Spatial Vision, Vol. 15, No. 1, pp. 45–59 (2001)

Occlusion junctions do not improve stereoacuity

RAYMOND VAN EE [1,2,*], BARTON L. ANDERSON [1] and HANY FARID [1]

[1] *Department of Brain and Cognitive Sciences, Massachussetts Institute of Technology, USA*
[2] *Utrecht University, Helmholtz Institute, Princeton Plein 5, 3584 CC, Utrecht, The Netherlands*

Received 20 September 2000; revised 24 December 2000; accepted 1 January 2001

Abstract—Occlusion geometry gives rise to interocular shifts in the positions of binocularly viewed contour junctions. Since these shifts do not give rise to normal binocular disparities, they have been called 'pseudodisparities'. Previous work has shown that the unmatched contour segments of a partially occluded contour at occlusion junctions can be used to recover the geometry of the occluding surface through the construction of 'illusory' contours. Here, experiments were performed to determine whether such junction shifts could enhance stereoscopic depth detection when the relative disparity between the contours was below threshold. Our results showed that stereoscopic depth detection does not improve when pseudodisparity is present. We conclude that the visual system is less sensitive to pseudodisparity than to conventional disparity information. We suggest that the primary role of pseudodisparity is to overcome conditions of camouflage.

Keywords: Binocular vision; stereopsis; stereoacuity; occlusion.

1. INTRODUCTION

The horizontal displacement of our two eyes gives rise to systematic spatial differences in the projected images of a three-dimensional (3D) scene. There are two broad kinds of image features that are generated in binocular viewing: features visible to both eyes, and features visible to only one eye. Features visible to both eyes give rise to binocular disparity. Since the invention of the stereoscope by Wheatstone (1838), it has been known that disparity can generate vivid percepts of three-dimensional structure. The vast majority of research conducted on stereoscopic processing has focused on understanding how the visual system senses and uses binocular disparity to recover a three-dimensional representation of a scene. However, not all regions of a scene project to both eyes. Monocularly visible scene fragments can arise from partial occlusion of

*To whom correspondence should be addressed. E-mail: r.vanee@phys.uu.nl

surfaces that intersect along a depth discontinuity. Indeed, in many natural viewing conditions, both kinds of features can be present in adjacent image regions.

Occlusion junctions generated by the crossing of features in different depth planes can be shifted in a variety of different directions in the two eyes' images (Anderson and Julesz, 1995). Malik, Anderson, and Charowhas (1999) referred to these junction shifts as pseudodisparity, and derived a quantitative expression relating the shift in junction position to the orientations and relative disparity of the component contours. In other words, pseudodisparities can theoretically be used to recover the depth of a feature relative to the depth of a second feature that partly occludes the first feature. The prefix 'pseudo' is intended to connote that these junctions do not correspond to the same surface regions in the projected 3D scene. It stresses that the monocularly visible regions that are brought about by the junctions do not give rise to corresponding features and consequently they do not give rise to normal disparities.

Thus, the presence of pseudodisparity in binocular images is a geometric fact, but its role in visual processing has been explored in only a limited number of geometric contexts. Anderson (1994) and Anderson and Julesz (1995) used a single contour stimulus to show that the presence of a monocularly visible contour segment can be used to construct an 'illusory' (Gulick and Lawson, 1976) occluding contour. The same stimulus was used by Malik *et al.* (1999), who demonstrated that the perceived depth and orientation of the illusory occluding contour could be predicted by the geometric equations relating pseudodisparity to the orientations and depths of the component contours. In other words, monocular features along a partially occluded contour provide information that the visual system uses to construct occlusion geometry when the occluding contour does not generate any image contrast in either eye (i.e. when it is camouflaged). It remains unclear if monocularly visible contour segments play a more general role in reconstructing scene geometry, or whether the primary function of these features is to overcome camouflage.

The ability to utilise pseudodisparities for the formation of illusory contours implies that the visual system must possess mechanisms that allow it to distinguish between image regions that are generated by a common source in the world, and those that arise from multiple surfaces that intersect along occlusion boundaries. This problem cannot be resolved on the basis of purely local spatial information. An example illustrating this ambiguity is depicted in Fig. 1. In order to overcome this ambiguity, the visual system must take into account global stimulus features. In the studies described herein, a zero-disparity random-dot surround is presented to define epipolar lines specifying the direction in which globally correct matches are present. In principle, these epipolar lines define which features should and should not be matched. Figure 2 illustrates that the corresponding features are horizontally shifted in the two half-images. In the discussion that follows, we will use the term disparity to refer to the shifts in the plane of the half-images; these lie in an epi-polar plane and are purely horizontal. Note that the unmatched contour segments in this image are accompanied by a junction shift with a vertical component. This suggests

