Interpolation across surface discontinuities in structure from motion

ASAD SAIDPOUR, MYRON L. BRAUNSTEIN, and DONALD D. HOFFMAN
University of California, Irvine, California

Interpolation across orientation discontinuities in simulated three-dimensional (3-D) surfaces was studied in three experiments with the use of structure-from-motion (SFM) displays. The displays depicted dots on two slanted planes with a region devoid of dots (a gap) between them. If extended through the gap at constant slope, the planes would meet at a dihedral edge. Subjects were required to place an SFM probe dot, located within the gap, on the perceived surface. Probe dot placements indicated that subjects perceived a smooth surface connecting the planes rather than a surface with a discontinuity. Probe dot placements varied with slope of the planes, density of the dots, and gap size, but not with orientation (horizontal or vertical) of the dihedral edge or of the axis of rotation. Smoothing was consistent with models of 2-D interpolation proposed by Ullman (1976) and Kellman and Shipley (1991) and with a model of 3-D interpolation proposed by Grimson (1981).

Wallach and O'Connell (1953) demonstrated that if subjects are shown the two-dimensional (2-D) shadow cast by a clear glass sphere that has small opaque dots on its surface and rotates about a vertical axis, then subjects perceive a spherical surface. A similar perception is commonly reported for computer-generated displays in which dots move about on a computer screen in a manner consistent with their being projections of dots on a rotating sphere (see, e.g., Braunstein, 1966; Braunstein & Andersen, 1984). Apparently human vision is adept at inferring 3-D structure from the 2-D motions of projected features. This process of inferring 3-D structure from 2-D motions, referred to as *structure from motion* (SFM) after Ullman (1979), has been studied extensively since the advent of computer-generated displays. (For reviews, see Braunstein, 1978, 1983, 1988.)

In many displays of SFM in which dots alone are projected on the computer screen, observers report seeing more than just a 3-D structure for the dots. They report, in addition, that they perceive a continuous surface in 3-D passing through, or near, the dots. Such perceptions of "subjective surfaces" on the basis of sparse collections of dots are not limited to SFM displays. Subjective surfaces have also been reported, for instance, with stereograms. (See Julesz, 1971; Julesz & Frisby, 1975; Stevens & Brookes, 1987.) Uttal, Davis, Welke, and Kakarala

(1988) found that subjects could discriminate among eight quadratic and cubic surfaces in stereograms containing as few as four dots on a surface. Grimson (1981) presented random dot stereograms of a half cylinder. In these stereograms a vertical strip, parallel to the major axis of the half cylinder and in the center of the display, was devoid of feature points. Grimson found that subjects interpolated a surface even through the region devoid of feature points. Surface interpolation has also been reported with line drawings depicting 3-D surfaces (Stevens & Brookes, 1987). And in the case of 2-D curves having gaps, visual interpolation leads to the well-known phenomenon of illusory contours (see, e.g., Brady & Grimson, 1981; Halpern & Salzman, 1983; Kanizsa, 1979; Kellman & Shipley, 1991; Ullman, 1976).

In a previous study of surface interpolation in SFM, Saidpour, Braunstein, and Hoffman (1992) used stimuli similar to Grimson's half cylinders. The displays were orthographic projections of dots on the surface of half cylinders rotating back and forth ±19° about a vertical axis through the center of the cylinder. A single dot, called the *SFM probe dot*, was placed in a vertical strip devoid of feature points along the major axis of the cylinder (the *gap*). The probe dot moved in phase with all other points on the cylinder, but its amplitude (projected velocity) could be adjusted by the subject. The amplitude determined the simulated distance of the probe dot from the axis of rotation, and the probe dot appeared to move in depth when the amplitude was adjusted. The subject's task was to adjust the depth of the probe dot until it appeared to be on the surface of the object. For a range of dot densities and gap widths, subjects placed the probe dot well in front of a cylindrical interpolation, suggesting that a slightly bulging surface was interpolated. It was concluded that no current theory of surface interpolation could account for this result.

Our previous investigation, like Grimson's, was limited to surfaces with continuous, indeed constant, curvature. A critical issue in theoretical accounts of surface interpolation is the handling of discontinuities. Some theories of surface interpolation minimize a functional that interpolates smooth surfaces (e.g., Brady & Horn, 1981; Grimson, 1981). When presented with a set of discrete 3-D feature points taken from a surface that has discontinuous edges or sharp corners, the functional interpolates a surface that is smooth even though the feature points are consistent with the presence of a sharp corner. Therefore, surface discontinuities pose a special problem for these theories of interpolation: it is difficult to determine when it is not appropriate to interpolate a smooth surface and when it is appropriate to allow a discontinuity (especially when there is noise in the data—see, e.g., Blake, 1984; Blake & Zisserman, 1986, 1987; Horn, 1983; Terzopoulos, 1982, 1986; Weiss, 1990). In what follows, we will focus on the interpolation theories of Grimson (1981), Kellman and Shipley (1991), and Ullman (1976), comparing our experimental results with the predictions of their theories. Therefore we will briefly review their theories.

Ullman (1976) developed a model for interpolating 2-D subjective contours between two edges. The contours are composed of the arcs of two circles tangent to the two edges, meeting smoothly, and minimizing the total curvature. Ullman's model provides a good fit to subjects' performance in interpolating subjective contours. Kellman and Shipley (1991) give a similar construction. But instead of using two arcs of circles, they use one arc and one line, with the arc tangent to one edge and the line tangent to the other, and require the arc and the line to intersect smoothly. In one special case, Kellman and Shipley's and Ullman's models both interpolate an arc of a circle—namely, when the two edges are so oriented that their tangents, when extended, intersect at a point equidistant from the two edges. Of course both models are proposed for interpolating gaps in plane curves, not gaps in surfaces, and the processes for interpolation in 3-D may differ from those used in 2-D. Nevertheless, the models can be used for interpolating through gaps in a special class of surfaces called ruled surfaces. These are surfaces such as planes and cylinders, which are formed by rigidly translating a plane curve along a straight line. If the plane curve has a gap, the resulting ruled surface will also have a gap. To interpolate this gap in the surface one can first use Kellman and Shipley's or Ullman's model to interpolate the gap in the plane curve. One can then translate the interpolated plane curve along a straight line to generate the interpolated surface. As we will describe shortly, the stimuli in our experiments depict ruled surfaces with gaps, so that the interpolations predicted by these two models can be compared with our experimental results.

Grimson (1981) proposed a model of 3-D surface interpolation based on the "quadratic variation" functional. If a surface S is described as a set of points

$$S = \{(x,y,z) \mid z = f(x,y)\}$$

for some fixed function $f: \Re^2 \to \Re$, then the quadratic variation functional θ is given by

$$\theta(f) = \left\{ \int \int (f_{xx}^2 + 2f_{xy}^2 + f_{yy}^2) dx dy \right\}^{1/2},$$

where the subscripts denote partial derivatives. Grimson proposed that if depth information is available only at discrete points in the visual field, human vision interpolates through those points a surface f, which minimizes the functional θ. Intuitively, such a minimizing surface is the smoothest surface which interpolates the points. Grimson's approach has several properties that bear on its plausibility as a model of human visual interpolation. First, the shape of the interpolated surface is not viewpoint invariant: if the points whose depth is known are rotated rigidly in space, the surface interpolated through the points by Grimson's functional will change its shape (and not just its orientation in space). Second, the discrete implementation Grimson gives for his approach requires that interpolation be done on a grid of regularly spaced points. Finally, Grimson's approach interpolates surfaces that do not have orientation discontinuities. As the density of the interpolation grid increases, Grimson's functional can more closely approximate an orientation discontinuity, but it will not, by itself, interpolate a surface f that has an orientation (or depth) discontinuity. Although all three of these properties deserve careful experimental investigation, it is to this last property of Grimson's functional that our experiments are most directly related.

