disparity lead to quite different expectations about psychometric functions for depth judgments of the grating. The horizontal disparity, which is the same for all of the plaids appearing in the experiment, provides no basis for explaining differences in perceived depth between the three display types. However, if projected disparities formed the effective metric, then the psychometric functions for Parallel and Orthogonal Transitive displays should be laterally displaced but otherwise similar (Chai & Farell, 2009; Farell, Chai, & Fernandez, 2009). The psychometric function for Non-Transitive displays, by contrast, might be expected to have a distinctive shape. To be judged as 'Nearest', the grating would have to appear nearer than the plaid with an orthogonal disparity direction. To be judged as 'Farthest', it would have to appear farther than the plaid with a parallel disparity direction. In the former case, depth matches for grating-plaid pairs occur when the grating's disparity is near zero and in the latter case when the grating's disparity magnitude equals the plaid's (Chai & Farell, 2009). Thus, Non-Transitive psychometric functions might be expected to be distinctively shallow, reflecting a paucity of both 'Grating Nearest' and 'Grating Farthest' judgments for grating disparities between 0° and 20°.

It is conceivable, too, that depth judgments for Non-Transitive displays might be distinctive by being mediated by horizontal disparities, as a means of resolving the disparity inconsistency. This would account for the similarity of 'Nearest' and 'Farthest' probabilities in Fig. 5C. The resulting Non-Transitive psychometric function would then be expected to fall between the two Transitive functions.

The probabilities of the grating receiving a “Nearest” and “Farthest” judgment as a function of the grating’s disparity appear in Fig. 6. These probabilities, shown separately for the three display types, correspond to the vertical height of the points in the three-variable plots of Figs. 5 and S1. They ignore variation among the points in the horizontal direction, which reflects differences between the plaids. The fits through the points are maximum-likelihood cumulative Gaussian functions.

By inspection, the psychometric functions for Non-Transitive displays are similar to those for the two Transitive display types, having similar slopes and intermediate lateral positions. Thus, the inconsistent pairwise depth relations in Non-Transitive displays do not appear to impede observers’ judgments of the grating’s depth relative to the plaids or to diminish sensitivity to disparity modulations. In fact, significant violations of monotonicity (Kendall’s *Tau*) in the psychometric functions for “Nearest” and “Farthest” judgments of the grating within individual runs were more frequent in Transitive than in Non-Transitive displays ($p < 0.05$).

“Nearest” and “Farthest” psychometric functions were combined (after inverting the former) to find the points of subjective equality (PSEs). The PSE, estimated by the 50% point on the psychometric function, is the disparity at which the grating has the same apparent depth as the plaids. The disparity at which the “Nearest” and “Farthest” curves cross within each plot of Fig. 5 gives a close visual approximation. The PSE is greater for Parallel Transitive displays (phase disparity of 20.6°, averaged across observers) than for Orthogonal Transitive displays (8.4°). For Non-Transitive displays it has an intermediate value (17.1°).

We compared the psychometric functions from the three conditions using 5000-iteration Monte Carlo simulations (Wichmann & Hill, 2001) on combined ‘Nearest’ and ‘Farthest’ data. These

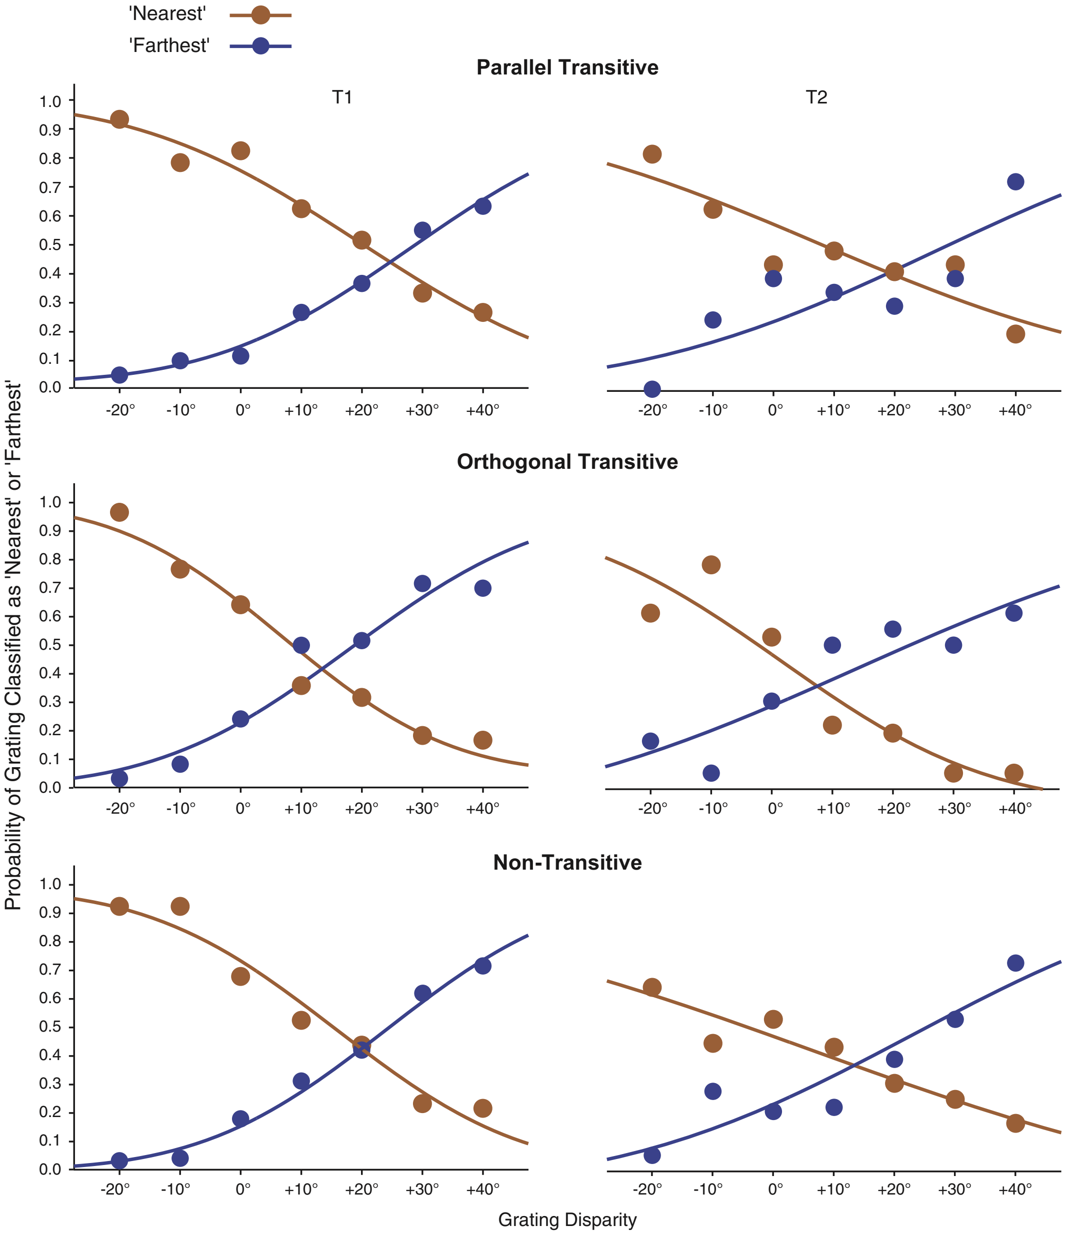

Fig. 6. Psychometric functions for the three conditions of Experiment 1. The probabilities are those that the grating was classified as ‘Nearest’ (red dots) or ‘Farthest’ (blue dots). The curves fitted through the data points are least-square cumulative Gaussian functions. Data for the two observers are shown separately. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

