

Human Factors: The Journal of the Human Factors and Ergonomics Society

<http://hfs.sagepub.com/>

Surface Textures Improve the Robustness of Stereoscopic Depth Cues

Geb Thomas, Joseph H. Goldberg, David J. Cannon and Steven L. Hillis

Human Factors: The Journal of the Human Factors and Ergonomics Society 2002 44: 157

DOI: 10.1518/0018720024494766

The online version of this article can be found at:

<http://hfs.sagepub.com/content/44/1/157>

Published by:



<http://www.sagepublications.com>

On behalf of:



Human Factors and Ergonomics Society

Additional services and information for *Human Factors: The Journal of the Human Factors and Ergonomics Society* can be found at:

Email Alerts: <http://hfs.sagepub.com/cgi/alerts>

Subscriptions: <http://hfs.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations: <http://hfs.sagepub.com/content/44/1/157.refs.html>

>> [Version of Record](#) - Jan 1, 2002

Surface Textures Improve the Robustness of Stereoscopic Depth Cues

Geb Thomas, The University of Iowa, Iowa City, Iowa, **Joseph H. Goldberg**, Oracle Corporation, Red Wood Shores, California, **David J. Cannon**, The Pennsylvania State University, University Park, Pennsylvania, and **Steven L. Hillis**, The University of Iowa, Iowa City, Iowa

This research develops design recommendations for surface textures (patterns of color on object surfaces) rendered with stereoscopic displays. In 3 method-of-adjustment procedure experiments, 8 participants matched the disparity of a circular probe and a planar stimulus rendered using a single visible edge. The experiments varied stimulus orientation and surface texture. Participants more accurately matched the depth of vertical stimuli than that of horizontal stimuli, consistent with previous studies and existing theory. Participants matched the depth of surfaces with large pixel-to-pixel luminance variations more accurately than they did surfaces with a small pixel-to-pixel luminance variation. Finally, they matched the depth of surfaces with vertical line patterns more accurately than they did surfaces with horizontal-striped texture patterns. These results suggest that designers can enhance depth perception in stereoscopic displays, and also reduce undesirable sensitivity to orientation, by rendering objects with surface textures using large pixel-to-pixel luminance variations.

INTRODUCTION

A wide variety of computer-mediated tasks require users to accurately perceive and discriminate among the three-dimensional (3D) positions of rendered objects. Depth, the spatial component in the direction away from the viewer, is among the most challenging components of clear and reliable spatial communication. Monocular depth cues, such as linear perspective, interposition, and textural gradients, assist the viewer in determining relative and absolute depth of objects in a graphical user interface. Stereoscopic displays, such as helmet-mounted displays and shutter glasses (e.g., CrystalEyes, StereoGraphics, Inc.) can add binocular disparity and convergence depth cues to reinforce the perception of depth. The convergence depth cue depends on the difference in gaze angle required to bring the fixated object to the center of each retinal image. The

binocular disparity depth cue depends on disparate retinal images created by the lateral separation of the viewer's eyes. Stereoscopic displays have been applied in computer-aided design software, factory simulators, entertainment software, and training simulators (Durlach & Mavor, 1995) but have not yet gained wide acceptance and pervasive use. Even for applications in which a slight enhancement to the operator's geometric understanding of a robotic environment can prevent expensive accidents or save lives (Cannon & Thomas, 1997), 3D rendering technologies are seldom used.

More effective use of available design parameters may improve the effectiveness of such displays. Currently, few guidelines for stereographic displays identify desirable properties of rendered objects or suggest how important these properties may be to user comprehension of object geometry. The developers' documentation provided by StereoGraphics Corporation

(Akka, 1998; Lipton, 1997), for example, describes the geometry of stereoscopic displays, appropriate use of graphics calls, and hardware setup, without mentioning how displayed objects might be adapted to improve presentation. Of Lipton's 13 design recommendations for stereoscopic displays, none mentions the importance of surface detail to achieve robust, effective presentation of relative object depth.

Although stereoscopic displays can enhance depth presentation, many subtle factors (e.g., object size, exposure duration, hue, intensity, texture, and viewing distance) can influence the viewer's ability to reliably interpret the information presented (Patterson, Moe, & Hewitt, 1992). Because it is difficult to predict the effectiveness of each stereoscopic cue, some researchers have proposed that each application should be individually investigated (Hsu, Pizlo, Chelberg, Babbs & Delp, 1996). The approach adopted in this research attempts to identify a general technique to improve the robustness of perceived depth in stereoscopic display applications, thereby improving their effectiveness and utility.

Stereoscopic depth perception is a relative measure of depth; the perceived absolute depth of a feature in a stereoscopically displayed image requires an estimate of both the feature's stereoscopic disparity and the absolute depth of the fixation point. When using a stereoscopic display to deduce the relative depth between a target and the observer's fixation point, the visual system compares the visual angle between the fixation and target points as observed separately by the left and right eyes. The difference in the perceived visual angles is referred to as *stereoscopic disparity*. Generally the disparity entails both horizontal and vertical components. As the viewer's eyes are separated horizontally, the horizontal component of the stereoscopic disparity is typically several orders of magnitude larger than the vertical component. Objects closer than the fixation point have a negative, or crossed, disparity. Objects farther than the fixation point have a positive, or uncrossed, disparity. Cormack and Fox (1985) and Koenderink and van Doorn (1976) reported the mathematical relationship between disparity and the relative depth of an object.