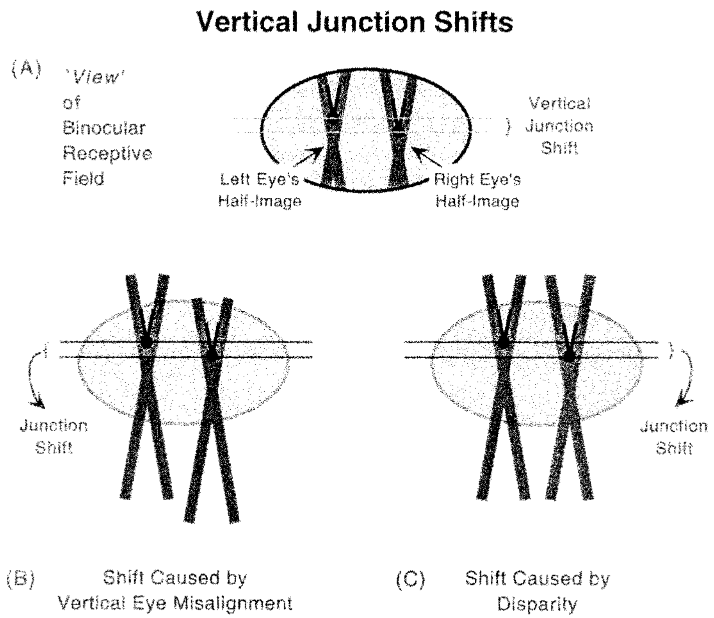

Figure 1. The identification of corresponding features, and thereby the specification of disparity is not a trivial matter. Assume that a binocular receptive field signals the presence of two junctions that have a relative vertical shift (A). The identification of corresponding features in panel (A) is ambiguous if we disregard global features. The reason for the ambiguity is that the input to the binocular receptive field can be created by two, globally different, stimulus configurations. Panel (B) demonstrates that the receptive field input can be created by vertical ocular misalignment without involving a real depth between the bars. This type of shift is caused by a single shifted 'X'. Panel (C) illustrates that the receptive field input can also be created by a junction shift that involves a relative depth. Thus, in order to identify corresponding features, global stimulus features must be taken into account.

that these local, vertical shifts may provide critical information that the visual system uses to distinguish between interocular displacements caused by occlusion junctions, and those caused a single, common surface (cf. Anderson and Julesz, 1995; Malik *et al.*, 1999).

In this paper, we investigated whether depth detection benefits from pseudodisparities. More specifically, the experiments described herein were conducted to determine whether the presence of pseudodisparities can cause a subthreshold disparity difference to become perceptually salient. It is worth stressing that a demonstration of the possible beneficial effect of pseudo-disparity requires that the normal binocular disparity between features is so small (subthreshold) that it does not evoke a perceived depth difference. In other words, to isolate the effect of pseudodisparity on depth detection, normal disparity should be below detection threshold.

There exist numerous environmental situations wherein the magnitude of pseudodisparity will be larger than the magnitude of disparity, and hence could potentially serve as a means of resolving depth differences that could not be detected by conventional disparity mechanisms. To see this, consider a natural viewing

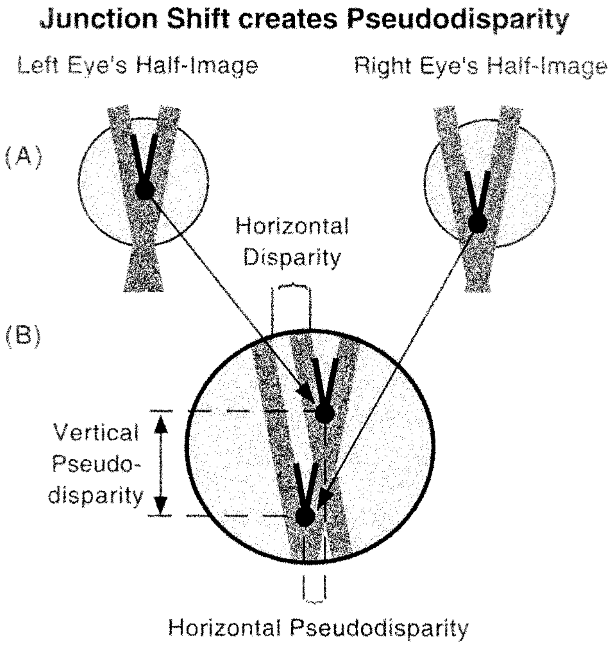

Figure 2. In panel (A), a junction is highlighted by a dot. In panel (B), the two views of panel (A) are superimposed. The two bars are horizontally shifted relative to one another in the two half-images of the stereogram. This horizontal shift is known as binocular disparity. The junctions are both horizontally and vertically shifted relative to one another. Such junction shifts have been coined pseudodisparity. The magnitude of pseudodisparity depends on the orientation of the features. In our example, the vertical junction shift is much larger than the horizontal shift, and the horizontal shift is half the horizontal disparity.