showed that PSEs for the Parallel and Orthogonal Transitive conditions differed significantly, with the former being larger for both observers. For T1 the difference was 12.6° where the 95% CI around an expected difference of zero was approximately ±4.3°; for T2 the corresponding value was 11.8° (CI: ±8.7°). The PSEs for Parallel Transitive and the Non-Transitive functions also differed significantly for observer T1 (5.7°, CI: ±3.5°) though not for observer T2 (1°, CI: ±8.9°), with the Transitive PSE being larger in each case. For Orthogonal Transitive and the Non-Transitive functions the PSEs differed significantly for both observers (6.9°, CI: ±3.5° for T1, and 10.8°, CI: ±8.3° for T2), the Non-Transitive PSE being larger. None of the slopes of the three pairs of psychometric functions differed significantly for either observer.

3.3. Discussion

The disparity directions within a Transitive display affected the grating's apparent depth relative to the plaids. When the grating and the plaids had the same disparity direction, the grating required a relatively large disparity to appear at the same depth as the plaids. When they had orthogonal disparity directions, a smaller grating disparity produced a depth match. Non-Transitive displays, however, did not produce the differing depth functions found between Parallel and Orthogonal Transitive displays. The perceived depth between the grating and a plaid in Non-Transitive displays was similar whether their disparity directions were parallel or orthogonal, with the PSE falling between the Parallel and Orthogonal Transitive PSEs. Thus, depth judgments of Non-Transitive displays appear quantitatively similar to those of the two types of Transitive displays and intermediate between them.

Despite being quantitatively intermediate, the Non-Transitive depth function is not an average of the Parallel and Orthogonal Transitive functions. Plaids with parallel and orthogonal disparity directions *individually* produce very similar depth functions when judged relative to the grating in Non-Transitive displays (Fig. 5C). The grating, whatever its disparity, appeared at approximately the same depth relative to each plaid. Thus, observers have resolved the inconsistent depth of the component stimulus pairs contained within Non-Transitive displays. The price of this resolution is an inconsistency between the results for Transitive and Non-Transitive displays: In the former, parallel and orthogonal disparities produce substantial differences in perceived depth; in the latter, their perceived depths are indistinguishable.

Parallel and orthogonal plaids have very different projected disparity values, but their perceived depths differ only in Transitive displays. In Non-Transitive displays, where their perceived depths are equivalent, projected disparities seem unable to account for the data. And while horizontal disparity components were equal for the two Non-Transitive plaids, they were equal as well for the two types of Transitive displays. Hence, horizontal disparity cannot explain differences in perceived depth between Transitive and Non-Transitive displays or between Parallel and Orthogonal Transitive displays. The results of Experiment 1 seem beyond explanations in which each stimulus contributes an independent disparity value to the calculation of depth. In order to preserve the independence of stimulus disparity contributions, it would be necessary to default to unparsimonious notions, in particular the possibility that different disparity components were used for judging different display types (for example, projected disparities for Transitive displays and horizontal disparities for Non-Transitive displays). The following two experiments tested such explanations. Experiment 2 used selective attention to a subset of display stimuli in order to isolate the contributions of individual relative disparities, and Experiment 3 assessed the contributions of horizontal disparity components.

4. Experiment 2

In Experiment 1 disparity direction affected perceived depth differently in Transitive and Non-Transitive displays. Experiment 2 attempted to determine how pairwise disparity relations contributed to this result. We instructed observers to judge the depth of a relevant subset of stimuli and to ignore the irrelevant stimuli. This allowed us to compare depth judgments across Transitive and Non-Transitive displays that differed only by the disparity direction of their irrelevant stimuli. If irrelevant stimuli do not affect observers' judgments of the depth of relevant stimuli, we would conclude that the disparity values of individual stimuli within the display are represented and available for perceptual processing independently of other disparity values present in the display. This result would suggest that the inconsistency between the data for Transitive and Non-Transitive displays in Experiment 1 had its source in a decision-level resolution of conflicting non-transitive depth relations, rather than a representational-level interference between individual disparity values. To preview, the results turned out quite differently.

4.1. Methods

4.1.1. Stimuli

Displays were composed of a central grating and four surrounding plaids configured in an 'X' pattern (Fig. 7). The relevant stimulus subset consisted of a grating and one of the diagonally-arranged pairs of plaids. Observers were instructed to attend to and judge the grating and the relevant plaids and to ignore the irrelevant plaids. In the 5-stimulus displays, the relevant stimuli formed a symmetrical region centered on the intended point of fixation. In light of the observer's task in this experiment, this configuration has an advantage over the 3-stimulus displays used in Experiment 1, for it eased concerns about adherence to the attentional instructions. The 3-stimulus displays would have required attending to a region lateral to fixation, which is more difficult subjectively and might lead to fixation compromises.[2]

The displays fell into the two main categories: Transitive and Non-Transitive. All four plaids of the Transitive displays had identical disparities. Non-Transitive displays contained plaids on one diagonal that had a disparity direction orthogonal to those on the other diagonal; the disparities were the same within a diagonal.

Both Transitive and Non-Transitive displays could be further classified into Parallel and Orthogonal sub-classes. All the plaids in Parallel Transitive displays had the same disparity direction as the grating. All the plaids in Orthogonal Transitive displays had the same direction, which was orthogonal to the disparity of the grating. Stimulus relevance distinguished the two Non-Transitive conditions: In Parallel Non-Transitive displays, the grating and the relevant plaids had the same disparity direction, while in Orthogonal Non-Transitive displays the grating and the relevant plaids had perpendicular disparity directions. There were, in addition, two control conditions with just three stimuli, one grating and two plaids. Their purpose was to measure the effect of the presence of irrelevant plaids. The two plaids were identical to the relevant plaids in 5-stimulus displays, with directions either parallel or orthogonal to that of the grating. All disparity directions appeared equally often at each display position, balancing vertical disparities across retinal locations and canceling depth signals that might arise from interactions between disparities and locations (Matthews et al., 2003).

---

[2] We nonetheless used 3-stimulus versions of the displays to replicate Experiment 2 with three of the four observers, obtaining data that were the same in all important respects to those reported here.