Figure 1 shows schematically the stimuli used in our experiments. The stimuli each simulated two planes separated by a gap; these planes are indicated in the figure by heavy solid lines. (All surfaces in this figure are depicted as being perpendicular to the plane of the figure.)
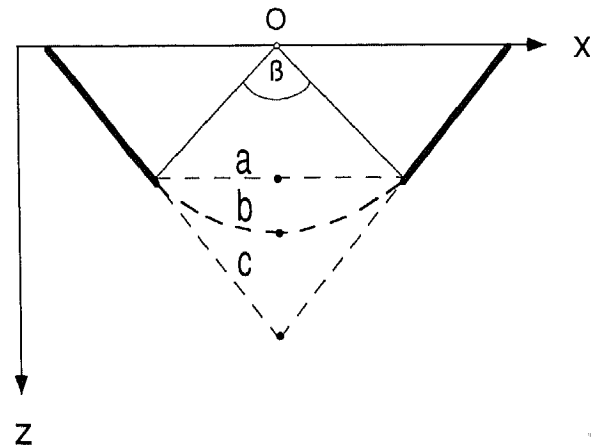

Figure 1. Three surfaces that might be interpolated through a gap in a dihedral angle (see text). The angle β, measured in degrees, specifies the gap size. The axis of rotation is vertical and passes through the origin, O.

In the stimuli for our experiments, the visible feature points are pseudorandomly placed only on the two slanted planes; there are no feature points in the gap. The figure also illustrates three (of many possible) surfaces that subjects might interpolate through the gap. These are indicated by dashed lines. Surface *a* is a plane connecting the edges of the gap. Surface *b* is the surface interpolated by Kellman and Shipley's (1991) and Ullman's (1976) algorithms. Surface *c* extends the two planes until they meet at the center of the gap, resulting in an orientation discontinuity.

The purpose of the present study was to provide quantitative data on surface interpolation in displays of SFM with stimuli like those shown schematically in Figure 1. Do subjects interpolate surfaces with a discontinuity, as in surface *c*, or do they interpolate a smooth surface? In a series of three experiments, subjects were shown displays of structure from motion depicting stimuli as schematized in Figure 1. As mentioned earlier, feature points were pseudorandomly placed on the two planes, but no feature points were placed in the gap between the planes. A single dot, the SFM probe dot, was placed in the center of this gap, and subjects manipulated the apparent depth of this dot until it appeared to lie on the surface that they perceived to be passing through the gap. In the first experiment, the planes and gap were oriented vertically, in such a way that the two planes, if extended, would intersect in a vertical line. (Henceforth we use the phrase *dihedral angle* to refer to the surface obtained by extending these two planes until they meet in a line of intersection, as illustrated by surface *c* of Figure 1. We use the phrase *dihedral edge* to refer just to the line where the planes would intersect.) The planes and the gap rotated together rigidly about a vertical axis (see Figure 1). We examined how variations in size of the gap, density of the dots, and slant of the planes affected the placement of the probe dot. In the second experiment we studied the effects of the orientation (horizontal or vertical) of the dihedral edge and of the axis of rotation; the purpose was to determine whether the results of the first experiment could be generalized to other combinations of edge and axis orientation. The third experiment was a partial replication of the first, except that the depth of the dihedral edge was matched across different slants. This was to ensure that the effects of plane slant on probe dot placement in our previous experiments were really due to variations in slant and not simply to variations in the speed of the probe dot required to place it on the dihedral edge.

EXPERIMENT 1

Method

Subjects. The subjects were one of the authors (A.S.) and 2 graduate students from the University of California, Irvine. Except for A.S., the subjects had no knowledge of the purposes of the experiment and were paid for their participation. Normal or corrected-to-normal acuity of at least 20/40 (Snellen eye chart) was required in the eye used throughout the experiment.

Stimuli. The stimulus displays consisted of white circular dots (0.5-mm radius) on a dark background. The dots in a display were

orthographic projections of points on the surfaces of two slanted planes; the two planes, if extended, would intersect along a vertical edge. The two planes almost always appeared to be pointed toward the observer. (Subjects apparently prefer to see convex objects, at least in these SFM displays, which had a vertical axis of rotation.) The surface of each dihedral angle was divided into a number of equal-area square cells with one-dot randomly assigned coordinates within the boundaries of each cell. The number of cells was determined by the density level. The *texture density*—the density of dots—was uniform in the projection as well as on the surface. (An orthographic projection of a uniform random texture on a slanted plane has uniform density in the image: A uniform random 2-D texture that is compressed in one dimension remains uniform in both dimensions.)

As illustrated in Figure 2, a gap was created in a vertical strip centered on the dihedral edge. This was done by removing dots from the display within that vertical strip. The probability that a dot would be included in the display (i.e., not removed) decreased linearly, starting from 1.0 at the edge of this strip and falling to zero within the central one half of the gap (see Figure 2). (Note that this probability could vary within an individual cell.) A gradual decrease in density was used to prevent a sharp density change at the edge of the gap; a sharp change might lead to the perception that the gap was produced by an occluding object (Grimson, 1981).

The gap sizes reported in this paper are based on the line at which the probability of dot placement was .5. Specifically, gap size is defined as the angle between the following two vectors. Both vectors lie in a plane perpendicular to the dihedral edge. Both vectors originate from the same point on the axis of rotation. One vector terminates at a point on the right plane of the dihedral angle where
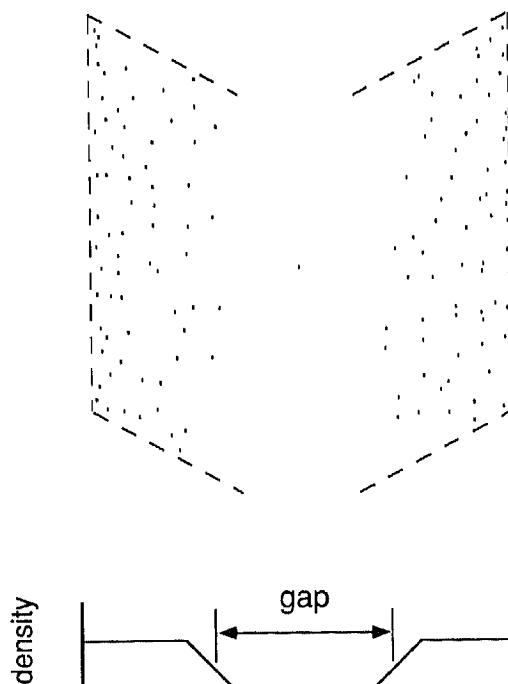