condition in which an observer is looking at two trees at a distance of 20 m. Two branches that differ in depth by 20 cm create a horizontal disparity of 6.6 arcsec. For most untrained observers, this value is about half the threshold of disparity processing, and no depth difference would be experienced on the basis of disparity (cf. Howard and Rogers, 1995). However, if one of the branches is almost vertical and the other has a relative orientation that differs from the first branch by 11 deg, the vertical pseudodisparity will be as large as 30 arcsec. This magnitude is about twice the horizontal disparity detection threshold. In such situations the detection of relative depth might benefit from the processing of pseudodisparities. More generally, the quantitative relationship between pseudodisparity, relative disparity, and the orientations of the occluding and occluded contours is given by (Malik *et al.*, 1999):

$$\text{pseudodisparity}_{\text{Horizontal}} = \frac{-\delta \cdot \cos\alpha \cdot \sin\beta}{\sin(\alpha - \beta)}, \tag{1}$$

$$\text{pseudodisparity}_{\text{Vertical}} = \frac{-\delta \cdot \sin\alpha \cdot \sin\beta}{\sin(\alpha - \beta)}, \tag{2}$$

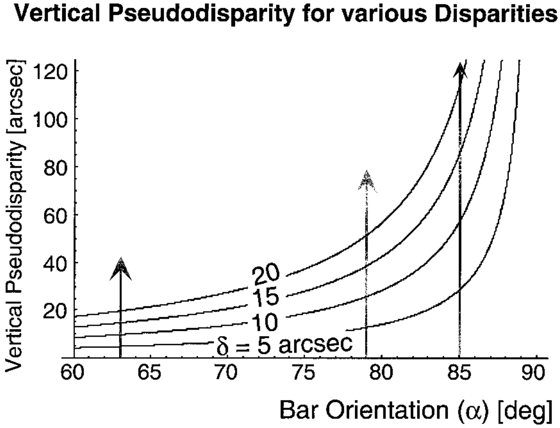

Figure 3. According to equation (4), the vertical pseudodisparity depends both on the bar orientation (α) and the relative disparity (δ) between the bars. This graph illustrates how the vertical pseudodisparity varies with bar orientation for a family of relative disparities. The gray arrows highlight the orientations (63, 79 and 85 deg) where we conducted threshold measurements. Most of the experiments described in this paper were conducted at a bar orientation of 79 deg. The thresholds that we found in the experiments ranged from 5 to 20 arcsec.

where δ denotes the horizontal disparity (see Fig. 1C) and α are β denote the bar orientation in degrees (see Fig. 4B). For the bar configuration that we used in our occlusion conditions β equaled $180 - \alpha$ and the equations of Malik *et al.* yield:

$$\text{pseudodisparity}_{\text{Horizontal}} = \frac{1}{2} \cdot \delta, \tag{3}$$

$$\text{pseudodisparity}_{\text{Vertical}} = \frac{\tan \alpha}{2} \cdot \delta. \tag{4}$$

Figure 3 illustrates how vertical pseudodisparity varies with orientation (α) for a family of relative disparities (δ). Note that the range of possible vertical pseudodisparities can be large: when the orientations of the half-occluding bars approach horizontal, vertical pseudodisparity approaches zero. But when the orientations of the two bars approach vertical, vertical pseudodisparity approaches infinity.

2. EXPERIMENT 1: THE ROLE OF ORIENTATION

2.1. Methods

2.1.1. Observers. Four subjects took part in the experiments. The subjects had normal (observers BA, DS, MS) or corrected-to-normal (JE) vision. In Experiment 1, subjects DS, JE and MS participated. In Experiment 2, BA, JE and MS participated. In Experiment 3, JE and MS participated. All observers were naïve

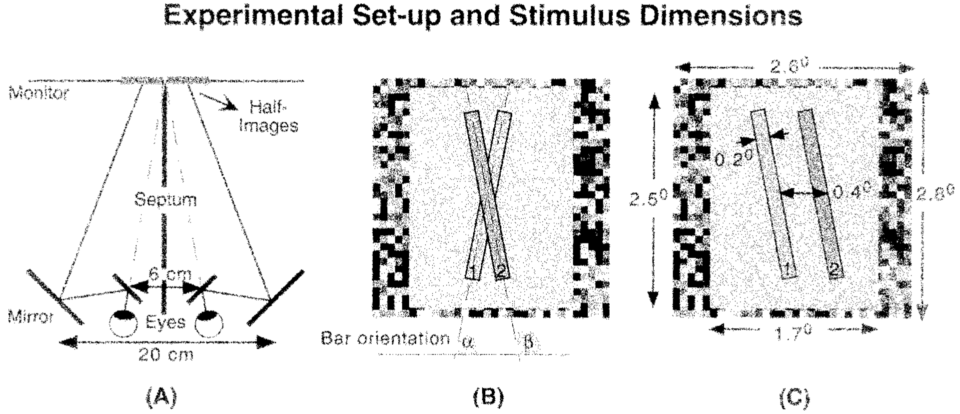

Figure 4. Panel (A) illustrates a top-view of the experimental set-up. The stereogram half-images were presented side by side on a monitor. The viewing distance was 100 cm. The mirror set-up and a septum were used to make the half-images visible to only one eye. In each session we presented occlusion (B) and non-occlusion (C) configurations. α and β denote the orientation of bar 1 and 2, respectively. Both bar 1 and bar 2 always had opposite orientations ($\alpha = 180 - \beta$; $\alpha < 90$) in the occlusion configuration, and identical orientations ($\alpha = \beta$; $\alpha > 90$) in the non-occlusion configuration. Bar 2 is the target bar whose disparity was varied depending on the subject's performance in a 2AFC procedure.

with respect to the purposes of the experiment except for observer BA (one of the authors).