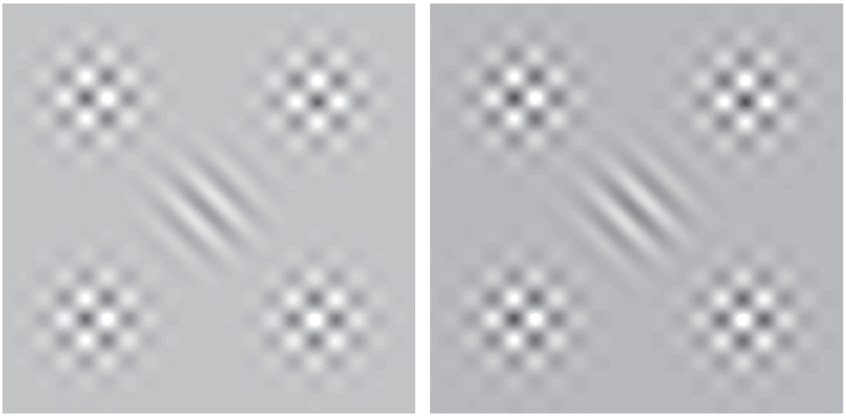

Fig. 7. Monocular view of a typical display in Experiment 2. The central grating was oriented at +45° or −45° and had variable disparity. The two plaids along each diagonal had identical disparities. The grating and two identical plaids were designated as task-relevant for the duration of a block of trials, to be attended and their depth order judged. The remaining plaids were irrelevant and to be ignored.

For three of the observers, the disparity of the plaids, whether relevant or irrelevant, had a direction of either +45° or −45° and a magnitude of 1.67 arcmin (equivalent to a phase disparity of 20° for the grating and the components of plaid). For the remaining observer, the plaids had a disparity direction of either +135° or −135° and a magnitude of 1.25 arcmin (equivalent to a phase disparity of −15°—negative to match the polarity of the central grating with an identical disparity vector). Plaid disparity values were constant within blocks of trials. The grating had an orientation of 45° or 135° and a disparity magnitude that varied across trials.

The center-to-center distance between the grating and a plaid was 2.5° of visual angle and the horizontal and vertical inter-plaid spacing was just over 3.5°. Unlike Experiment 1, the contrast of the grating, 0.1, equalled the contrast of each 1-D component of the plaids, giving the plaid double the contrast of the grating. The grating was centered on the fovea in this experiment; its salience was approximately matched to that of the more peripheral plaids at this new contrast ratio.

4.1.2. Procedure

Observers were instructed to judge the central grating as “Near” or “Far” relative to the pair of relevant plaids flanking it. They were informed at the start of each block of trials which diagonal pair of plaids was relevant; this pair remained relevant throughout the block of trials, in order to discourage switching errors and hysteresis effects that might accompany trial-to-trial selection of relevant stimuli. Each diagonal pair was relevant equally often. Observers were made aware that the relevant plaids had the same disparity and, though non-contiguous, could be judged as a perceptual unit. As in Experiment 1, only the disparity of the grating and the absolute phases of the stimuli varied across trials within a trial block. The grating disparity was controlled by a constant-stimulus procedure that selected among five disparity values.

After attaining nonius alignment, the observer initiated the experimental trial with a click of a mouse. Stimuli were presented for 174 ms. Responses were made by clicking one of two on-screen buttons labeled “Near” and “Far” that appeared shortly after offset of the stimulus display. The response designated the perceived depth of the grating relative to the relevant plaids. No feedback was given about individual responses.

For each of the four display types, each observer’s data came from 8 to 12 runs of 50 experimental trials, each of which was preceded by a handful of warm-up trials. These runs were divided evenly between displays having gratings with orientations of 45° and 135° and relevant plaids arranged along the major and minor diagonals. The remaining procedural details followed those of Experiment 1.

4.1.3. Plaid-plaid depth judgments

An assumption behind the construction of Non-Transitive displays is that plaids with equal horizontal disparities will appear at equal depths within displays containing 1-D stimuli. To verify this assumption, depth judgments for a pair of plaids were collected using a modified version of the 3-stimulus control display. Methods and results appear in the Supplementary Material (Section S2).

4.1.4. Observers

Of the four observers in this experiment, three were naive as to the purposes of the study. One of these (T1) had run in Experiment 1. The remaining observer was one of the authors.

4.2. Results

Fig. 8 shows PSEs for the 5-stimulus displays of Experiment 2. The plotted values are the grating disparities required to obtain perceptual depth matches between the grating and the relevant plaids. Each panel shows PSEs for the four display conditions defined by display type (Transitive or Non-Transitive) and the relative disparity directions of the grating and the relevant plaids (Parallel or Orthogonal). The plaids against which the grating was judged had a fixed disparity magnitude that was equivalent to that of a grating with a phase disparity of +20°, except in the case of observer T1, for whom the corresponding value was −15°. Data for each of the four observers are shown in a separate panel.

PSEs for Parallel Transitive displays were greater than those for Orthogonal Transitive displays, with PSEs for Non-Transitive displays falling between the two Transitive cases. PSEs for Parallel and Orthogonal Transitive displays differed significantly for all four observers ($p < 0.01$ for T3, T4, and T5; $p < 0.03$ for T1, by Monte Carlo simulations, with similar results by t-tests). The four-plaid Transitive displays were similar to those for the respective two-plaid displays, as seen in Fig. 9. Thus, the presence of irrelevant stimuli had no systematic effect on judgments of relevant Transitive stimuli.

However, the presence of irrelevant stimuli did affect judgments of Non-Transitive displays. Fig. 8 shows that PSEs for the Parallel Non-Transitive displays were smaller than those for Parallel Transitive displays and the reverse was the case for Orthogonal displays ($p < 0.05$ in all cases with the exception of Orthogonal PSEs for observer T1, which were not significantly different,