When rendered for stereoscopic display, objects with large, untextured surfaces (surfaces without color or luminance variations) typically provide ineffective stereoscopic depth cues. In these instances, the visual system has few unique, discernible features at which to sample stereoscopic disparity, as illustrated in Figure 1. Figure 1a presents a graphical display of a model of a table set against a wall. The scene is illuminated by an infinitely distant light source. In a stereoscopic display, the wall and floor surfaces are identical in the left and right images of the stereoscopic pair. Because there are no features from which disparity can be sampled, the depth of the wall and floor are stereoscopically ambiguous. Only objects with nonhorizontal edges, such as the left and right sides of the table and the table's legs, exhibit an unambiguous stereoscopic depth. The image still contains some depth cues, such as foreshortening and occlusion, but much of the benefit of stereoscopic disparity is lost. Figure 1b presents the same image rendered using surface textures. In this image, horizontal disparity can be measured on every surface from almost any viewpoint. The depth of all the objects in Figure 1b will be less ambiguous, regardless of the user's viewpoint. The surface textures also reinforce the perspective depth cue by providing texture gradients, further supporting accurate perception of the geometric relationships.

The optical system in a stereoscopically rendered image requires distinguishable locations at which the visual system can measure disparity (Yeh & Silverstein, 1992). Most stereo-matching algorithms assume the presence of sampling positions (Julesz, 1971); however, few studies have analyzed the effect of various texture patterns on depth perception. Harris and Parker (1992) found that depth discrimination performance becomes consistent when more than 100 dots are used in stereogram stimuli. Pérez-Martínez (1995) found that textures composed of line figures, such as circles and triangles, improved observers' discrimination of differences in form, whereas random-dot patterns improved observers' discrimination of size. Todd, Norman, Koenderink, and Kappers (1997) noted that a granite texture improved the accuracy of observers' estimates of the




Figure 1. An image of a table in a room illuminated by an infinitely distant light source (a) without surface textures and (b) with surface textures. Without texture, the depth of most of the surfaces is stereoscopically ambiguous. With surface textures, the stereoscopic depth is unambiguous, regardless of viewpoint.

shape of abstract, smooth surfaces. Buckley, Frisby, and Mayhew (1989) found that clear-cut texture boundaries between objects assist the visual system in fixing the location of depth contours, or depth creases, among objects. The experiments presented here show that surface textures reduce variability in response error related to stimulus orientation. The studied surface textures, composed of random noise and line patterns, provide additional sampling positions to improve disparity-matching performance.

The debilitating effect of stimulus orientation on stereoscopic perception is readily observed with linear stimuli – that is, stimuli dominated by long parallel edges, such as needles, rods, and taut strings. Ogle (1955) showed that stereoacuity, a measure representing the smallest detectable difference in disparity, varies pro-

portionally with the cosine of the angle of inclination, measured with respect to horizontal. The greater the angle of the stimuli from the vertical position, the larger the depth interval needed for depth discrimination. In the present work we comply with the convention followed by Blake, Camisa, and Antoinetti (1976) and measure angles with respect to the horizontal, in which case stereoacuity varies with the sine of orientation.

Figure 2 illustrates the relationship between stimulus orientation and disparity. The top portion of Figure 2 illustrates an experimental setup in which two cameras view a sphere and a narrow, tilted rod. Both cameras are aligned so that the sphere appears in the center of each image. Because the cameras are in different locations, each camera views the rod in a different position relative to the sphere. Combining




Figure 2. Two cameras record slightly different views of a rod (oriented at θ° with respect to horizontal) when focusing on the more distant sphere. Summing the images produces the double image of the rod. The insert shows the length of the rod, L , the vertical projection of the rod ($L \sin \theta$), the horizontal separation of the double rod images (γ), and their perpendicular separation ($\gamma \sin \theta$).

the images from both cameras results in an image with two rods. The horizontal separation between the two rods in the combined image is the disparity of the rod. The disparity depends on the rod's position relative to the sphere as well as on the distance between the cameras and the sphere. The disparity does not depend on the orientation of the rod. For a person viewing the same scene under the same conditions, however, the rod's perceived disparity would depend on its orientation.

Ebenholtz and Walchli (1965) also measured stereoscopic acuity as a function of orientation. In their experiments, participants judged which of two parallel lines was perceived as closer when the lines were presented at different orientations. Ebenholtz and Walchli found that stereoscopic thresholds correlated with the sine of the stimulus orientation, and they have proposed two explanations for this effect, the first of which is based on a geometric observation. The visual system should be able to detect the horizontal disparity of the tilted stimulus, which is independent of stimulus orientation. Ebenholtz and Walchli noted that the perpendicular disparity (the shortest

distance between the images of the line in the left and right views), denoted by γ' in Figure 2, varies with the stimulus orientation, corresponding to the pattern observed in stereoscopic thresholds. Consequently, they suggested that the visual system measures the perpendicular, rather than horizontal, separation between tilted stimuli. A similar observation was made by Blake et al. (1976), who showed that the vertical projection of the stimulus bars, L' , also varies with the stimulus orientation consistent with the trend observed with stereoscopic thresholds. Blake et al. suggested that the vertical projection is related to the accuracy of the disparity perception process.

Ebenholtz and Walchli (1965) proposed a second, quite different, hypothesis suggesting that the stereo-processing system may rely on directionally sensitive sensors, as in the visual cortex of cats and monkeys (Hubel & Wiesel, 1968). Disparity signals for receptors tuned to lines with small orientation angles may provide stimulation for stereoscopic depth perception that is less effective than that of identical disparities produced by nearly vertical stimuli. This hypothesis suggests why different stimuli may elicit different stereoscopic sensitivities, given that different regions of the visual cortex respond differently. This hypothesis does not, however, suggest why stereoscopic sensitivity should vary with object orientation or why the variation would follow a sine pattern.

Other features of textures are also likely to influence depth matching in a stereoscopically presented image. In most current rendering libraries, programmers must select an image to serve as the surface texture. Current human factors and programming literature provides little guidance regarding desirable image features. Pixel-to-pixel intensity variation in the image will ultimately determine whether or not unique stereoscopic disparity sampling positions are salient in the rendered image. Also, lighting conditions in the physical viewing environment will affect image contrast, and the geometric and lighting modeling assumptions will affect the texture visibility. Striped textures present a different challenge to the viewer of a stereoscopic display because these textures may be susceptible to the same perceptual sensitivity to orientation presented by a single line.