Figure 2. A random dot display simulating a dihedral angle with a gap at its center and a probe dot in the gap (top), and the density profile for three regions of the dihedral angle (bottom). (The actual displays were white dots on a black background.)
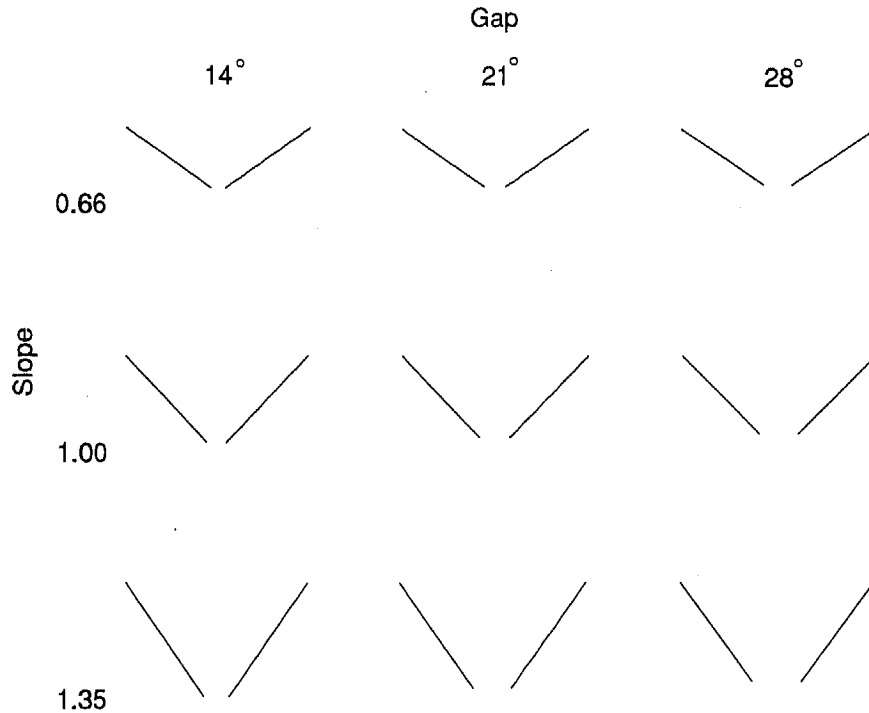
Figure 3. Top view of the simulated dihedral angle with the nine combinations of slope and gap studied in Experiment 1.

the probability of dot placement is .5. The other vector terminates at the corresponding point on the left plane. The combinations of gap size and plane slant used in our experiments are illustrated in Figure 3.

The visible portion of the dihedral angle was clipped to a window 1,200 pixels square. The simulated size of the dihedral angle was sufficiently large that its outer edges never appeared inside the clipping window during oscillation. The simulated depth of these outer edges was 21 pixels in front of the axis of rotation. The simulated depth of the dihedral edge was 396, 600, or 810 pixels, depending on the slope of the planes (see Design).

A single dot, the probe dot, was placed at the horizontal center of the gap. Its vertical position was assigned randomly on each trial within the middle 40% of the gap (see Figure 2), and its initial depth was assigned randomly to be between 50 and 200 pixels in front of or behind the dihedral edge. Subjects could adjust the probe dot to be as near as 98 pixels and as far as 1,638 pixels from the axis of rotation; these limits prevented the dot from being placed too close to the axis, and thus remaining relatively stationary, or too far from the axis, and thus disappearing outside the clipping window. The probe dot oscillated rigidly and in phase with the dihedral angle.

The dihedral angle oscillated continuously about a vertical axis through an angle of ±19°, starting randomly either from +19° or −19°. (The orientation in which the dihedral edge was at the center of the screen was defined as 0°.) The dihedral angle accelerated sinusoidally from the ±19° positions to the 0° position in 0.6 sec (reaching a maximum rotation speed of 31.7°/sec) and then decelerated sinusoidally from the 0° position to the ±19° positions (reaching a minimum speed of 0°/sec). It then reversed direction and repeated this pattern. The sinusoidal function was used to prevent sudden changes of direction at the ends of the oscillation cycle. The dihedral angle continued to oscillate until the subject responded. The frame rate was 40 frames/sec.

Design. We examined three independent variables: slope of the dihedral angle, texture density, and gap size. The slopes were 0.66, 1.00, and 1.35, resulting, respectively, in dihedral angle magnitudes of 113°, 90°, and 73° and in dihedral edge locations at 396, 600, and 810 pixels from the axis of rotation. The texture density levels were 0.21 dots/cm², 0.55 dots/cm², and 1.03 dots/cm² for the smaller slope; 0.26 dots/cm², 0.46 dots/cm², and 1.04 dots/cm² for the middle slope; and 0.22 dots/cm², 0.49 dots/cm², and 1.03 dots/cm² for the larger slope. (The density variations resulted from partitioning dihedral angles with different visible surfaces into equal-area square cells and placing one dot in each cell.) The gap size, measured from the axis of rotation (see Figure 2), was 14°, 21°, or 28°. Each subject responded 10 times to each of the 27 conditions (3 slopes × 3 densities × 3 gaps).

Apparatus. The displays were presented on a 21-in. CRT monitor (Xytron, Model AB21) with a P4 phosphor under computer (Vax Station II) control. The total resolution of the system was 4,096 × 4,096 pixels. The subjects viewed the display monocularly in a darkened room from a distance of 87 cm through a tube-and-mask arrangement that limited the field of view to a square area within the borders of the stimuli. A 0.5 neutral-density filter was inserted in the tube so that any traces on the CRT would not be visible to the observer. The visual angles subtended by the dihedral angle (both height and width) were 6.75°. The visual angle subtended by the projected gap varied with the slope of the dihedral angle and the angle of rotation. For gap sizes of 14°, 21°, and 28°, the visual angles at 0° rotation were 0.51°, 0.74°, and 0.95° for the 0.66 slope; 0.74°, 1.05°, and 1.35° for the 1.00 slope; and 0.96°, 1.35°, and 1.70° for the 1.35 slope (see Figure 3).

The subject's response device consisted of a button and a three-position (double-throw momentary) switch. Moving the switch forward (away from the observer) decreased the amplitude of the dot's oscillation in the image, simulating movement of the dot closer to the axis of rotation (away from the observer if the angle was seen

as convex). Moving the switch backward increased the oscillation amplitude, simulating motion of the dot away from the axis of rotation. When the button was pressed, the computer recorded the amplitude of the probe dot, coded as a distance from the dihedral edge along a line connecting the probe dot to the axis of rotation.

Procedure. Each subject participated in six sessions. The first session was practice. The other five sessions were experimental and began with 2 practice trials followed by a random sequence of 54 trials (27 conditions × 2 repetitions) in two blocks separated by a 5-min rest period. The practice trials were randomly selected from the 27 experimental conditions (3 slopes × 3 gaps × 3 density levels). The intertrial interval was 3 sec.

The subjects were instructed to "move the middle dot until it is on the surface of the object" by using the three-position switch, and to press the button when they were satisfied with their adjustment. They were given an opportunity to practice with the switch and to observe the resulting changes in motion of the probe dot. The room was darkened 2 min before trials began. All subjects were run without feedback in both the practice and the experimental trials.