2.1.2. Apparatus. Observers viewed stereogram half-images on a Radius high-resolution monitor at a viewing distance of 100 cm. A pixel subtended 40 × 40 arcsec. The half-images of the stereogram were presented side by side on the monitor. The half-image of one eye was made invisible to the other eye by using a mirror set-up and a black septum between the center of the stereogram on the monitor and the forehead. Figure 4A depicts a top-view of the experimental set-up. The room was dimly lit. The head was stabilized with a chin and forehead rest.

2.1.3. Stimuli. Two classes of stimuli were used. In the first (occlusion) configuration, the bars overlapped, creating occlusion junctions (Fig. 4B); in the second (control) configuration, the bars did not overlap (Fig. 4C). The bar width and height subtended 0.2 deg and 2.2 deg, respectively. Each half-image of the stereogram subtended 1.7 × 2.5 deg. The surrounding random dot pattern provided a frame of reference for the stabilization of vergence and cyclovergence and defined epi-polar lines for establishing correspondence. The random dot pattern subtended 2.6 × 2.8 deg and its noise elements subtended 6 × 6 arcmin. The angular separation between the centers of the two half images was 3.25 deg, and the half-images were separated by a black divider subtending 0.65 deg. Since stereoacuity based on horizontal disparity degrades at lower luminance contrast (Cormack *et al.*, 1991), the bars were displayed at reduced contrast so as to avoid ceiling effects.

Specifically, the luminance of the bars and the background was 20 and 65 cd/m^2, respectively (Michelson contrast was 0.53). The bars were presented with a pedestal disparity of 2.5 arcmin, causing them to be perceived in front of the surrounding noise pattern. The purpose of this pedestal disparity was to prevent subjects from judging the depth of the target bar relative to the noise pattern, rather than between each other. In experiment 1, the bar orientations were varied across trials from 20 to 85 degrees with respect to horizontal. The shape and texture of the bars were varied in experiments 2 and 3. The angular separation between the bars in the control (parallel) condition was 0.4 deg. At this separation there is no depth attraction or substantial depth repulsion between the bars (Westheimer and Levi, 1987).

Using antialiasing techniques, we were able to construct sub-pixel disparity steps of 1/8 of a pixel (5 arcsec). More specifically, in order to create an N $\times$ N image shifted by 1/8 of a pixel, the image was first created at a resolution of 8N $\times$ 8N. This higher resolution image was shifted by one pixel, and then repeatedly low-pass filtered (blurred) and sub-sampled by a factor of two producing the desired N$\times$N image after three iterations. The one-pixel shift of the higher resolution image results in an effective shift of 1/8 of a pixel in the final N $\times$ N image. Similarly, if the higher resolution image is shifted by four pixels, then the effective shift in the final image is 1/2 of a pixel.

2.1.4. Task and procedure. As explained in the Introduction, the insight motivating the experiments is that to isolate the effect of pseudodisparity on depth detection, normal disparity should be below detection threshold. A staircase procedure was used to determine the threshold for perceiving depth differences between the two patterns. We measured the smallest disparity that subjects were able to detect as a depth difference for bar orientations of 20, 63, 79 and 85 deg. The latter three orientations are highlighted in Fig. 3. The horizontal disparity of the target bar (bar 2 in Fig. 4) was varied using a 1-up/2-down staircase, yielding a threshold that represents the 70.7% correct point on the psychometric function. Criterion problems (bias) were eliminated by employing a successive 2 alternative-forced-choice (2AFC) procedure. When absolute thresholds were measured, observers reported which of the two intervals contained a depth difference between the two bars. When increment thresholds were measured, observers were required to report the interval containing the largest depth step. Step size in the staircases was initially 4/8 pixels, but was reduced to 2/8 pixels after the second reversal, and to 1/8 pixel (5 arcsec) after the fourth reversal. The staircase was terminated after 28 reversals and the threshold was determined from the average of disparities for the last 24 reversals. The occlusion and non-occlusion configurations were presented intermixed within every session. In the first experiment, the bars in both configurations could have one of three different orientations. Staircases for the resulting 6 conditions were randomly interleaved. Both intervals in the 2AFC procedure were presented for 1000 ms. The noise frame was visible for 500 ms in between the presentation of the two intervals to maintain information for the control of eye posture and to define epipolar lines. The noise frame was visible after the target was presented until the subject responded.