In this case, the orientation of a striped surface texture should affect an observer's depth-matching accuracy. In addition, striped patterns sometimes induce the "wallpaper effect" (Jordan, Geisler, & Bovik, 1990), in which the visual system, rather than determining disparity by matching corresponding points on the same stripe in each retinal image, miscalculates disparity by matching similar points from adjacent stripes. This systematic mistake can cause a parallel-line pattern to appear either closer or farther away than it actually is by an offset related to the angular separation of parallel stripes.

Regardless of the reasons underlying the sinusoidal relationship, unless a stereoscopic application constrains stimulus orientation, the display's practical benefit will be diminished by inaccurate perceptions of object depth whenever the object is viewed with an orientation that obscures disparity cues. Constraining the orientations of viewed geometric objects, however, places a severe restriction on versatile design and display systems. The results presented herein suggest that adding appropriate surface textures to rendered objects eliminates this difficulty and substantially improves the effectiveness of stereoscopic displays.

OBJECTIVES AND HYPOTHESES

The present study investigated observers' stereoscopic depth-matching ability under varying conditions of stimulus depth, orientation, and texture mapping. Unlike prior research, these experiments used flat surface stimuli, rather than linear stimuli, to impart useful design recommendations in computer-rendered 3D objects. The experimental task required participants to employ stereoscopic depth cues to place a circular probe (or cursor) at the same perceived depth as the stimulus object.

Three experiments were conducted to model disparity-matching ability. The objective of Experiment 1 was to measure depth disparity error when adjusting the probe to the depth of a single edge of the planar stimulus, with the edge presented at different orientations. The results of this experiment will provide designers with a quantifiable basis for depth disparity accuracy when selecting, moving, or placing

objects with a single visible edge in a stereoscopic display. The objective of Experiment 2 was to investigate the degree, if any, of depth-matching improvement gained by adding surface texture to the horizontal comparison of planar objects. Subtle textures with small pixel-to-pixel contrasts should be less effective than similar textures with large pixel-to-pixel contrasts, given that subtle textures provide fewer easily discernable features at which to sample disparity. Experiment 3 further investigated the sources of depth-matching error by introducing striped stimuli, intended to increase stereoscopic depth perception error via the wallpaper effect. The objective of this experiment was to quantify the expected depth-matching error attributable to the introduction of this texture feature. In addition, this experiment allowed a broader interpretation of the sensitivity of depth estimation to differences in texture patterns.

GENERAL METHODS

Apparatus and Stimuli

In each experiment, shutter goggles and a stereoscopic display using a field-sequential rendering technique (CrystalEyes, StereoGraphics Corporation) were used to present participants with a circular probe at one depth and a stimulus at another. In each trial, the participant increased or decreased the disparity of the circular probe until it appeared to be at the same depth as the stimulus. The probe maintained a consistent size in each view regardless of its disparity. A fixture positioned the stereoscopic goggles so each participant's head was 75 cm from the center of the monitor. The fixture also prevented interference from ambient light and stray reflections on the monitor surface. All stereoscopic disparities were generated assuming that (a) the participant's eyes were 75 cm from the center of the screen, (b) the participant had an interocular separation of 6.5 cm (Boff and Lincoln, 1988), and (c) the participant fixated on the center of the display. Throughout the experiments, disparity was calculated as $\text{disparity} = 2 \cdot \arctan(J/D_f) - 2 \cdot \arctan(J/D_t)$, in which J is interocular separation, D_f is the distance to the fixation point, and D_t is the distance to the target (Cormack & Fox, 1985).

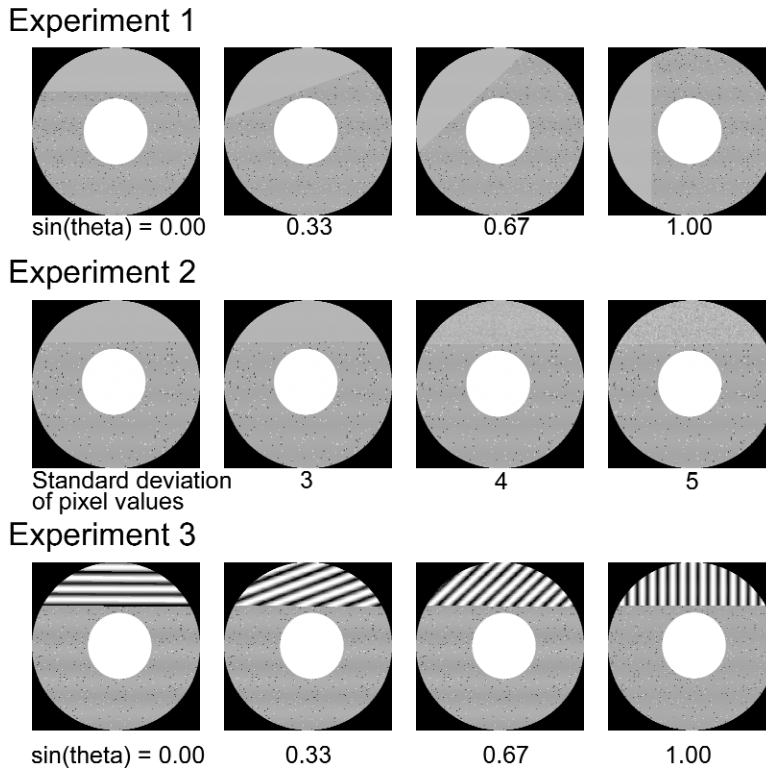


Figure 3. The stimuli used in these experiments arranged by rows for Experiments 1, 2, and 3. Note that the stimuli were presented without green, resulting in shades of purple.