Results

Constant error. The probe dot was placed inside the dihedral angle on 98% of the trials, and between the dihedral edge and a plane connecting the edges of the gap (i.e., between surfaces a and c in Figure 1) on 88% of the trials. An analysis of variance (ANOVA) was conducted for the three slopes, three gaps, and three density levels. The dependent variable was constant error (CE)—the signed distance between the probe dot and the dihedral edge—averaged over the 10 replications in each condition. The main effect of slope was significant $[F(2,4) = 75.73, p < .01, \omega^2 = .641]$. As the slope increased, subjects placed the probe dot farther inside the dihedral angle (see Figure 4). The mean CEs for the three slope levels, 0.66, 1.00, and 1.35, were respectively -28, -61, and -103 pixels. A post hoc analysis (Tukey's honestly
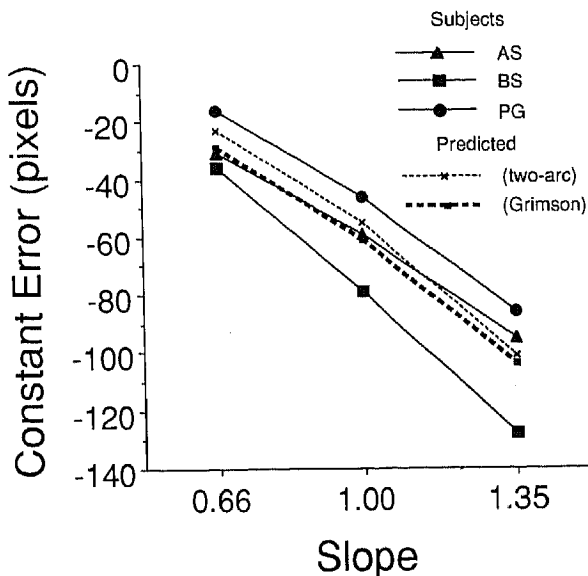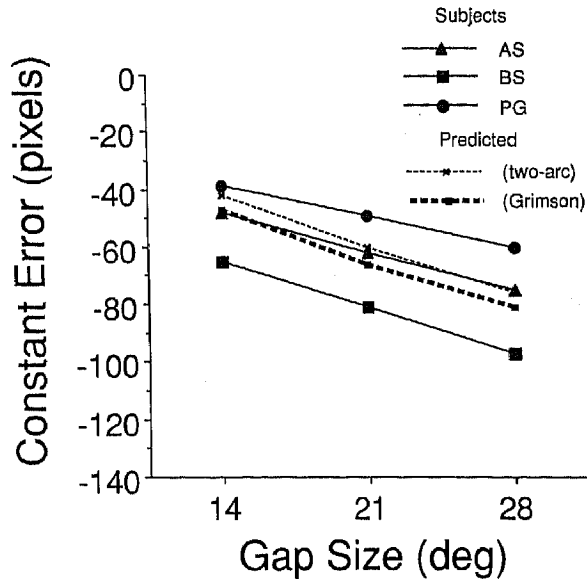

Figure 5. Constant error as a function of gap size for 3 subjects in Experiment 1. See text for explanation of gap size. (Dashed lines show predicted values; see text.)

significant difference [HSD] test) showed the CEs at the three slope levels to be significantly different from one another. There was a significant main effect of gap $[F(2,4) = 59.05, p < .01, \omega^2 = .08]$. The mean CEs for the gap levels of 14°, 21°, and 28°, respectively, were -51, -64, and -77 pixels (see Figure 5). A post hoc analysis showed all levels to be significantly different from one another. There was a significant interaction between slope and gap $[F(4,8) = 20.23, p < .01, \omega^2 = .028$; see Figure 6]. There was also a significant main effect of density $[F(2,4) = 7.27, p < .05, \omega^2 = .046]$. As density increased, subjects' probe dot placements moved closer toward the dihedral edge (see Figure 7). The mean CEs for the three density levels, respectively, were -75, -64, and -53 pixels. The mean CE for the lowest density level was significantly different from the mean for the highest density level. There were no other significant interactions. All subjects reported seeing a 3-D object all of the time and a surface consisting of two planes connected by a smoothly curved region most of the time.

As noted earlier, for a given 3-D gap size the projected gap width varies with the slope of the planes forming the dihedral angle. To examine possible effects of the projected gap width, we selected our levels of gap and slope to include different combinations of gap size and slope that produced the same projected gap widths. Specifically, projected gap widths of approximately 130, 170, and 240 pixels were each produced with two different combinations of 3-D gap and slope. Figure 8 shows CE as a function of projected gap width. Note that the curves for slope are separated at each of the three matched projected gap widths, indicating that the effects of slope on CE results cannot be attributed to variations in the projected gap.

Figure 4. Constant error as a function of slope for 3 subjects in Experiment 1. (Dashed lines show predicted values; see text.)

We compared the CE results to predictions derived from 2-D interpolation models (Kellman & Shipley, 1991; Ullman, 1976) and from a 3-D interpolation model (Grimson, 1981). The first two models have no free parameters; Grimson's model has one free parameter—grid density. As we described earlier, to apply the 2-D models we considered a horizontal cross-section of the dihedral angle display, which is simply an angle with a gap at its apex. For the symmetrical stimuli studied, the same prediction is made by both models. Figure 6 shows the predictions derived from the 2-D models, together with the observed CEs, for the nine combinations of slope and gap size. These predictions are also shown, together with individual subject results, in Figure 4 for the three levels of slope and in Figure 5 for the three levels of gap size.

Grimson's model determines the depth coordinates of feature points at fixed positions on a 2-D grid. To apply Grimson's model, a regular grid must be placed on the image. We selected grid densities that placed grid elements at the edges of the gap and in the center of the gap. This allowed Grimson's model to predict subjects' placements in depth of the probe point, which was in the center of the gap, as a function of slope and gap width. The requirement that a grid element be located at each edge of the gap and in the center of the gap constrains the maximum separation of grid elements to be half of the gap width. Because of this constraint, the grid densities for almost all stimulus conditions were higher than the random dot densities used in the experiment. In these cases, we used the same grid density in Grimson's model to predict gap and slope effects for all three density levels

Figure 7. Constant error as a function of density level for 3 subjects in Experiment 1.
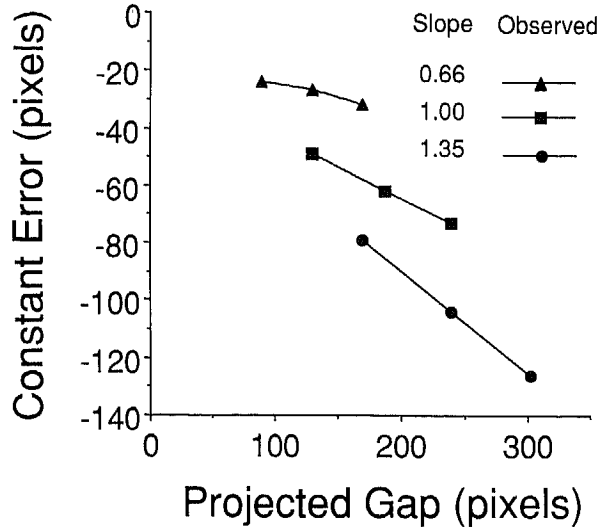
Figure 8. Constant error as a function of projected gap width in Experiment 1.