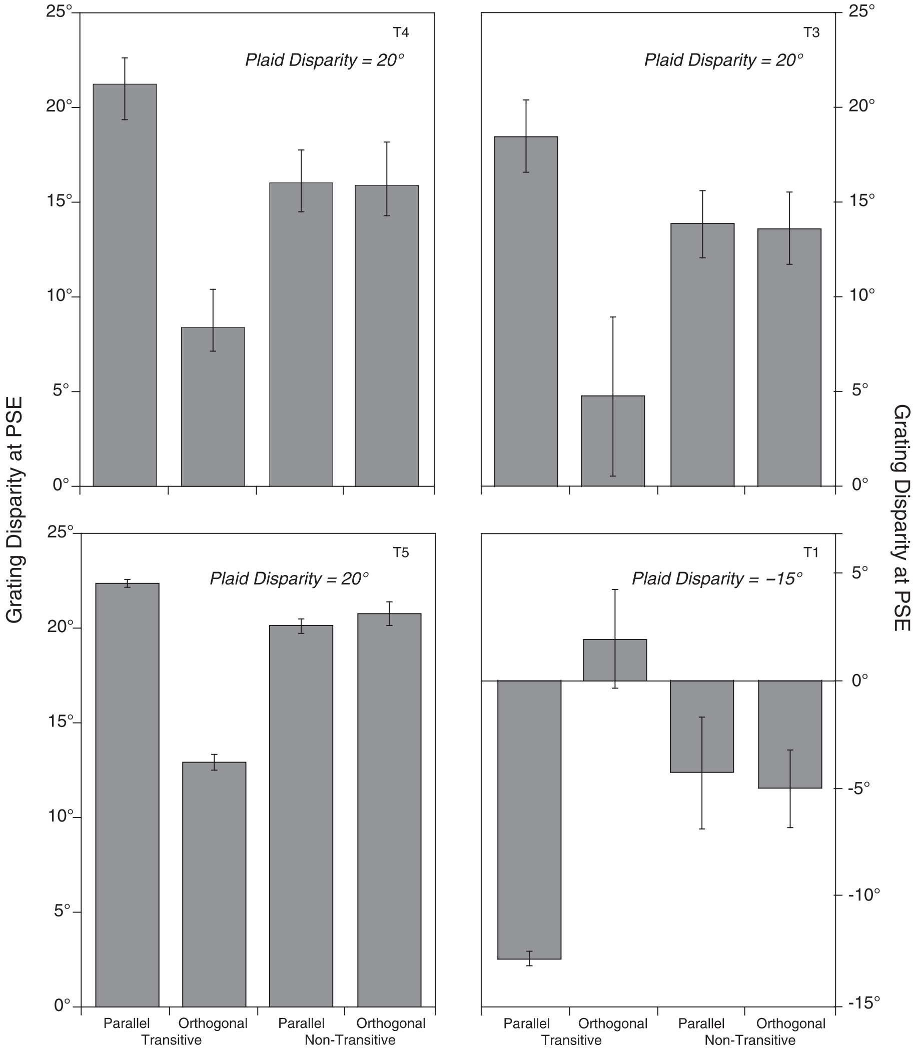

Fig. 8. PSEs for the Transitive and Non-Transitive conditions of Experiment 2. Grating phase disparities that resulted in a perceived depth match between the grating and the relevant plaids are shown for each of the 4 observers. The disparity direction of relevant plaids was either parallel or orthogonal to the grating disparity. Plaids' phase disparities were +20° for observers T3, T4, and T5, and −15° for observer T1. Error bars show ±1 s.e.m.

$p = 0.058$). PSEs for Parallel and Orthogonal Non-Transitive displays were essentially identical for each observer. Thus, the disparity direction of the relevant plaids did not affect depth perception of Non-Transitive displays; swapping parallel and orthogonal plaids as the relevant stimuli conserved observers' depth judgments. This means that attentional selection had no effect. Irrelevant plaids contributed to these judgments as much as relevant plaids did.

Projected disparity predicts PSEs of +20° (or −15°) for Parallel Transitive displays and 0° for Orthogonal Transitive displays. The Orthogonal Transitive data of observers T4 and T5 deviate considerably from zero (Fig. 8). The difference was in the direction of seeing the grating as nearer than expected from its disparity, mirroring results of Experiment 1 and apparent to some extent in most observers' data in all experiments (it is also opposite in direction from the bias we observed in earlier experiments (Chai & Farell, 2009; Farell, Chai, & Fernandez, 2009), where gratings and plaids had a center-surround configuration.) On the other hand, all the observers show a significant difference between PSEs for Parallel and Orthogonal Transitive displays and no difference between the Parallel and Orthogonal conditions of Non-Transitive displays. PSEs for the two Non-Transitive conditions consistently fell between those of the two Transitive conditions. It thus appears that all the plaids, whether attended or unattended, contributed to depth judgments of subsets of Non-Transitive stimuli. As in Experiment 1, the disparities of individual stimuli do not appear to contribute as discrete variables to the decision process.

Given that attention had no effect, a horizontal disparity calculation of depth is consistent with the equality of PSEs for the Parallel and Orthogonal conditions within Non-Transitive displays. This is because the horizontal disparities of all Non-Transitive plaids were the same. However, horizontal disparities were the same in all conditions of the experiment, so they cannot explain the difference in PSEs between Parallel and Orthogonal Transitive displays or between Transitive and Non-Transitive displays. An alternative is that the observed PSEs may reflect the pooling of the individual projected disparities from all plaids in the display. Experiment 3 tested this idea under conditions in which horizontal disparities produce contrasting expectations.

5. Experiment 3

Experiment 3 manipulated the disparity directions of the plaids in 5-stimulus displays. This freed the displays from the constraint operating in the first two experiments whereby all the plaids had the same horizontal disparity. As a result, the PSEs expected if horizontal disparities were the effective signal for depth were distinct from those expected if projections onto a stimulus disparity axis

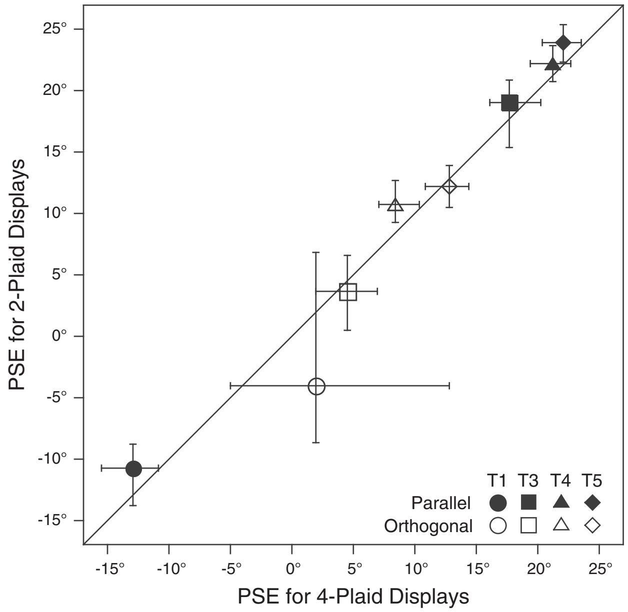

Fig. 9. PSEs for 2-plaid parallel and orthogonal displays, shown with respect to PSEs for 4-plaid Parallel and Orthogonal Transitive displays. The 4-plaid data are those from Fig. 8. Error bars enclose the 95% confidence intervals.

were the signal instead. The observer's task was the same as in Experiment 2.

5.1. Methods

The pair of plaids along one diagonal of the 5-stimulus displays had disparity directions of either +45° or −45°, as in the previous experiment. The other pair had a disparity direction drawn from the set {0°, −15°, +15°, −45°, +45°}, with the latter two values selected to complement the sign of the disparity of the first-mentioned pair. Thus, the two pairs of plaids differed in disparity direction from a minimum of 30° to a maximum of 90°. They shared the same disparity magnitude, 1.67 arcmin (equivalent to 20° phase disparity). The grating was oriented either at 45° or at 135° and had a disparity that varied from trial to trial. In separate blocks of trials each of the diagonal grating-plaid subsets within each display was designated as relevant to the task. This yielded eight pairs (relevant and irrelevant) of plaid disparity directions, two of which were identical to the Non-Transitive displays of Experiment 2. For each of these eight pairs, observers carried out 6 or 8 blocks of 50 trials, plus practice trials, divided among 5 constant stimulus values. The procedure followed that of Experiment 2. Two observers from that experiment, T3 and T5, served in this one.