Stimuli, as shown in Figure 3, were presented within a 15.2° circular aperture at the center of the display. The stimulus was presented with a disparity of -19.0 or 20.0 arcmin from the aperture center, equivalent to a depth of 5 cm in front of or 5 cm behind the monitor, respectively. The central circular probe had a diameter of 6.0° . The probe was presented with a randomly selected disparity representing a depth within 1 cm of the stimulus. The stimulus and probe appeared in front of a lightly textured purple background with a stereoscopic disparity of 35.0 arcmin at its center, equivalent to a depth of 10 cm behind the monitor (Figure 4). On a scale from 0 to 255, 0.5% of the background pixels had red-green-blue (RGB) values of (255, 0, 255) and another 0.5% had RGB values of (0, 0, 0). The remaining pixels had RGB values of 125, 0, 125. Although the probe could be directed to positions farther from the viewer than the background (resulting in large errors in stimulus depth estimation), this caused a depth cue conflict because

the stereoscopically distant probe occluded the stereoscopically nearer background. This situation usually occurred only in trials in which the stereoscopic depth of the stimulus was ambiguous.

The colors and intensities of the probe, stimulus, and background were selected to reduce the effect of ghosting. Ghosting occurs when portions of the image intended for one eye are visible in the image intended for the other eye. This imaging artifact produces blurred edges that cause eyestrain (Yeh & Silverstein, 1990) and permit participants to strategically use ghost width for matching, rather than disparity.

A Silicon Graphics Iris Indigo workstation employing the Inventor software library was used to render and animate the stereoscopic test images. A 34.3×27.4 cm HL7965Kw-SG Mitsubishi monitor displayed the probe, stimulus, and background with luminances of 1.00, 0.49, and 0.36 cd/m^2 , respectively. Luminance was measured without shutter glasses.

Participants

Nine male students at the Pennsylvania State University, ranging in age from 22 to 36 years, participated in either two or three of these experiments. Each participant was screened for visual criteria, similar to those of Yeh and Silverstein (1990). A Bausch and Lomb Orthorator was used to measure the near and far visual acuity (Snellen 20/20 or better), near and far vertical phoria (<0.5 prism diopters of deviation), near lateral phoria (-6.0 to 0.0 prism diopters), far lateral phoria (-1.66 to 2.33 prism diopters), stereoscopic acuity (27 arc seconds or less), and color vision (3 of 4 correct in the Orthorator test) of each participant.

Procedure

Each participant was instructed to proceed steadily through the experiment and to concentrate on making consistently accurate responses. Each completed a training session of eight trials, including an example from each condition, prior to each experiment. During each experiment, participants were permitted to take a break at any time, but they were encouraged to steadily continue through the 15-min experiment.

Participants employed the up and down arrow cursor keys on a keyboard to adjust the probe's disparity in increments corresponding to a forward or backward motion of 0.07 cm. The right and left arrow cursor keys provided coarse disparity control corresponding to move-

ments of 0.2 cm. After positioning the probe, participants pressed the space bar to advance to the next trial, which immediately presented the stimulus and probe at new positions. The trials were not timed.

Experimental Designs

Four repetitions of fully crossed (within-subject) factorial designs were performed in all three experiments. In each experiment, each of the 8 participants completed 32 trials: four randomized sets of the four levels of the main independent variable crossed with two stimulus disparities.

Statistical Procedures

Statistical analysis of the root mean square of the disparity error (RMSE) was conducted for each participant under each condition. Huynh and Feldt (1976) adjusted F statistics (designated in the following as F_{adj}) were used to test for the main effects of the two within-subject factors (depth and the main independent variable) and their interaction. As there were only 8 participants, adjusted F tests were used rather than Hotelling's T^2 multiple analysis of variance (MANOVA) test, because the adjusted F test is more appropriate for small sample sizes (Kirk, 1982, p. 261.) A significant interaction test was followed by post hoc comparison and regression analysis, whereas a nonsignificant interaction test was followed by main effects tests and then by appropriate post hoc MANOVA tests. It was not necessary to assume equal variances for the different experimental conditions or equal correlations between the different experimental conditions for any of the tests.

EXPERIMENT 1: STIMULUS ORIENTATION

Method

Disparity matching was conducted for an untextured stimulus at four orientations and two depths. The four orientations are illustrated in the first row of Figure 3, with sine values of .00, .33, .67 and 1.00, corresponding to angles of 0°, 19.27°, 42.1°, and 90°, respectively. Horizontal stimuli (sine = 0) demonstrated an extremely ambiguous depth, given that the




Figure 4. The stereoscopic position of the aperture, stimulus, and background for far stimuli conditions. The user moved the probe forward and backward until it appeared to be at the same depth as the stimulus.

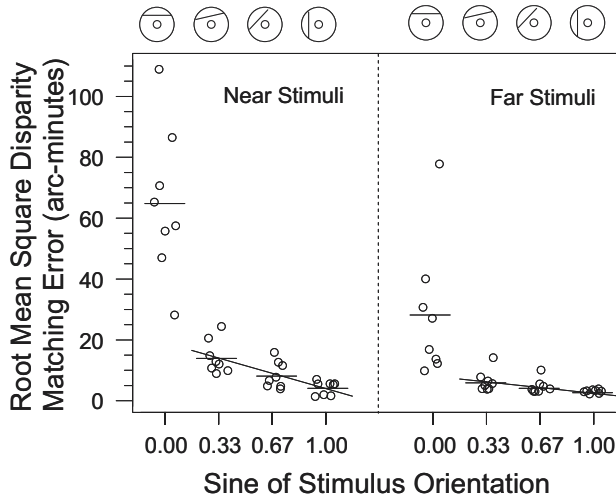


Figure 5. RMS disparity-matching error for an untextured, planar stimulus edge for each research participant at four stimulus orientations and two depths. Fitted lines for the three nonhorizontal orientations at each depth are superimposed. The data have been randomly perturbed in the x direction to improve legibility.

stereoscopic images in the pair were nearly identical, differing only in foreshortening created by perspective transformation. This difference led to a 1-pixel vertical difference in the lower corners of the stimulus. The three nonhorizontal conditions presented unambiguous stereoscopic depth cues and verified the sinusoidal hypothesis. The nearly ambiguous horizontal condition measured worst-case, effectively random, performance.