in our experiment. The relationship between grid density and gap and slope in applying Grimson's model thus precluded use of this algorithm to estimate density effects for our stimuli. Table 1 shows the grid density that we used in applying Grimson's model for each of the 27 combinations of gap, slope, and density. For three stimulus conditions, marked with asterisks in Table 1, the random dot densities were higher than the grid densities determined by the method described above and a closer match could be obtained by the double-density method: We doubled the grid density so that, in addition to grid elements
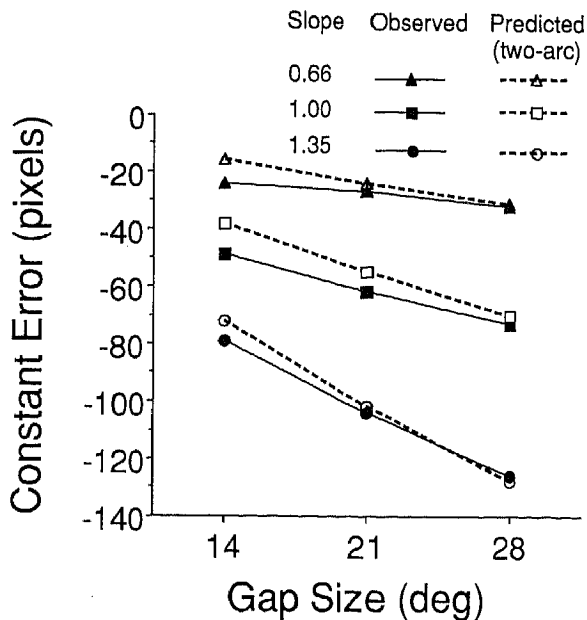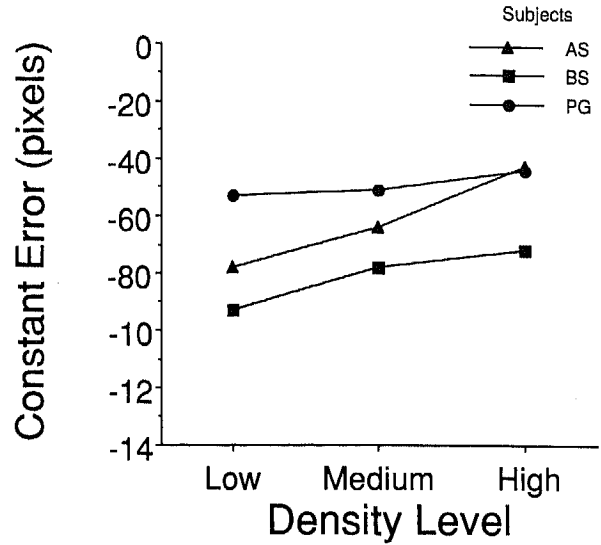

Figure 6. Constant error as a function of slope and gap size in Experiment 1. Dashed lines show values predicted using Ullman's (1976) or Kellman and Shipley's (1991) model.

Table 1
Grid Densities and Constant Errors (CE) in Experiment 1 Compared With Grimson's Model

		Slope										
		0.66				1.00				1.35		
		Density		CE		Density		CE		Density		CE	
Gap	Density	Stimulus	Model	Observed	Model	Stimulus	Model	Observed	Model	Stimulus	Model	Observed	Model
14°	low	5	25	−34	−20	6	17	−57	−44	6	13	−90	−77
	medium	8	25	−21	−20	8	17	−52	−44	9	13	−77	−77
	high	11	25	−18	−20	12	17	−37	−44	13	13	−70	−77
21°	low	5	17	−33	−29	6	11	−67	−63	6	9	−115	−108
	medium	8	17	−25	−29	8	11	−65	−63	9	9	−105	−108
	high	11	17	−22	−29	12	11	−53	−63	13	19*	−90	−101
28°	low	5	13	−40	−37	6	9	−88	−80	6	7	−146	−136
	medium	8	13	−32	−37	8	9	−73	−80	9	15*	−128	−126
	high	11	13	−25	−37	12	19*	−59	−71	13	15*	−104	−126

*Double-density method (see text).

at each edge of the gap and in the center of the gap, there were elements halfway between the edges and the center.

Table 1 also lists the observed CEs for the 27 conditions. We can see that Grimson's model tends to underestimate the magnitude of the CE for the lowest density condition and overestimates it for the highest density condition. To illustrate the fit of Grimson's model to the observed CEs, in a manner similar to our illustration of the fit of Kellman and Shipley's and Ullman's models (which does not make different predictions for different densities), we averaged the predicted CEs in Table 1 across densities. Figure 9 shows that the average fit of the Grimson model over density levels is similar to the fit of Kellman and Shipley's and Ullman's models.

Variable error. In an ANOVA we examined the effects of the three independent variables on the standard deviations (*SDs*) of the subjects' 10 repeated responses

for each of the 27 conditions. The main effect of slope was significant [$F(2,4) = 10.20, p < .05, \omega^2 = .278$]. As slope increased, subjects' *SD* also increased. The mean *SDs* for the three slope levels, 0.66, 1.00, and 1.35, were respectively 10.9, 16.8, and 21.6 pixels. A post hoc analysis, Tukey's HSD, showed the slope levels 0.66 and 1.35 to be significantly different from each other. There was a significant main effect of density [$F(2,4) = 16.01, p < .05, \omega^2 = .195$]. As density increased, subjects' *SD* decreased. The mean *SDs* for the three density levels were 21.4, 15.4, and 13 pixels, respectively. The mean *SD* for the lowest density level was significantly different from the mean of the middle and highest density levels. The main effect of gap was not significant [$F(2,4) = 2.211, p > .05$]. Also, none of the interactions were significant.

Discussion

The principal finding of Experiment 1 was that subjects appear to interpolate a smooth surface, not a surface with discontinuities. Subjects' placements of the probe dot and their verbal descriptions of the appearance of the surface were both consistent with this interpretation. Furthermore, the responses of all 3 subjects closely followed the predictions of both 2-D interpolation models and Grimson's 3-D interpolation model. The effects of slope and gap size were consistent with the predictions of both types of models.

Before one can reach a general conclusion about smooth versus discontinuous interpolation, it is necessary to demonstrate that these results are not specific to a vertically oriented dihedral edge or to a vertical axis of rotation. This issue is addressed in our next experiment, in which two different dihedral angle orientations and two different axes of rotation were studied.

EXPERIMENT 2

In several recent studies of curvature discrimination in displays of SFM (Cornilleau-Pérès & Droulez, 1989; Norman & Lappin, 1992), it has been found that discrimination between curved and flat surfaces is best when the axis of rotation and the direction of greatest surface curvature
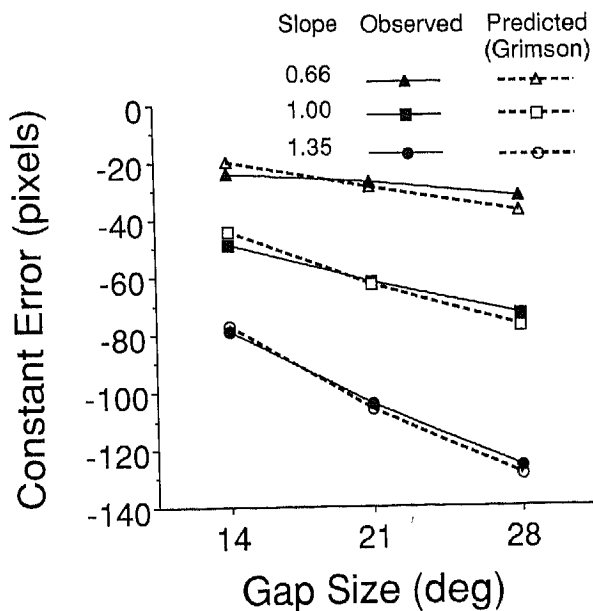

Figure 9. Constant error as a function of gap size in Experiment 1. Dashed lines show values predicted using Grimson's (1981) model.

are parallel. In Experiment 1, if one smooths over the dihedral edge, the direction of greatest curvature is horizontal; thus the axis of rotation and the direction of greatest surface curvature are perpendicular. The studies just cited suggest that, for discriminating between flat and curved surfaces, this arrangement is not optimal. One might argue, therefore, that the tendency to interpolate a smoothly curving surface is due, in part, to this suboptimal arrangement; this tendency might be reduced under conditions more favorable to discriminating between curved and flat surfaces. To examine this possibility, as well as to control for any effects specific to a horizontal slant or a vertical axis of rotation, we included in Experiment 2 all four combinations of vertical and horizontal dihedral edges with vertical and horizontal axes of rotation.