5.2. Results

Fig. 10 plots the PSEs observed in Experiment 3 against three measures of plaid disparity: (1) the disparities of relevant plaids (top row); (2) the disparities of the irrelevant plaids (middle row); and (3) the pooled disparities of all plaids, with equal weight given to each (bottom row). Each of these disparities was measured after being projected onto two axes, the grating-disparity axis (left column) and the horizontal disparity axis (right column). Each data point gives the mean of PSEs from all displays yielding the corresponding plaid disparity value plotted on the abscissa. Linear functions have been fit through the data points.

The top row of Fig. 10 shows PSEs as a function of the disparities of relevant plaids. The grating-axis projections of relevant plaids (Fig. 10A) produced functions whose slopes are 0.25 and 0.26 for observers T3 and T5, respectively. Only the latter is significantly different from zero ($t(3) = 2.58$ and 4.86, respectively, $p < 0.05$). The slopes for the horizontal-axis components of the relevant plaid disparities (Fig. 10B) were similar, 0.30 and 0.17 for T3 and T5, and not significantly different from zero.

The projected disparities of irrelevant plaids (Fig. 10C) produced functions with slopes of 0.29 and 0.23 for observers T3 and T5, respectively. Neither value differs significantly from zero ($t(3) = 1.21$ and 2.79, $p > 0.05$). Thus, the projected disparities of relevant and irrelevant plaids had practically identical effects on observers' depth judgments. Horizontal-axis components of the irrelevant plaid disparities (Fig. 10D) produced slopes of −0.26 and −0.06 for T3 and T5, respectively, neither of which differs significantly from zero. The differences in slope polarity between relevant and irrelevant horizontal disparities seem attributable to the difference in sign between their correlations (+0.24 and −0.16, respectively) with the pooled projected disparities, which are discussed next.

The abscissa values of the data of the lower left panel of Fig. 10 come from averaging the grating-axis projections of relevant and irrelevant plaid disparities. PSEs are well described as a linear function of these values ($r^2 = 0.96$ for each observer), the slopes of which (0.66 and 0.72) are considerably nearer 1.0 than those derived from relevant or irrelevant plaid disparities separately. Slopes for the combined function are steeper by a factor of about 2.8 than the slope for relevant plaids alone, a difference that is significant for both observers ($t(3) = 3.90$ and 5.22 for T3 and T5, respectively, $p < 0.005$). By contrast, the horizontal disparity of the relevant and irrelevant plaids combined (Fig. 10F) yielded slopes of −0.08 and 0.20, neither of which differs significantly from zero.

To summarize, relevant plaid disparities had only a weak correspondence with observed PSEs, no more strongly predictive of observers' depth judgments than irrelevant plaid disparities. This matches the results of the previous experiment. Pooling relevant and irrelevant plaid disparities was considerably more predictive of PSEs. This pooling gave equal weighting to relevant and irrelevant plaid disparities, consistent with their equivalent individual effects on observers' judgments. The combined projected disparities from relevant and irrelevant plaids underestimated the absolute value and overestimated the gain of the grating disparity required for a depth match. Nevertheless, these matches show that the underlying computation is sensitive to two-dimensional disparity vectors rather than only to their horizontal components and that this sensitivity is global, responding to a pooled disparity measure from all 2-D stimuli in the display, regardless of their task relevance. Thus, the depth perceived between a constituent pair of stimuli, one 1-D and the other 2-D, cannot be predicted from their disparities, for it depends also on the disparities of the other stimuli present in the display.

6. General discussion

We expect transitivity to apply to stereoscopic depth, as we expect it to apply generally. If each of the several stimuli appearing in a stereoscopic display has a known disparity, then we expect that we could rank these disparities to correspond with their depth order as given by perception and that this ranking will be pairwise transitive. Such a correspondence requires parameters linked by a common metric. Horizontal disparity supplies the common metric for perceived depth in most laboratory studies and theories of stereopsis. However, two 1-D stimuli, if they have different orientations and identical horizontal disparities, will have different effective disparities: The two stimuli can appear at different depths relative to the same reference stimulus (Chai & Farell, 2009; Farell, 2006; Farell, Chai, & Fernandez, 2009; Ito, 2005). By dissociating perceived depth from horizontal disparity, 1-D stimuli allow us