Results and Discussion

Figure 5 presents the RMSE values for each participant at each depth and orientation. A significant interaction was evident between depth and orientation, $\text{adj-}F(3, 21) = 18.34$, $p < .01$. The RMSE of the ambiguous horizontal stimulus was significantly larger than the other three stimuli at each depth, $F(1, 7) > 8.13$, $p < .03$ (each comparison), indicating that participants had difficulty estimating the ambiguous disparity of horizontal stimuli. The RMSE for near stimuli was larger than for far stimuli for all except the vertical orientation, $F(1, 7) > 9.27$, $p < .02$ (each comparison), indicating that the sensitivity of the disparity-matching task varies with depth.

Although the ambiguous horizontal condition created a significant lack-of-fit for the linear model with respect to RMSE versus stim-

ulus orientation at each depth, $F(2, 6) = 24.27$ and 5.19 , $p > .01$ and $> .05$ (near and far, respectively), no significant nonlinearity was detected, $F(1, 7) < 0.79$, $p > .20$ (near and far, respectively) for the three unambiguous conditions. The fitted least-squares lines from the three nonambiguous orientations were $\text{RMSE} = 19.03 - 15.04 \sin(\theta)$ for the near stimuli and $\text{RMSE} = 7.96 - 4.87 \sin(\theta)$ for the far stimuli. The R^2 values were $.52$ and $.25$, respectively, when fit to the individual RMSE values and $.9923$ and $.9996$, respectively, when fit to the means from each orientation. The fitted least-squares estimates were identical regardless of whether individual or mean data were employed for fit. The sinusoidal sensitivity to stimulus orientation, originally observed from the stereoscopic threshold of linear stimuli such as wires and needles, thus also occurs for the RMS disparity-matching error of the edges of untextured, extended stimuli. The ambiguous horizontal condition did not conform to the trend indicated by the other orientations, suggesting that a transition region from the sinusoidal to the ambiguous region lies somewhere between 0° and 19.27° of stimulus orientation.

Stimulus orientation is an important predictor of disparity-matching accuracy and must be considered in stereoscopic rendering of untextured surfaces. The next experiment investigated

whether surface textures can mitigate the relative insensitivity to horizontal edges.

EXPERIMENT 2: RANDOM TEXTURES

Method

Four texture patterns were applied to the surface of a horizontal stimulus by defining pixel RGB values of $X, 0, X$, with X distributed normally about a mean of 128 (Figure 3, middle row). Gaussian pixel variations of this type have been frequently used to describe noise in image-processing applications (Gonzalez & Wintz, 1987). The pixels on the textured surface were presented with a spatial frequency of 15.6 pixels/° and a luminance of $-0.3 + 0.0066 X$ cd/m². Standard deviations of 3, 4, and 5 were used when generating pixel values for the smooth, intermediate, and coarse textures, respectively. A standard deviation of 1 corresponded to an average pixel-to-pixel luminance contrast of approximately 1.2%, using the equation luminance contrast = $(L_{max} - L_{min}) / L_{max}$. In the fourth (control) condition, all pixels used RGB values of 128, 0, 128, equivalent to a standard deviation of zero.

Results and Discussion

Mean RMSE values, computed for participant, depth, and texture, are plotted as a

function of orientation in Figure 6. The mean RMSE of each of the three textured conditions was significantly less than that of the control condition, $F(1, 7) > 11.70, p < .02$ (for each texture and depth), indicating that participants could make more precise depth matches using textured, as opposed to untextured, horizontal edges in a stereoscopic display.

A significant interaction was evident between depth and texture, $adj-F(3, 21) = 3.48, p < .04$. The mean RMSE values for the three far, textured stimuli conditions did not significantly differ, $adj-F(2, 14) = 1.83, p > .21$, for which the overall mean RMSE was 2.22. Thus all three textures were approximately equally effective in reducing the ambiguity of the horizontal stimulus depth. In contrast, the mean RMSE values for the three near, textured stimuli significantly differed, $adj-F(2, 14) = 4.35, p < .05$, and demonstrated a significantly decreasing linear trend, $F(1, 7) = 6.06, p < .05$. Although no significant quadratic component existed, $F(1, 7) = 1.61, p > .24$, Figure 6 suggests a tapering off beyond the intermediate texture. A pair-wise comparison between the intermediate and coarse textures in the near condition was, however, not significant, $F(1, 7) = 0.52, p > .49$. Overall, for the near stimuli the intermediate and coarse textures provided slightly more perceptual benefit than that provided by the smooth texture.

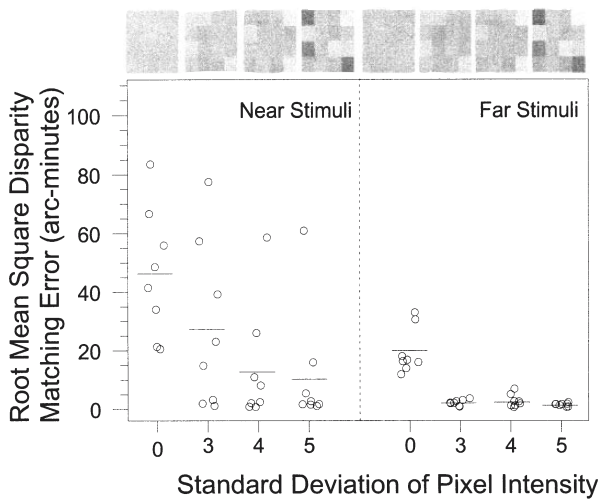


Figure 6. RMS disparity-matching error for a horizontal planar stimulus at two depths with four random textures differing in the standard deviations of the Gaussian distribution of their pixel intensities. The icons at the top of the figure have been exaggerated somewhat to indicate the variation among the textures. The data have been randomly perturbed in the x direction to improve legibility.