Method

The method was the same as in Experiment 1, except as noted below.

Subjects. The subjects were one of the authors (A.S.) and two graduate students who had no knowledge of the purposes of the experiment and were paid for their participation. One of these subjects (B.S.) had participated in Experiment 1; the other (J.B.) had not.

Design. We examined three independent variables: slope and orientation of the dihedral angle, and orientation of the axis of rotation. The slopes were 0.66, 1.00, or 1.35, as in Experiment 1. The orientation of the dihedral edge was either vertical, as in Experiment 1, or horizontal. The axis of rotation was either vertical, as in Experiment 1, or horizontal. The intermediate gap size, 21°, and the highest density level used in Experiment 1 were used for all conditions in the present experiment. Each subject responded 10 times to each of the 12 conditions.

Procedure. Each subject participated in one practice session and two experimental sessions. Each experimental session began with 2 practice trials followed by a random sequence of 60 trials (12 conditions × 5 repetitions) in 4 blocks. There were 1-min rest periods after each block. The subjects were instructed to respond only when the object was facing toward them. This instruction was used because some depth reversals were reported when the stimuli were used in pilot studies. The rest of the instructions were the same as in Experiment 1.

Results

The subjects placed the probe dot inside the dihedral angle on 100% of the trials and between the dihedral edge and a plane connecting the edges of the gap on 95% of the trials. An ANOVA of the CEs for the three slopes, two axes of rotation, and two orientations of the dihedral angle again yielded a significant main effect of slope $[F(2,4) = 191.06, p < .01, \omega^2 = .93]$. As the slope increased, subjects placed the probe dot further inside the dihedral angle, with mean CEs of -31, -67, and -110 for the three slopes (see Figure 10). Post hoc comparisons (Tukey's HSD test) found the difference between all three slopes to be significant. The main effects of orientation of the axis of rotation $[F(1,2) = 4.74]$ and orientation of the dihedral angle $[F(1,2) = 3.39]$ were not significant. There were no significant interactions. Figure 10 shows predicted CEs based on Kellman and Shipley's and
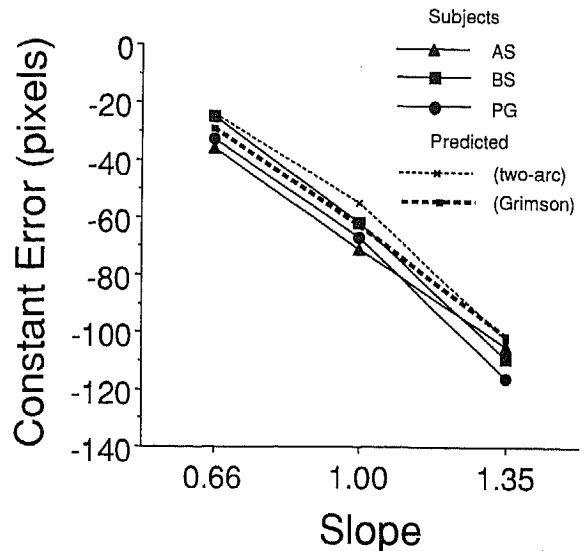
Figure 10. Constant error as a function of slope for 3 subjects in Experiment 2. Dashed lines show predicted values.

Ullman's models and on Grimson's model. The latter predictions are from the high-density rows in Table 1.

A separate ANOVA was conducted on SDs for the same three independent variables. There was a significant main effect of slope $[F(2,4) = 49.82, p < .05, \omega^2 = .434]$. As slope increased, SDs increased, with the mean SDs of 7.94, 13.41, and 16.73 for the three slopes. Post hoc comparisons (Tukey's HSD test) revealed the difference between all three slopes to be significant. The main effects of orientation of the axis of rotation $[F(1,2) < 1]$ and orientation of the dihedral angle $[F(1,2) < 1]$ were not significant, and there were no significant interactions.

Discussion

The principal finding of Experiment 2 was again that subjects tend to interpolate smooth surfaces, rather than surfaces with discontinuities. This tendency appeared in all four combinations of surface orientation with axis orientation. Moreover, this experiment replicated another result of Experiment 1: subjects' judgments were consistent with the models of Kellman and Shipley, Ullman, and Grimson.

EXPERIMENT 3

In both Experiments 1 and 2, the simulated depth of the outer edges of the dihedral angle was kept constant across the experimental conditions. As a result, the simulated distance of the dihedral edge from the axis of rotation varied with the slope of the planes forming the angle. For this reason the speed of the probe dot, if placed on the dihedral edge, would be different in the different slope conditions. Since most of the variance in these experiments was accounted for by the slope, one could ar-

Predicted
Slope   Observed   (two-arc)   (Grimson)

0.66   —▲—   ----△----   --△--
1.00   —■—   ----□----   --□--
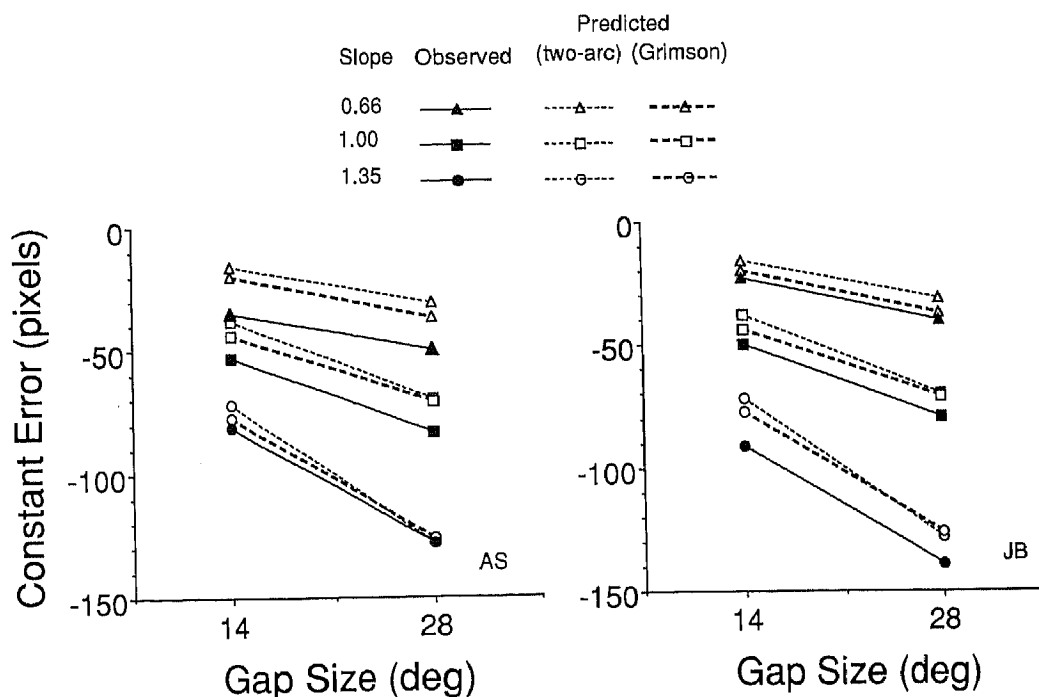1.35   —●—   ----⊖----   --⊖--

Figure 11. Constant error as a function of slope and gap size in Experiment 3. Dashed lines show predicted values.

gue that the slope effects are really due to these differences in probe dot speed. In Experiment 3 we examined this possibility by matching, across levels of slope, the distance between the dihedral edge and the axis of rotation.