2.2. Results and discussion

The results of Experiment 1 are presented in Fig. 5. These results were averaged over three naive subjects (DS, JE, MS) for both the occlusion and non-occlusion configurations; individual data were similar to those depicted in Fig. 5. It can be seen that the thresholds in the occlusion configuration were equal to or higher than those in the non-occlusion configurations, despite the presence of pseudodisparity. For the 79 and 85 deg bars, the thresholds were nearly identical to the non-occlusion configuration. At the shallower contour orientations of 20 and 63 degrees, the thresholds increased relative to the non-occlusion configuration. This decrease in sensitivity may be due in part to the increase in the distance between the corners of the bars (Westheimer and Levi, 1987). Depth thresholds for oblique parallel bars (our control configuration) are in general agreement with those previously reported. It has been shown that disparity thresholds increase as the orientation of the target bars approaches horizontal (Ebenholtz and Walchli, 1965; Blake *et al.*, 1976; van Ee and Schor, 2000). For orientations greater than 60 deg, depth thresholds are relatively constant, falling to around 10 arcsec. For orientations less than 60 deg, thresholds increase precipitously. However, for the 20 deg orientation, the thresholds we found are lower than those previously reported. This is likely due to the fact that we used short lines within a well-defined frame of reference. Van Ee and Schor (2000) showed that a frame of reference and the visibility of the endpoints of the bar help to establish correspondence, increasing sensitivity to relative depth specified by disparity.

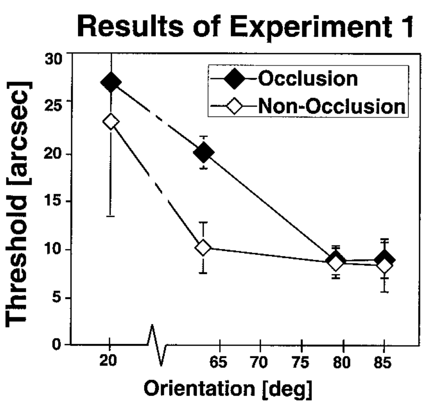

Figure 5. Results of Experiment 1. The threshold represents the smallest disparity step that can be detected as a depth difference. Thresholds in the occlusion configuration are generally larger than in the non-occlusion configuration. Both configurations show an increase of thresholds when the bar orientation gets smaller. The error bars represent one standard deviation in the mean across three subjects. There is no evidence that the occlusion configuration improves the precision of stereoscopically perceived depth.

Thus, the results of Experiment 1 did not provide any support for the hypothesis that pseudodisparity enhances sensitivity to relative depth. One possible explanation for this negative result is that observers interpret the occlusion configuration as a single 'X' shaped object lying in a single depth plane instead of two bars at different depths. There are at least two potential reasons why no benefit in stereoscopic sensitivity could be found from the presence of pseudodisparity. One possibility is that pseudodisparity is not used to resolve depth ordering, but rather, only plays a role in breaking camouflage (by generating 'illusory' contours). Alternatively, it is possible that our attempts to define epipolar lines with random-dot patterns were not effective, causing a basic ambiguity as to the cause of the junction shift (i.e. whether it is a disparity or a pseudodisparity). Experiment 2 was performed in an attempt to provide a less ambiguous stimulus for matching to test this second possibility.

3. EXPERIMENT 2: MONOCULAR SEGMENTATION CUES AND CURVED STIMULI

3.1. Methods

In Experiment 2, the letters 'S' and 'c' were used as stimuli, which were configured with and without occlusion junctions (Fig. 6). The benefit of these shapes is that their familiarity may aid in the segmentation of these contours into discrete objects, and the curvature potentially provides a less ambiguous stimulus for determining the matching direction of the constituent contours.

The methods were identical to those used in Experiment 1, except the two letters were used instead of the two bars. More specifically, a 2AFC staircase procedure was used to determine the 70.7% correct threshold. The relative disparity of the 'c' was varied, and the (pedestal) disparity of the 'S' was constant at 2.5 arcmin. As shown in Fig. 6, there were two non-occlusion configurations. In the first, the 'c' was placed at the bottom-right of the 'S', and in the second, the 'c' was placed in closer proximity to the 'S' (top-left). These configurations were chosen to ensure that thresholds in the occlusion configuration did not critically depend on the positions of the two letters. The occlusion configuration used in this experiment is illustrated in Fig. 6C.