TABLE 1: Mean RMSE of Textured (Experiment 2) Stimuli Minus Vertical, Untextured Stimuli (Experiment 1)

Depth	Texture	RMSE Difference	SD	<i>t</i>	<i>p</i>
Near	Smooth	23.05	28.16	2.31	.054
Near	Intermediate	9.50	19.36	1.39	.208
Near	Coarse	7.16	19.99	1.01	.344
Far	Smooth	-0.85	1.19	-2.03	.082
Far	Intermediate	-0.25	2.14	-0.33	.751
Far	Coarse	-1.54	0.93	-4.70	.002

Table 1 presents the differences between the textured mean RMSE values from Experiment 2 and the vertical orientation mean RMSE values from Experiment 1. For near stimuli, variation was not significant, except for a difference with regard to the marginally significant smooth texture ($p = .054$). For the far stimuli, textured mean errors were all less than vertical stimulus errors, attributable only to the fact that coarse texture RMSEs were significantly less than the vertical orientation RMSE ($p = .002$). Thus, the best texture performance compares favorably with vertical stimulus performance.

This experiment demonstrated that even a subtle texture can significantly improve depth perception in a stereoscopic display. The benefit of texture roughness is significant for average pixel-to-pixel luminance contrasts as small as 3.6% for far stimuli and 5.2% for near stimuli. The third experiment explored linear symmetry, another feature of surface textures that influences depth perception in a stereoscopic display.

EXPERIMENT 3: ORIENTED TEXTURES

The sample texture used in Experiment 3, illustrated in the third row of Figure 3, consisted of a 1.07 cycle/° sinusoidal grating pattern using the same four orientations as in Experiment 1. This pattern was investigated to see if participants demonstrated a similar RMSE pattern to that observed for the rotated stimulus edge.

Results and Discussion

Eleven depth-matching trials using the vertical orientation resulted in errors that were

much larger than the corresponding values for the two intermediate orientations. These outliers were tightly clustered around disparity values offset by 56.45 arcmin, which corresponds to the visual angle between similar features in the 1.07 cycle/° sinusoidal grating. This error pattern is typical of the wallpaper illusion, in which the left-eye and right-eye images are 360° out of phase with each other. The remainder of this analysis employs data corrected for this offset.

Figure 7 presents the average RMSE values at each depth. Error matching did not differ between near and far depths ($p > .10$) or among the Depth × Orientation levels ($p > .10$). Matching error from the horizontal line pattern was greater than that in each of the other line patterns at each depth, $F(1, 7) > 26.13$, $p < .01$, for each comparison. However, there were no significant differences among nonhorizontal line patterns for either the near or far conditions, $\text{adj-}F(2, 14) = 0.66$ and 0.22 , $p > .53$ and $> .66$, near and far, respectively. Unlike Experiment 1, no evidence was observed supporting a Depth × Orientation interaction or any relationship between matching error and texture orientation. The nonhorizontal line pattern linear slope for error versus orientation was not significant for either near or far conditions ($p > .50$). The depth matching was equally accurate for both stimulus depths, which was not the case for untextured or randomly textured stimuli.

Pair-wise comparisons of the six RMSE means in this experiment with the corresponding means from Experiment 1 show only one significant difference. The mean RMSE for the second orientation near stimuli in Experiment

1 was significantly greater than the second orientation in Experiment 3, $F(1, 7) = 6.18, p < .05$, indicating that parallel-line textures either remove the sinusoidal effect of orientation or render the sinusoidal pattern too subtle to detect in this experiment. This condition may be a result of the redundant stimulus information presented by the repeating pattern or by additional vertical disparity differences created by perspective transformation. However, this analysis was performed after the data were corrected for the wallpaper effect. The uncorrected data yielded a multimodal distribution resulting from participants' occasional, exceptionally inaccurate, estimates of vertical stimulus disparity. Although the results of the corrected parallel-line pattern were somewhat better than those of the rotated stimulus, it would be imprudent to recommend the use of parallel-line textures, given the hazards of the wallpaper illusion. An irregular-striped texture, however, should achieve the benefits of rich texture without risking the effects of wallpaper illusion.

OVERALL DISCUSSION

Three experiments investigated observers' ability to match computer-generated objects in projected depth as a function of the direction of depth disparity, object orientation, and

object texture. Experiment 1 demonstrated that depth-matching accuracy is sensitive to both stimulus orientation and the direction of the depth disparity (crossed vs. uncrossed). Experiments 2 and 3 demonstrated that the presence of textures on the stimulus surface can mitigate this sensitivity and improve the effectiveness of stereoscopic applications.

The RMS of disparity error varied linearly with stimulus orientation, at least over the range from 90° (vertical) to 19.5° (relatively horizontal). In the horizontal condition, the images in the stereoscopic pair were almost identical except for subtle, 1-pixel differences at the corners of the stimulus. Many participants ignored these subtle horizontal stimuli cues and placed the probe at unexpected positions, such as the background depth, the far clipping planes, or the screen depth.

The observed influence of orientation on depth matching was largely predicted by existing theories; however, the magnitude of the errors for unambiguous conditions (3–15 arcmin) was much larger than expected from the 1 arcmin error performance described in the existing literature.