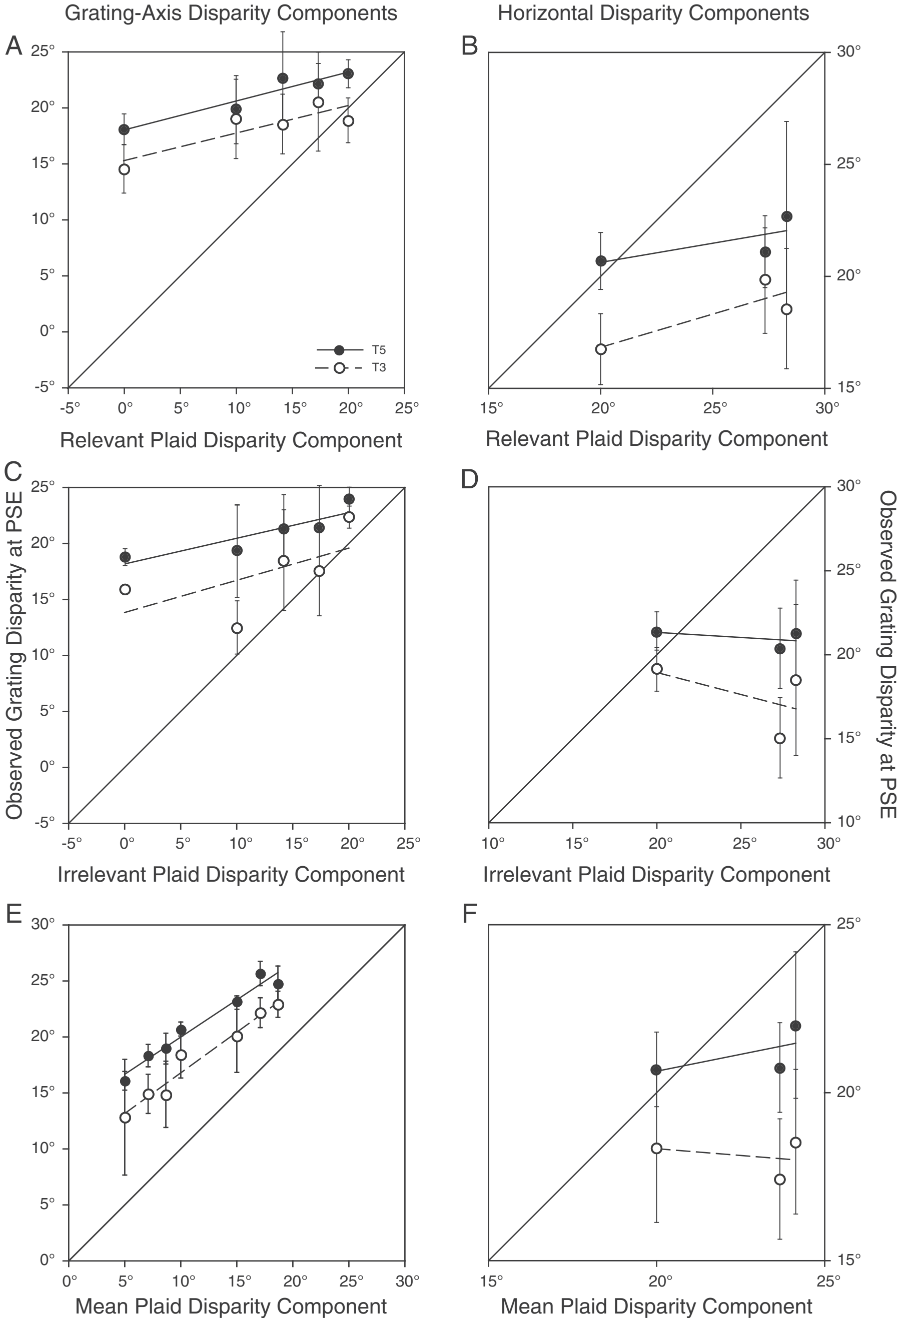

Fig. 10. PSEs as a function of the magnitude of plaids' disparity in the grating's disparity direction (left column) and the horizontal disparity direction (right column) for the two observers in Experiment 3. The plaid disparities plotted are those of the relevant plaids (A and B), the irrelevant plaids (C and D), and the mean of all plaids (E and F). Predicted data points would fall on the diagonal line with slope of 1.0. Error bars are ±1 s.e.m.

to create a set of stimulus pairs whose depth relations, taken together, violate transitivity, despite the normal appearance of each pair when viewed individually. We asked here what depth is seen when the non-transitive pairs are displayed at the same time. Is the depth inconsistency resolved perceptually? If so, how is it resolved, and what depths are seen?

Experiment 1 used displays containing a grating and two plaids, each with an oblique disparity direction, to investigate these questions. In Transitive displays the two plaids had identical disparities and in Non-Transitive displays their disparities were identical in magnitude but directed perpendicularly, each deviating equally from horizontal. If horizontal disparities alone mediated the depth appearance, all the displays of Experiment 1 (and Experiment 2) would have appeared identical in depth, for all the plaids had the same horizontal disparity. If the various display types appeared to differ in depth, then the depth-from-disparity calculation must have been sensitive to disparity direction in two dimensions. But then the violation of transitivity would have to be resolved in novel ways or else be expressed as perceptual ambiguities.

With the plaids' disparities held constant, we measured how perceived depth changed with the grating's disparity magnitude. In Experiment 1 we found that perceived depth changed at the same rate in Non-Transitive displays as in Transitive displays, but psychometric functions were displaced laterally, giving each of the three display types a different PSE. PSEs for Parallel Transitive displays were significantly larger than those for Orthogonal Transitive displays, PSEs for Non-Transitive displays being intermediate. However, data for Non-Transitive displays showed no difference between parallel and orthogonal disparity directions. Thus, these displays showed no evidence of the inconsistent depth relations among their constituent pairs. These same effects were observed in Experiment 2, where only stimulus pairs of one type, parallel or orthogonal, were relevant to the judgment. In fact, the disparity of irrelevant plaids contributed to the judgment of the relevant grating-plaid pair as much as the disparity of the relevant plaids did. Thus, selective attention to stimulus subsets failed to affect judgments of depth. Experiment 3 showed that this result held regardless of the disparity direction difference of the plaids.

While observers resolved the pairwise inconsistencies within Non-Transitive display, their judgments across Transitive and Non-Transitive displays remained inconsistent. In Transitive displays the judged depth between gratings and plaids depended on their relative disparity directions; in Non-Transitive displays, disparity direction had no effect. Thus, the perceived depth of relevant stimuli varied with the context provided by irrelevant stimuli. For instance, one can take a grating that appears in the same depth plane as several plaids and move it, without changing its disparity or orientation, to another display containing plaids with different disparity directions but the same disparity magnitudes and the same horizontal disparity components as in the first display, and the grating will now appear in a different depth plane from the plaids. But if one moves it to a third display in which half the plaids were taken from the first display and half from the second, one can adjust the grating's disparity magnitude so that all the stimuli appear in the same depth plane.

6.1. The effective disparity axis

We favor an account in which the between-display inconsistency in depth judgments arises because the disparity axis used under the conditions of the experiments is not horizontal, but instead is determined by the disparities of the stimuli themselves. This effective disparity axis is used for calculating depth; it is not necessarily the same as or defined by the axis along which binocular correspondences are made. (An analogous case arises when horizontal disparity determines the perceived depth of a stimulus with an oblique disparity direction.) This stimulus-dependent disparity axis varies across the display types used here.

Observers' depth judgments of a grating and a plaid can be approximated by the relative magnitude of disparity components along the disparity axis of either stimulus. One case calculates the orthogonal projection of the plaid's disparity onto the grating's disparity axis, a common axis for comparing the two disparities (Fig. 1). The other case calculates the disparity components in the direction of the plaid's disparity, again allowing the two disparities to be compared along a common axis and yielding the same disparity ratio. Here, the grating's disparity is expressed by the intersection of its constraint line and the plaid's disparity axis. The two calculations are equally consistent with the data of the present experiments and those of previous studies of grating-plaid pairs (Chai & Farell, 2009; Farell, Chai, & Fernandez, 2009).

By bringing a second plaid into the display, Experiment 1 raised the question of the role of pairwise comparisons in multi-stimulus depth judgments. Non-Transitive displays have a crucial role in exploring this issue. Surprising evidence came in Experiments 2 and 3 from the failure of attention to select between relevant and irrelevant plaid disparities, even those with orthogonal directions. This shows that the task of comparing the depths of a grating and a plaid may engage more disparities than those of two stimuli. The disparities of all the plaids present in the display appear to be pooled. These plaids were iso-eccentric and identical except for position, phase, and disparity direction in our experiments, leaving unanswered questions about the generality of the pooling and the possibility of attentional selection among less similar stimuli.

Judging relative depth along a stimulus disparity axis is related by a simple rotation to the standard depth calculation from horizontal disparities. If observers in our experiments compared disparities projected onto the grating's disparities axis, then a global rotation of the display that brings the grating to vertical makes its disparity axis horizontal and allows the disparity comparisons to be expressed in terms of horizontal disparity components. Thus, the grating disparity axis in the unrotated display takes on the role usually assigned to the horizontal. The same applies to a single-plaid disparity axis or to the average of the plaid disparity axes, if either of these were used instead of the grating's axis as the reference disparity direction. We can use the term *relative disparity axis* (RDA) to refer to the axis on which disparity components are compared to yield relative depth judgments.