Method

The method was the same as in Experiment 1, except as noted below.

Subjects. The subjects were one of the authors (A.S.) and a graduate student (J.B.) who had no knowledge of the purposes of the experiment and was paid for participating. Both subjects also participated in Experiment 2 (A.S. after Experiment 3, J.B. before Experiment 3).

Stimuli. The stimuli were the same as in Experiment 1, except that two values were selected for the distance between the axis of rotation and the dihedral edge, and these two values were used for all three slope conditions. This was accomplished by shifting the dihedral angle along the z-axis (depth axis) as was required in order to achieve the desired distances (see Design). The edges of the dihedral angle closest to the axis of rotation were clipped by an amount proportional to the shift along the z-axis, in order to keep the visible areas of the dihedral angles comparable to those of Experiment 1.

Design. We examined three independent variables: slope (0.66, 1.00, and 1.35), gap size (14° and 28°), and the distance between the dihedral edge and the axis of rotation (600 or 810 pixels). The density was the same as the highest density level used in Experiment 1. Each subject responded 10 times to the 12 conditions.

Procedure. Each subject participated in one practice session and two experimental sessions. The subjects were instructed as in Experiments 1 and 2. The practice session consisted of two repetitions of the 12 conditions presented in a random order, with a rest period halfway through the session. The experimental sessions started with 2 practice trials, followed by 60 randomly ordered trials (12 conditions × 5 repetitions) divided into 2 blocks, separated by a 5-min rest period.

Results

The probe dot was placed inside the dihedral angle on 100% of the trials, and between the dihedral edge and a plane connecting the edges of the gap on 87% of the trials. An ANOVA of the CE data showed a significant main effect for slope [$F(2,2) = 44.59, p < .05, \omega^2 = .714$]. As the slope increased, subjects placed the probe dot farther inside the dihedral angle, with mean CEs of -37, -66, and -110 for the three slopes. (The individual subject means were -42, -68, and -104 for A.S. and -31, -64, and -115 for J.B.) Post hoc comparisons (Tukey's HSD test) showed the differences between all three slopes to be significant.

The main effect of gap size was also significant [$F(1,1) = 383,526.94, p < .01, \omega^2 = .198$]. The mean CEs were -55 and -86 pixels for gaps 14° and 28°, respectively. (The individual subject means were -56 and -87 for A.S. and -55 and -86 for J.B.) There was a significant interaction between slope and gap [$F(2,2) = 584.52, p < .01, \omega^2 = .035$] (see Figure 11). There were no other significant interactions. The main effect of the distance of the dihedral edge from the axis of rotation was not significant [$F(1,1) = 2.10$]. The mean CEs were -70 and -72 for simulated dihedral edge positions of $z = 600$ and $z = 810$. Figure 11 shows predicted CEs based on Kellman and Shipley's and Ullman's models and on Grimson's model. The latter predictions are from the high-density rows in Table 1.

An ANOVA of the SD results showed a significant main effect only for slope [$F(2,2) = 89.17, p < .05, \omega^2 = .302$]. SDs increased with slope, with mean SDs of 9.86,

11.07, and 16.07 pixels, respectively, for the three slopes. Post hoc comparisons (Tukey's HSD test) showed the mean *SD* for the 1.35 slope to be significantly different from the other two levels, 0.66 and 1.00. There were no significant interactions.

Discussion

Although the distances between the dihedral edges and axis of rotation used in Experiment 3, unlike Experiments 1 and 2, were independent of the slope conditions, subjects' judgments were consistent over all three experiments. In Experiment 3, the main effect of slope remained significant and accounted for over 70% of the variance. Also, there were no significant effects involving the distance of the dihedral edge from the axis of rotation, indicating that slope effects in Experiments 1 and 2 were not due to variations in probe dot velocity resulting from variations in position of the dihedral edge.

GENERAL DISCUSSION

All three experiments suggest that, for the displays described here, human vision prefers to interpolate a smooth surface; it does not interpolate a surface with an orientation discontinuity even though such an interpolation would be consistent with the displayed planes. Subjects placed the SFM probe dot inside the angle formed by the two planes on almost every trial in every condition. They placed it farther inside when the planes forming the dihedral angle were more sharply slanted, when the gap was larger, and when the texture density was reduced. The effects of slope and gap were consistent with predictions derived from Kellman and Shipley's (1991) model or from Ullman's (1976) model, which make the same predictions for our stimuli. (These models are concerned with 2-D curvature interpolation and cannot be used to predict 3-D density effects.) The effects of slope and gap were also consistent with Grimson's (1981) 3-D interpolation model. Grimson's model predicts that the subjective surface will be closer to the dihedral edge with increased density, and our results for density are consistent with that prediction. Grimson's model, however, predicts CEs larger than those we obtained for the lower density conditions and smaller than those we obtained for the high-density conditions.

In applying Grimson's (1981) model to our data, we based the grid spacing on the projected gap width, which is a function of the 3-D gap size and the slope of the planes forming the dihedral angle. The density of grid elements thus did not match exactly the densities used in our displays. Even when the densities were similar, it is not clear that matching a random texture with the same number of regularly spaced grid elements is appropriate. Since the fit of Grimson's model to our data can be made better or worse by adjusting the grid density in the model, we did not attempt to make quantitative comparisons between the fit of this model to the data and the fit of Kellman and Shipley's (1991) model or Ullman's (1976) model. Figures 6 and 9 indicate that the fits are qualitatively sim-ilar when we average the observed data and the Grimson predictions over density.

Placement of the probe dot did not depend on the relationship between the orientation of the axis of rotation and the orientation of the planes forming the dihedral angle. This type of relationship, which has been found to be relevant in studies of curvature discrimination (Cornilleau-Pérès & Droulez, 1989; Norman & Lappin, 1992), does not appear to be relevant to surface interpolation. Interpolation was not affected significantly by whether the direction of slant was horizontal or vertical or by whether the axis was horizontal or vertical.

The consistent indication across all three experiments that subjects interpolate a smooth surface, rather than one with an orientation discontinuity, is evidence against another theoretically possible form of interpolation. Subjects could have perceived the planes forming the dihedral angle as continuing unaltered into the gap and could have placed the probe dot at the intersection of these planes. The SFM displays provided sufficient information, theoretically, to recover the 3-D structure of the visible portion of the planes. There were always more than three views, and at least four points present that were not on the same plane, thus permitting application of Ullman's (1979) theorem. Although a model that recovered the structure of the planes and found their intersection could have placed the probe dot accurately on the simulated object, it would not have matched the behavior of our subjects.

Our results suggest that, at least for the SFM displays of dihedral angles studied here, the smooth interpolation which occurs in 3-D is similar to that which occurs in the interpolation of subjective contours in 2-D images (Kellman & Shipley, 1991; Ullman, 1976). This suggests that human vision might use general principles of interpolation across a variety of situations, although the specific processes for interpolation in 2-D and interpolation in 3-D may be different.