3.2. Results and discussion

The thresholds averaged over three subjects (BA, JE, MS) are shown in Fig. 6D. In the occlusion configuration the threshold was 13.2 ± 2.8 arcsec, and in the two non-occlusion configuration they were 11.8 ± 2.7 ('c' bottom-right) and 11.4 ± 2.5 arcsec ('c' top-left). The thresholds are slightly larger than in experiment 1. The increase in thresholds in both conditions suggests that this stimulus actually provided less effective spatial information in establishing binocular correspondence than that used in Experiment 1 (or at least provided less useful disparity information). Nonetheless, the results of this experiment were qualitatively the same as those

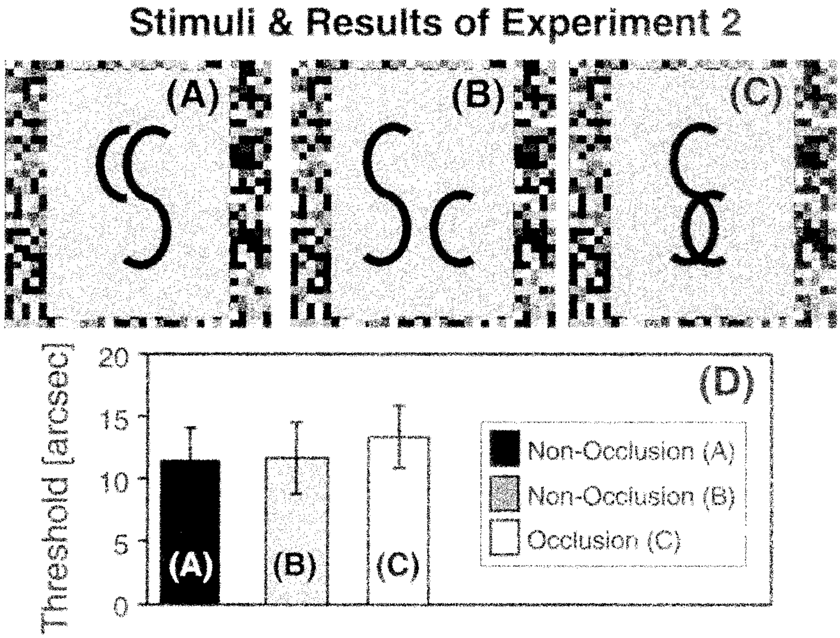

Figure 6. In Experiment 2 we added monocular segmentation cues by using more familiar objects like an 'S' and a 'c' and we used curved objects that consisted of image features that disambiguated the matching process. The depicted half-images are exact replications of the half-images that were presented in Experiment 2. In the non-occlusion configurations the 'c' is presented either to the top-left of the 'S' (A) or to the bottom-right of it (B). (C) represents the occlusion configuration. Panel (D) depicts the results for the three configurations. The thresholds in the two non-occlusion configurations (gray and black histogram bar) are slightly smaller than in the occlusion configuration (white histogram bar). The error bar is the standard deviation in the mean across three subjects.

observed in Experiment 1. In both experiments, we were unable to find any evidence that the pseudodisparity generated by occlusion junctions provides any additional information that the visual system uses to compute relative depth.

Thus, both Experiment 1 and 2 suggest that pseudodisparity of an occlusion junctions is either not used to resolve depth ordering, or that the stimuli used in these experiments were ineffective in defining the global geometry of the two views. Recall that the local junction shifts can be generated by either errors in eye alignment, or by local occlusion geometry, so there must be sufficient additional information in these two images to distinguish between pseudodisparity and disparity. In these first two experiments, we used a random dot surround in an attempt to define a coordinate space against which the junction shifts could be measured, thereby deducing that they were purely local and therefore due to occlusion geometry. However, the two bars that we used were untextured, which means that there was no information in the immediate vicinity of these junctions that unambiguously determined the relative positions of the two eyes. It may be that this lack of information limited performance in these displays. We therefore conducted a third experiment to determine whether a contribution of pseudodisparity to stereoscopic acuity could be observed in patterns containing local information about epipolar geometry.

4. EXPERIMENT 3: TEXTURE CUES AND INCREMENT THRESHOLD

4.1. Methods

In the third experiment we further tested the hypothesis that performance in the occlusion configuration was hampered in experiments 1 and 2 because of a lack of monocular segmentation cues. One possible reason for the negative results of experiments 1 and 2 is that the visual system could not determine which features were matchable and unmatchable. To overcome this problem, we added texture to the bars used in Experiment 1, since this stimulus apparently provided more effective depth information that the patterns used in Experiment 2. More specifically, we added texture (random dots) along either one or both of the bars, which were combined in a manner that caused the bars to appear either opaque or transparent (see Fig. 7). The purpose of the texture elements was to increase the information about both the (epipolar) matching geometry of the images, and to provide additional local information that the visual system could potentially use to identify the unpaired (non-corresponding) regions, and hence, the image displacements corresponding to pseudodisparities.

As an additional test, we also measured both increment and absolute thresholds. We reasoned that we might have failed to find an effect of pseudodisparity because the visual system could not parse the two bars into two separate surfaces. It has been conclusively demonstrated that unpaired contour terminators can generate vivid percepts of occlusion (Anderson, 1994; Anderson and Julesz, 1995; Malik *et al.*, 1999), so we knew that this information could be used in assessing occlusion geometry. We reasoned that even if pseudodisparities were not used to increase absolute stereoscopic sensitivity, pseudodisparities might be used to scale the *magnitude* of depth differences once unmatched contour segments were labeled as such. Indeed, Malik *et al.* (1999) demonstrated that the size of an unmatched contour segment does alter the perceived depth and orientation of the inferred (illusory) occluding surface. We therefore performed two experiments to determine whether pseudodisparities could influence increment detection thresholds, even if they are not used to boost absolute thresholds.