The RDA can be horizontal and usually is. A shift from the horizontal direction to a different, stimulus-defined direction in the presence of 1-D stimuli has an analog in spatial-domain reference stimuli. Contextual stimuli, even if irrelevant, can shift the direction of maximum stereo sensitivity away the fixation plane and toward a local stimulus-defined reference plane (Glennerster & McKee, 1999; Glennerster, McKee, & Birch, 2002; Mitchison & Westheimer, 1984; Petrov & Glennerster, 2006; Westheimer, 1979). But while it is easy to imagine benefits in calculating stereo depth with respect to a local reference plane, the advantages of calculating depth from non-horizontal disparities appear elusive. Binocular matching must have some tolerance for vertical disparities to accommodate their occurrence in the viewing environment, but this does not imply a role for them in the depth-from-disparity computation. So, why does the visual system not use horizontal disparities when judging the depth of 1-D stimuli? The lack of parsimony and especially the lack of veridicality of the resulting depth estimates argue against specialized mechanisms dedicated to the analysis of disparities along 1-D-stimulus-defined directions rather than horizontal disparities. One alternative is that our data might be an artifact of presenting conventional mechanisms with artificial stimulus parameters that have no, or highly improbable, real-world correlates. Another possibility is that the stimulus-defined RDA might be the general case. The veridicality of stereo depth estimates would then be an index of how closely the RDA approximates the horizontal. The approximation would depend on how close the typical or average disparity in the viewed scene is to the horizontal, with 1-D and 2-D stimuli and possibly other factors contributing different weightings to the average.

6.2. Salience, vertical disparities, and attention

Mitchison and Westheimer (1984) defined salience as the relative disparity between a stimulus and its neighbors with respect to a common reference (the fixation plane, in their calculations). Stimulus separation contributes via a weighting factor. The salience metric makes perceived depth context dependent, with neighboring stimuli contributing a disparity-contrast signal. This contrast signal depends on disparity and position alone, so salience discounts modulation by other factors—by attention, for example—and gives all nearby stimuli a role in any particular depth judgment. The equivalence of relevant and irrelevant stimuli in our studies fits this description well. For our displays, however, it is clear that disparity calculations do not use such a context-independent reference as the fixation plane. In addition to contributing relative disparity signals, contextual stimuli have effects equivalent to changing the direction along which disparities are compared.

The disparity axes used for calculating the depth between 1-D and 2-D stimuli are stimulus disparity axes. Horizontal axis values make no independent contribution. By contrast, disparity components along both cardinal directions contribute to depth calculations of 2-D stimuli, with vertical disparities taking a modulatory role. For instance, vertical disparities can calibrate horizontal disparities by providing an estimate of ocular viewing parameters (Gillam, Chambers, & Lawergren, 1988; Mayhew & Longuet-Higgins, 1982). Measurable effects of calibration require stimuli much larger than those presented here, and even then only partially compensate for effects of ocular posture (Rogers & Bradshaw, 1993). However, calibration implies that vertical disparities should be pooled to arrive at a global value (Bradshaw, Glennerster, & Rogers, 1996; Garding et al., 1995; Kaneko & Howard, 1997; Rogers & Bradshaw, 1993; Stenton, Frisby, & Mayhew, 1984; but see Rogers & Koenderink, 1986) and this has potential importance for interpreting the inability of observers to selectively process relevant stimuli. Perhaps the presence of vertical disparities triggered a pooling of disparity signals that overrides attentional partitioning. But a pooling of vertical disparities would still leave horizontal disparities, whether calibrated or not, as the principle factor contributing to, if not determining, perceived depth. Yet horizontal disparities did not make a significant contribution, so there is little reason to attribute the failure to selectively process relevant stimuli to the presence of vertical disparities in the display.

Stimulus position is a prime cue for attentional selection (e.g., Posner, 1980), yet selection failed in our experiments despite the static segregation of relevant and irrelevant positions throughout trial blocks. Attentional selection in our displays was not limited by the particular parameters of the plaids, for a single crucial change in a different parameter enables selective processing: When a plaid is substituted for the grating, the disparity of the irrelevant plaids has no effect on PSEs (Farell & Ng, 2014; see Supplement Section S3). Thus, the failure to select among the plaids in the present experiments is not due to an inherent property of plaid processing (e.g., coherence of plaid disparities) or of attention (e.g., a resolution too coarse to resolve individual plaid disparities). Instead, the failure of selection seems specific to the computation of depth between 1-D and 2-D stimuli.

6.3. *Offset and gain*

Horizontal-axis disparity projections do not accurately predict depth judgments in our experiments. Projected disparities do better, giving us a metric that is responsive to the sign of vertical disparity and linearly related to PSEs (Fig. 10E). However, PSEs differed systematically from the values expected from stimulus-axis components. They showed an overall displacement toward higher-than-predicted grating disparity values and a less-than-unity slope, with a value of only about 0.7 (Fig. 10E). A bias to see gratings as near relative to plaids, other things being equal, might have arisen because of a difference in retinal eccentricity or in effective contrast.

We had attempted to remove extraneous reference stimuli in order to require observers to judge the stimuli by comparing them directly with each other rather than indirectly via extraneous reference stimulus (see Exp. 1 *Methods*, Section 3.1.1). Yet there are remaining factors that could make projected disparities inexact proxies for the disparity metric that actually supports task performance. For example, the disparity-gradient limit (Burt & Julesz, 1980) might show directional selectivity. That is, the separation of stimuli within a display, in terms either of visual angle or carrier wavelengths, might have been sub-optimal for accurate coding of disparities with strongly divergent directions. However, the perceived depth separation between *plaids*, whether these are laterally adjacent (Farell, Chai, & Fernandez, 2010) or separated by another stimuli (Supplement Fig. S2), is consistent with a horizontal disparity metric despite a considerable difference in two-dimensional disparity direction. This does not suggest a disparity gradient limit.

Another approach to understanding the discrepancy between data and prediction is to consider a more nuanced specification of the grating disparity axis, one less rigidly tied to the nominal orientation of the stimulus. The nominal orientation is just the center of the grating's orientation band. Thus, there is an orthogonal disparity-axis band with a comparable width. Merging this line of thought with a loosening of assumptions, we can entertain hypotheses about which of the disparity directions within this band are used when an observer compares the depths of a pair of stimuli. One hypothesis is that stimulus disparity components are compared preferentially along an axis that minimizes the difference in disparity directions of the two stimuli. This inverts the logic of off-frequency listening in auditory detection and off-frequency looking in visual detection, and applies it in the disparity domain for the purpose of facilitating comparisons rather than escaping masking. Thus, for this purpose of comparing disparities, nominal and effective disparity directions might differ, with the effective directions being more similar across stimuli than the nominal directions. The size of the difference would vary with bandwidth and the difference in nominal values, following a sine function. It would be largest when the nominal difference was greatest (90°) and would contribute, possibly, to the less-than-expected gain in grating depth observed here. Both the disparity gradient and the direction band hypotheses are readily testable by direct measurement.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant Number NSF BCS-1257096. The authors gratefully acknowledge the contributions of the two anonymous reviewers, whose sharp-eyed and critical reading and numerous suggestions helped remove burrs and clarify bristly points.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.visres.2014.10.012.