The influence of depth on matching error was evident in both Experiments 1 and 2, in that matching accuracy was greater for far stimuli (with uncrossed disparities) than for

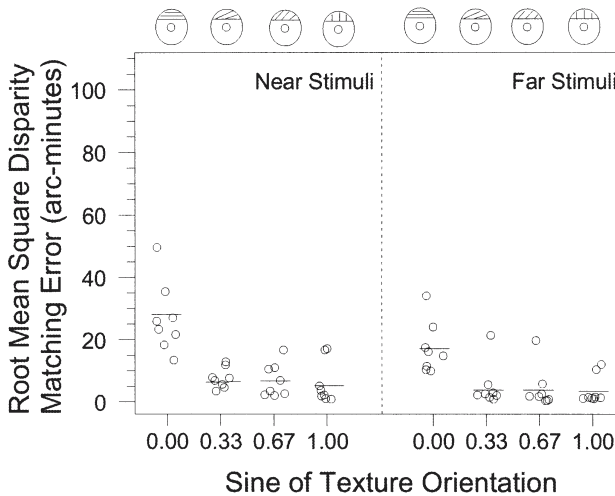


Figure 7. RMS disparity-matching error for a horizontal, planar stimulus at two depths with a surface texture of parallel lines positioned at four orientations. The data have been randomly perturbed in the x direction to improve legibility.

near stimuli (with crossed disparities). This effect was not significant in Experiment 3. Other researchers (Fox, 1985; Patterson, Cayko, Short, Flanagan, Moe, Taylor, & Day 1995; Patterson & Fox, 1984) have found that for small targets or short presentation periods, depth estimates for uncrossed disparities are less accurate than depth estimates of crossed disparities. In Patterson et al. (1992) this effect disappeared for a stimulus with a diameter of 5.5° . Parrish and Williams (1990) found that the apparent depth of larger stimuli was perceived less accurately for uncrossed disparities than for crossed disparities. The stimuli in this experiment were larger than 5.5° , so the main effect of depth is consistent with the emerging pattern found in other research, although more investigation into the effects of stimulus size on depth-matching accuracy is required to conclusively demonstrate this relationship.

A second explanation for the depth effect is that a cue conflict in which the ends of the stimulus are superimposed by the stereoscopically more distant aperture may reduce the precision of the depth estimates. In this experiment, the aperture was placed on the monitor surface. In the near condition, the stimuli were presented in front of the screen; however, the stereoscopic depth of the aperture appeared to the observer to be behind the depth of the stimulus. This conflict may have caused observers to place the depth of the objects closer to the monitor surface than the stereoscopic depth cues alone would indicate. Although the experiment could have been designed to avoid this conflict, the results would have been less relevant to application designers. In normal viewing conditions, the monitor edges naturally create a cue conflict when portions of an object fall outside the monitor's rendering area.

Although the Depth \times Orientation interaction observed in Experiment 1 does not resolve the debate with respect to the cause of the orientation effect, this significant effect casts a new light on the three theories proposed in the literature. Ebenholtz and Walchli's (1965) first hypothesis suggests that the sinusoidal orientation effect is caused by the lateral separation of the oriented stimulus rather than the horizontal disparity. The lateral separation may be controlled independent of orientation by ad-

justing the depth of the stimulus, suggesting a technique to support or refute this hypothesis by comparison of disparity-matching error for an array of depths and orientations. In the research results reported herein, the absolute value disparities of the near and far conditions differed by 5%, resulting in a variation of the lateral separation of the tilted stimulus edge also by 5%. Although the pattern of results in Experiment 1 conforms to the pattern predicted by Ebenholtz and Walchli's first hypothesis, the resolution of the experiment cannot reliably discern such a subtle difference.

Ebenholtz and Walchli's (1965) second hypothesis of different receptors is consistent with the Depth \times Orientation interaction because different receptors are used for crossed and uncrossed orientations. The hypothesis is not predictive, however, in the sense that unless the function of all receptors is specified, the behavior of the visual system cannot be predicted. The Blake et al. (1976) vertical projection hypothesis is not consistent with the observations, given that the vertical projection of the stimuli was consistent at both depths, whereas the matching error was significantly different. The hypothesis might be extended, however, to account for this difference by postulating different stereo mechanisms for use at different depths.

The lack of significance of the Depth \times Orientation interaction for crossed and uncrossed disparities observed in Experiment 3 is also interesting. This result suggests that when more information is available, the effect of orientation is more difficult to detect. No current hypothesis directly accounts for this situation; however, the result is consistent with the general recommendation of this research: Providing more information to the stereoscopic system improves the effectiveness of the stereoscopic depth cue and improves performance in stereoscopic applications. The ability of "good" surface textures to eliminate the effects of orientation in stereoscopic displays was clearly evident from Experiments 2 and 3.

Experiment 2 demonstrated that surface textures can largely eliminate the orientation effects by creating enough sampling positions that the greatest threat to accurate stereoscopic perception – the ambiguous horizontal case –

is perceived with RMS disparity error values similar to those for vertical targets.

CONCLUSIONS

Experiment 1 demonstrated that the effects of planar stimulus orientation on RMS disparity-matching error conform to the sinusoidal pattern of stereoscopic thresholds reported by Ogle (1955). The experiment also demonstrated that stimulus orientation is an important consideration in analyzing the effectiveness of modern spatial displays, indicating that ensuring the accuracy of relative depth judgments requires stereoscopic application designers to consider the orientation of large, untextured objects with straight edges. Because the magnitude of the effect also depends on the depth of the stimulus, designers should also control the viewing distance. User-controlled viewpoints and interactive models can make such orientation and depth constraints difficult or impossible to impose. However, without controlling the effects of planar orientation, stereoscopic displays can be unreliable and unpredictable tools for transferring geometric information.

Experiment 2 demonstrated how the effects of orientation might be overcome by incorporating even subtle textures with luminance contrasts greater than 3.6% on the surfaces of large objects. Experiment 3 defined some of the limitations of using striped textures to improve stereoscopic perception. The experiment indicated that parallel-line textures can improve viewer perception of the depth of a horizontal stimulus as long as the texture pattern is not horizontal and the pattern does not create a wallpaper illusion. Consequently, designers are advised to apply nonsymmetric textures to objects presented in stereoscopic displays.

The practical benefit of the observations described herein is that designers of stereoscopic display applications can improve the accuracy of user depth perception by incorporating textures on the surfaces of displayed objects. Textures should be selected according to the following criteria:

1. Textures should provide unique disparity-sampling positions across the surface of the displayed objects.
2. Useful textures may require only slight color variations. The experimental results indicate that pixel-to-pixel intensity variation with a luminance contrast of 3.6% and frequencies of 15.6 pixels/° provides satisfactory performance. Figure and background luminance contrasts, however, most likely play an important role in determining the lower limit for appropriate texture variations.
3. Patterns of identical parallel lines should be avoided because effective depth presentation is sensitive to viewing orientation and the wallpaper effect.