There were four texture conditions used for each of the occlusion and non-occlusion configurations. The target bar (2), whose disparity was varied in the 2AFC task, was either an opaque gray occluder, an opaque textured occluder, or a transparent textured bar (Fig. 7). The other bar (1) was either textured or non-textured. The conditions differ in the amount of information that is available for the identification of unpaired regions. The orientation of the bars was 79 deg with respect to horizontal. We examined increment thresholds at three different relative disparities (8, 16, and 32 times the absolute threshold).

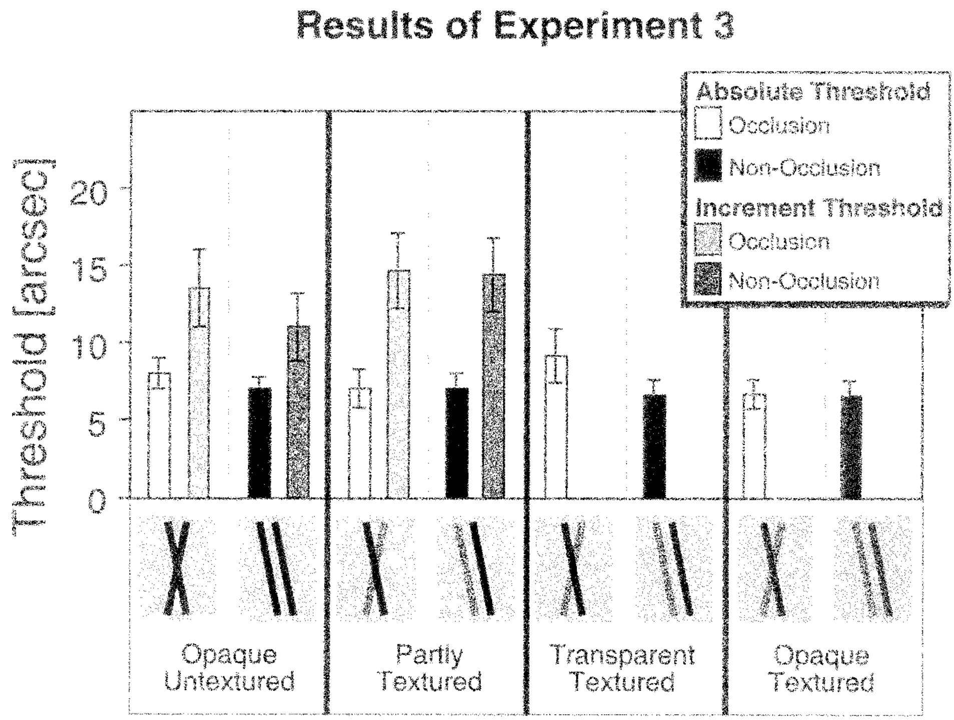

Figure 7. Results of Experiment 3. Icons representing the various stimulus conditions for both the occlusion and the non-occlusion configurations are presented along the horizontal direction. Absolute thresholds (white and black histogram bars) are generally smaller than increment thresholds (light and dark gray histogram bars). The error bars represent one standard error in the mean across the 48 reversals in the staircases of the two participating subjects. There is no significant difference in the obtained thresholds between the occlusion and the non-occlusion configuration. If anything, the obtained thresholds in the occlusion configuration (white and light gray histogram bars) are generally slightly larger than the thresholds in the non-occlusion configuration (black and dark gray histogram bars).

4.2. Results and discussion

Figure 7 shows the averaged stereoacuity across two subjects (JE, MS). As observed in Experiments 1 and 2, we found no evidence that pseudodisparity improves stereoacuity in either the absolute or increment threshold detection tasks. The texture information that was added did not improve stereoacuity. Increment thresholds were similar for each of the three different relative disparities (8, 16, and 32 times the absolute threshold) so we averaged them. Figure 7 shows that the resulting increment thresholds were clearly worse than the absolute thresholds, which can be expected from Weber's law. The most relevant aspect of these data is that the addition of monocular and binocular segmentation cues did not facilitate the detection of depth differences, even when the occlusion relationships were clearly perceptible.

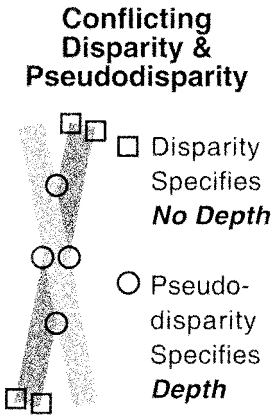

Figure 8. The open squares and disks highlight the locations of features that define disparities and pseudodisparities, respectively, of the occluded bar relative to the non-occluded bar. One is able to create configurations where the disparity signals are below threshold (indicating that there is no depth between the occluded and non-occluded bars), while the pseudodisparities indicate that there is a depth between the bars. Thus, in such configurations there is a conflict between the depth information provided by disparity and the depth information provided by pseudodisparity.