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, 1A*, 1266.

Blake, R., Camisa, J., & Antoinetti, D. N. (1976). Binocular depth discrimination depends on orientation. *Perception and Psychophysics, 20*, 113–118.

Bradshaw, M. F., Glennerster, A., & Rogers, B. J. (1996). The effect of display size on disparity scaling from differential perspective and vergence cues. *Vision Research, 36*, 1255–1264.

Burt, P., & Julesz, B. (1980). A disparity gradient limit for binocular fusion. *Science, 208*, 615–617.

Calabro, F. J., & Vaina, L. M. (2006). Stereo motion transparency processing implements an ecological smoothness constraint. *Perception, 35*, 1219–1232.

Chai, Y.-C., & Farell, B. (2009). From disparity to depth: How to make a grating and a plaid appear in the same depth plane. *Journal of Vision, 9*, 1–19.

Cumming, B. C. (2002). An unexpected specialization for horizontal disparity in primate visual cortex. *Nature, 418*, 633–636.

Delicato, L. S., & Qian, N. (2005). Is depth perception of stereo plaids predicted by intersection of constraints, vector average or second-order feature? *Vision Research, 45*, 75–89.

Durand, J., Celebrini, S., & Trotter, Y. (2007). Neural bases of stereopsis across visual field of the alert macaque monkey. *Cerebral Cortex, 17*, 1260–1273.

Ebenholtz, S., & Walchli, R. (1965). Stereoscopic thresholds as a function of head- and object-orientation. *Vision Research, 5*, 455–461.

Farell, B. (1998). Two-dimensional matches from one-dimensional stimulus components in human stereopsis. *Nature, 395*, 689–693.

Farell, B. (2006). Orientation-specific computation in stereoscopic vision. *Journal of Neuroscience, 26*, 9098–9106.

Farell, B., & Ahuja, S. (1996). Binocular disparity of 1- and 2-D contours. *Investigative Ophthalmology and Visual Science, 37*, S284.

Farell, B. & Ng, C. J. (2014). Attention to pattern depth depends on pattern dimensionality. In *Visual sciences society conference*, St. Petersburg, FL, May, 2014.

Farell, B., Chai, Y.-C., & Fernandez, J. M. (2009). Projected disparity, not horizontal disparity, predicts stereo depth of 1-D patterns. *Vision Research, 49*, 2209–2216.

Farell, B., Chai, Y.-C., & Fernandez, J. M. (2010). The horizontal disparity direction vs. the stimulus disparity direction in the perception of the depth of two-dimensional patterns. *Journal of Vision, 10*(25), 1–15.

Farell, B., & Li, S. (2004). Seeing depth coherence and transparency. *Journal of Vision, 4*, 209–223.

Fennema, C. L., & Thompson, W. B. (1979). Velocity determination in scenes containing several moving objects. *Computer Graphics and Image Processing, 9*, 301–315.

Friedman, R. B., Kaye, M. G., & Richards, W. (1978). Effect of vertical disparity upon stereoscopic depth. *Vision Research, 18*, 351–352.

Garding, J., Porrill, J., Mayhew, J. E. W., & Frisby, J. P. (1995). Stereopsis, vertical disparity and relief transformations. *Vision Research, 35*, 703–722.

Gillam, B., Chambers, D., & Lawergren, B. (1988). The role of vertical disparity in the scaling of stereoscopic depth perception: An empirical and theoretical study. *Perception and Psychophysics, 44*, 473–483.

Glennerster, A., & McKee, S. P. (1999). Bias and sensitivity of stereo judgements in the presence of a slanted reference plane. *Vision Research, 39*, 3057–3069.

Glennerster, A., McKee, S. P., & Birch, M. D. (2002). Evidence for surface-based processing of binocular disparity. *Current Biology, 12*, 1–20.

Howard, I. P., & Rogers, B. J. (2002). *Seeing in depth* (Vol. II). Toronto: I. Porteous.

Ito, H. (2005). Illusory depth perception of oblique lines produced by overlaid vertical disparity. *Vision Research, 45*, 931–942.

Kaneko, H., & Howard, I. P. (1997). Spatial properties of shear disparity processing. *Vision Research, 37*, 2871–2878.

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.

Matthews, N., Meng, X., Xu, P., & Qian, N. (2003). A physiological theory of depth perception from vertical disparity. *Vision Research, 43*, 85–99.

Mayhew, J. E. W., & Longuet-Higgins, H. C. (1982). A computational model of binocular depth perception. *Nature, 297*, 376–378.

Mitchison, G. J., & Westheimer, G. (1984). The perception of depth in simple figures. *Vision Research, 24*, 1063–1073.

Morgan, M. J., & Castet, E. (1997). The aperture problem in stereopsis. *Vision Research, 39*, 2737–2744.

Ogle, K. N. (1955). Stereopsis and vertical disparity. *Archives of Ophthalmology, 53*, 495–504.

Patel, S. S., Bedell, H. E., & Sampat, P. (2006). Pooling signals from vertical and non-vertically orientation-tuned disparity mechanisms in human stereopsis. *Vision Research, 46*, 1–13.

Patel, S. S., Ukwade, M. T., Stevenson, S. B., Bedell, H. E., Sampath, V., & Ogmen, H. (2003). Stereoscopic depth perception from oblique phase disparities. *Vision Research, 43*, 2479–2792.

Pelli, D. G., & Zhang, L. (1991). Accurate control of contrast on microcomputer displays. *Vision Research, 31*, 1337–1350.

Petrov, Y., & Glennerster, A. (2006). Disparity with respect to a local reference plane as a dominant cue for stereoscopic depth relief. *Vision Research, 46*, 4321–4332.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32*, 3–25.

Quaia, C., Sheliga, B. H., Optican, L. M., & Cumming, B. G. (2013). Temporal evolution of pattern disparity processing in humans. *The Journal of Neuroscience, 33*, 3465–3476.

Rogers, B. J., & Bradshaw, M. F. (1993). Vertical disparities, differential perspective and binocular stereopsis. *Nature, 361*, 253–255.

Rogers, B., & Koenderink, J. (1986). Monocular aniseikonia: A motion parallax analogue of the disparity-induced effect. *Nature, 322*, 62–63.

Stenton, S. P., Frisby, J. P., & Mayhew, J. E. W. (1984). Vertical disparity pooling and the induced effect. *Nature, 309*, 622–623.

Westheimer, G. (1979). The spatial sense of the eye. *Investigative Ophthalmology and Visual Science, 18*, 893–912.

Wichmann, F. A., & Hill, N. J. (2001). The psychometric function: II. Bootstrap-based confidence intervals and sampling. *Perception and Psychophysics, 63*, 1314–1329.