Application of these recommendations should effectively eliminate the impact of viewing position and orientation on the presentation of depth in a stereoscopic application.

ACKNOWLEDGMENTS

This research was partially supported by a grant from Sandia National Laboratories.

REFERENCES

- Akka, B. (1998). *Converting existing applications*. San Rafael, CA: StereoGraphics Corp.
- Blake, R., Camisa, J., & Antoinetti, O. (1976). Binocular depth discrimination depends on orientation. *Perception and Psychophysics*, 20, 113-118.
- Boff, K. R., & Lincoln, J. E. (1988). *Engineering Data Compendium: Human perception and Performance*. Wright-Patterson Air Force Base, OH: 1, 48-49.
- Buckley, D., Frisby, J., & Mayhew, J. E. (1989). Integration of stereo and texture cues in the formation of discontinuities during three-dimensional surface interpolation. *Perception*, 18, 563-588.
- Cannon, D., & Thomas, G. (1997). Virtual tools for supervisory and collaborative control of robots. *Presence Journal*, 6(1), 1-28.
- Cormack, R., & Fox, R. (1985). The computation of retinal disparity. *Perception and Psychophysics*, 37, 176-178.
- Durlach, N. I., & Mavor, A. S. (Eds.). (1995). *Virtual reality - Scientific and technological challenges*. Washington, DC: National Academy Press.
- Ebenholtz, S. M., & Walchli, R. M. (1965). Stereoscopic thresholds as a function of head- and object-orientation. *Vision Research*, 5, 455-461.
- Fox, R. (1985). Interaction of image characteristics of stereoscopic forms during depth perception. (Tech. Report N14-001-85C-001) Arlington, VA: Office of Naval Research, Engineering Psychology Programs.
- Gonzalez, R. C., & Wintz, P. (1987). *Digital image processing*. Reading, MA: Addison-Wesley.
- Harris, J. M., & Parker, A. J. (1992). Efficiency of stereopsis in random-dot stereograms. *Journal of the Optical Society of America A*, 9, 14-24.
- Hsu, J., Pizlo, Z., Chelberg, D. M., Babbs, C. F., & Delp, E. J. (1996). Issues in the design of studies to test the effectiveness of stereo imaging. *IEEE Transactions on Systems, Man and Cybernetics - Part A: Systems and Humans*, 26, 810-819.

- Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology (London)*, *195*, 215–245.
- Huynh, H., & Feldt, L. S. (1976). Conditions under which mean square ratios in repeated measurement designs have exact F-distributions. *Journal of Educational Statistics*, *1*, 69–82.
- Jordan, J., Geisler, W., & Bovik, A. (1990). Color as a source of information in the stereo correspondence process. *Vision Research*, *30*, 1955–1970.
- Julesz, B. (1971). *Foundations of cyclopean perception*. Chicago: University of Chicago Press.
- Kirk, R. E. (1982). *Experimental design* (2nd ed.). Pacific Grove, CA: Brooks/Cole.
- Koenderink, J. J., & van Doorn, A. J. (1976). Geometry of binocular vision and a model for stereopsis. *Biological Cybernetics*, *21*, 29–35.
- Lipton, L. (1997). *Developers' handbook*. San Rafael, CA: StereoGraphics Corp.
- Ogle, K. N. (1955). Stereopsis and vertical disparity. *Archives of Ophthalmology*, *53*, 495–504.
- Parrish, R.V., & Williams, S.P. (1990). *Determination of depth-viewing volumes for stereo three-dimensional graphic display* (Tech. Report AVSCOM 90-B-016) Hampton, VA: Langley Research Center.
- Patterson, R., Cayko, R., Short, G.L., Flanagan, R., Moe, L., Taylor, E., & Day, P. (1995). Temporal integration differences between crossed and uncrossed stereoscopic mechanisms. *Perception & Psychophysics*, *57*(6), 891–897.
- Patterson, R., & Fox, R. (1984). Metacontrast making between cyclopean and luminance stimuli. *Vision Research*, *30*, 439–448.
- Patterson, R., Moe, L., & Hewitt, T. (1992). Factors that affect depth perception in stereoscopic displays. *Human Factors*, *34*, 655–667.
- Pérez-Martínez, D. (1995). Texture discrimination at the cyclopean retina. *Perception*, *24*, 771–786.
- Todd, J. T., Norman, J. F., Koenderink, J. J., & Kappers, A. M. (1997). Effects of texture, illumination, and surface reflectance on stereoscopic shape perception. *Perception*, *26*, 807–822.
- Yeh, Y., & Silverstein, L. D. (1990). Limits of fusion and depth judgment in stereoscopic displays. *Human Factors*, *32*, 45–60.
- Yeh, Y., & Silverstein, L. D. (1992). Spatial judgements with monoscopic and stereoscopic presentation of perspective displays. *Human Factors*, *34*, 583–600.

Geb Thomas is an assistant professor of industrial engineering at the University of Iowa and director of the Graphical Representation of Knowledge Lab. He received a Ph.D. in industrial engineering in 1996 from the Pennsylvania State University.

Joseph H. Goldberg is a Principal Research Scientist of Advanced User Interfaces at Oracle Corporation. He received his Ph.D. in industrial and operations engineering in 1985 from the University of Michigan.

David J. Cannon, is an associate professor of industrial and manufacturing engineering at the Pennsylvania State University. He received his Ph.D. in Mechanical Engineering in 1992 from Stanford University.

Stephen L. Hillis is director of the Statistical Consulting Center in the Department of Statistics and Actuarial Science at the University of Iowa. He received his Ph.D. in statistics in 1987 from the University of Iowa.

Date received: April 7, 1997

Date accepted: August 1, 2001