5. GENERAL DISCUSSION

We were unable to find any evidence that pseudodisparity is used to boost stereoscopic sensitivity when the disparities in the scene are below threshold. We did not encounter a condition in which thresholds in the occlusion configuration were smaller than in the non-occlusion configuration. Our current results provide insights into the functional role played by pseudodisparity in stereoscopic processing. It has been previously demonstrated that a vertical pseudodisparity offset can be used by the visual system in constructing an 'illusory' contour, even when only a single contour is present in the two eyes (Anderson 1994; Anderson and Julesz, 1995; Malik *et al.*, 1999).

A possible explanation for our failure to observe a result of pseudodisparity is that it is usually given small weight relative to the depth provided by disparity. In our stimuli, there was a conflict present between the depth provided by the pseudodisparity signal and the depth provided by the disparity signal. Consider a stimulus situation in our displays in which disparity detection is at threshold. At the endpoints of the bars and along the bars' contours, disparity computations are ineffective in providing depth information, simply because the disparity difference is too small to be resolved. However, at the occlusion junctions in the image, the vertical pseudodisparity indicates that a depth difference is present. This is illustrated in Fig. 8. In this figure, the locations of features of the occluded bar that define disparities and pseudodisparities are highlighted by the open squares and disks, respectively. When disparity detection is below threshold, the disparity information in the image specifies that no depth difference is present. Thus, any depth signal generated by pseudodisparity would have to veto this information and cause the bars to be separated in depth. Apparently, even if the vertical pseudodisparity signal was detected it did not override the depth signaled by disparity. This hypothesis is not inconsistent with the findings of Anderson (1994), Anderson and Julesz (1995) and Malik *et al.* (1999). In their stimuli, there was no conflict between the depth provided by the pseudodisparity signal and the depth provided by the disparity signal. Thus, there was simply no disparity signal present that could have specified the depth of the illusory contour, and no contradictory information specifying the absence of a depth difference.[1]

We suggest that the use of pseudodisparity is restricted to those conditions where binocular disparity is unspecified. From this perspective, the primary role of pseudodisparity is to break camouflage, and therefore serves primarily for the interpolation of surface properties at occluding boundaries.

Acknowledgements

RVE was supported by a fellowship of the Royal Netherlands Academy of Arts and Sciences and grant 1 R29 EY11487-01A1 to BA. We are grateful for the help received from Dr. M. Singh.

NOTE

1. Liu, Stevenson and Schor (1994) reported that quantitative depth perception could be elicited by a specially designed stereogram in which there seemed to be no luminance defined corresponding features. However, Gillam (1995) and Liu, Stevenson and Schor (1997) stated that matchable features were present in this stereogram.

REFERENCES

Anderson, B. L. (1994). The role of partial occlusion in stereopsis, *Nature* **367**, 365–368.

Anderson, B. L. and Julesz, B. (1995). A theoretical analysis of illusory contour formation in stereopsis, *Psychol. Rev.* **102**, 705–743.

Anderson, B. L. and Nakayama, K. (1994). Towards a general theory of stereopsis: Binocular matching, occluding contours, and fusion, *Psychol. Rev.* **101**, 414–445.

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

Cormack, L. K., Stevenson, S. B. and Schor, C. M. (1991). Interocular correlation, luminance contrast and cyclopean processing, *Vision Research* **31**, 2195–2207.

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

Gillam, B. (1995). Matching needed for stereopsis, *Nature* **373**, 202–203.

Gulick, W. L. and Lawson, R. B. (1976). *Human Stereopsis: A Psychophysical Approach.* Oxford University Press, New York.

Howard, I. P. and Rogers, B. J. (1995). *Binocular Vision and Stereopsis.* Oxford University Press, New York.

Liu, L., Stevenson, S. B. and Schor, C. M. (1994). Quantitative stereoscopic depth without binocular correspondence, *Nature* **367**, 66–69.

Liu, L., Stevenson, S. B. and Schor, C. M. (1997). Binocular matching of dissimilar features in phantom stereopsis, *Vision Research* **37**, 633–644.

Malik, J., Anderson, B. L. and Charowhas, C. E. (1999). Stereoscopic occlusion junctions, *Nature Neuroscience* **2**, 840–843.

van Ee, R. and Schor, C. M. (2000). Unconstrained stereoscopic matching of lines, *Vision Research* **40**, 151–162.

Westheimer, G. and Levi, D. M. (1987). Depth attraction and repulsion of disparate foveal stimuli, *Vision Research* **27**, 1361–1368.

Wheatstone, C. (1838). On some remarkable and hitherto unobserved phenomena of binocular vision, *Phil. Trans. Roy. Soc. London* **2**, 371–394.