Our finding of smooth interpolations in SFM is consistent with results reported by Collett (1985) for interpolation in stereo. Collett's stimuli were random dot stereograms depicting vertically oriented dihedral angles. The dots on the planes forming the angle were presented stereoscopically, but the dots in a region centered about the dihedral edge—the equivalent of the gap in our stimuli—were presented monocularly. The subject's task was to adjust a stereo probe dot in the gap until it appeared to be on the apparent surface. The results indicated that subjects interpolated a smooth surface between the two planes.

Our results should not be taken to imply that subjects are unable to discriminate between smoothly rounded and sharp corners in a 3-D display. Research with random dot stereograms (Norman, Lappin, & Zucker, 1991) and with SFM displays (Zucker & Iverson, 1987) indicate that subjects are quite good at that discrimination. Interpolating a surface between visible features is a very different task from discriminating between visible surfaces.

Our previous study of surface interpolation (Saidpour et al., 1992) used stimuli similar to those of the present study, with one difference: the dots in the previous study lay on the surface of a cylinder rather than on a dihedral angle. In that study, we found that subjects placed the probe dot significantly outside a true cylindrical interpolation. That study, like the present one, indicated that subjects interpolated a smooth surface through the gap. In that study, unlike the present one, the subjects' results were not consistent with Grimson's model. Also, had we applied Ullman's (1976) or Kellman and Shipley's (1991) model to a 2-D cross-section of the simulated surface, those models would not have predicted placement of the probe dot outside of the cylindrical interpretation. We do not have a specific explanation for this discrepancy, but there is an important difference between the two studies that may be relevant: The tangent to the visible surface at the edge of the gap is used in the interpolation process in both Kellman and Shipley's model and Ullman's model. Since the visible surface in the dihedral angle has constant slope, information for computing the tangent is available over the entire surface. The tangent to the cylinder, on the other hand, changes continuously. If the tangent at the gap were estimated from feature points over a larger area of the surface, a bias might result in the case of the cylinder that would not occur for the dihedral angles. This explanation is speculative, but it indicates that there is at least one important difference between cylinders and dihedral angles that may be related to interpolation.

One might question how a quantitative study of SFM interpolation can be meaningful when Todd and his collaborators (Todd, Akerstrom, Reichel, & Hayes, 1988; Todd & Bressan, 1990; Todd & Norman, 1991) have provided evidence that depth is recovered only up to an affine transformation in SFM displays. Rescaling of depth along the line of sight, however, would not necessarily affect the results that we obtained with our SFM probe. Since the probe dot moved rigidly with the SFM display (except while the subject was adjusting it), rescaling of depth along the line of sight, provided that it was a constant rescaling, would not change the relative positions in depth of the simulated dihedral angle and the probe dot. Although subjects may not perceive the simulated structure accurately in an SFM display, they do perceive a 3-D structure in which depth is systematically related to image variables (Liter, Braunstein, & Hoffman, in press) and the relative judgments obtained in the present research should be informative about interpolation within that perceived structure.

Finally, we should point out that SFM interpolation can be quite precise. In Experiment 1, the mean SD across replications was 2.2% of the depth in the simulated dihedral angles in the high-density condition and 2.7% of this depth overall.

To further advance our understanding of surface interpolation, we will need to study different types of surface discontinuities along with more complex shapes (e.g.,

shapes that vary in slant or curvature in more than one dimension). It should be especially useful to study nonsymmetric objects in which the magnitude of the slant or curvature changes from one side to the other of a region devoid of feature points. In this case, Kellman and Shipley's (1991) and Ullman's (1976) models, applied to the 2-D cross-section, would provide different predictions. It will also be important to examine interpolation with SFM combined with other sources of depth information, such as binocular disparity and shape from shading.

REFERENCES

BLAKE, A. (1984). Reconstructing a visible surface. In *Proceedings of AAAI Conference 1984* (pp. 23-26). Los Altos, CA: AAAI Press.

BLAKE, A., & ZISSERMAN, A. (1986). Invariant surface reconstruction using weak continuity constraints. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (pp. 62-67). Washington, DC: IEEE Computer Society Press.

BLAKE, A., & ZISSERMAN, A. (1987). *Visual reconstruction*. Cambridge, MA: MIT Press.

BRADY, M., & GRIMSON, W. E. L. (1981). *The perception of subjective surfaces* (Memo No. 666). Cambridge, MA: Massachusetts Institute of Technology, Artificial Intelligence Laboratory.

BRADY, M., & HORN, B. K. P. (1981). *Rotationally symmetric operators for surface interpolation* (Memo No. 654). Cambridge, MA: Massachusetts Institute of Technology, Artificial Intelligence Laboratory.

BRAUNSTEIN, M. L. (1966). Sensitivity of the observer to transformations of the visual field. *Journal of Experimental Psychology*, 72, 683-687.

BRAUNSTEIN, M. L. (1978). Perception of motion. In E. Carterette & M. Friedman (Eds.), *Handbook of perception* (Vol. 8, pp. 147-172). New York: Academic Press.

BRAUNSTEIN, M. L. (1983). Perception of rotation in depth: The psychophysical evidence. In N. I. Badler & J. K. Tsotsos (Eds.), *Motion: Representation and perception* (Proceedings of the ACM SIGGRAPH/SIGART Interdisciplinary Workshop on Motion, Representation and Perception, pp. 242-247). New York: North Holland.

BRAUNSTEIN, M. L. (1988). The empirical study of structure from motion. In W. N. Martin & J. K. Aggarwal (Eds.), *Motion understanding: Robot and human vision* (pp. 101-142). Hingham, MA: Kluwer Academic Publishers.

BRAUNSTEIN, M. L., & ANDERSEN, G. J. (1984). Shape and depth perception from parallel projections of three-dimensional motion. *Journal of Experimental Psychology: Human Perception & Performance*, 10, 749-760.

COLLETT, T. S. (1985). Extrapolating and interpolating surfaces in depth. *Proceedings of the Royal Society of London: Series B*, 224, 43-56.

CORNILLEAU-PÉRÈS, V., & DROULEZ, J. (1989). Visual perception of surface curvature: Psychophysics of curvature detection induced by motion parallax. *Perception & Psychophysics*, 46, 351-364.

GRIMSON, W. E. L. (1981). *From images to surfaces: A computational study of the human early visual system*. Cambridge, MA: MIT Press.

HALPERN, D. F., & SALZMAN, B. (1983). The multiple determination of illusory contours: 1. A review. *Perception*, 12, 281-291.

HORN, B. K. P. (1983). The curve of least energy. *Association for Computing Machinery Transactions on Mathematical Software*, 9, 441-460.

JULESZ, B. (1971). *Foundations of cyclopean perception*. Chicago: University of Chicago Press.

JULESZ, B., & FRISBY, P. (1975). Some new subjective contours in random-line stereograms. *Perception*, 4, 145-150.

KANIZSA, G. (1979). *Organization in vision*. New York: Praeger.

KELLMAN, P. J., & SHIPLEY, T. F. (1991). A theory of visual interpolation in object perception. *Cognitive Psychology*, 23, 141-221.

LITER, J. C., BRAUNSTEIN, M. L., & HOFFMAN, D. D. (in press). Inferring structure from motion in two-view and multi-view displays. *Perception*.

NORMAN, J. F., & LAPPIN, J. S. (1992). The detection of surface curvatures defined by optical motion. *Perception & Psychophysics*, **51**, 386-396.

NORMAN, J. F., LAPPIN, J. S., & ZUCKER, S. W. (1991). The discriminability of smooth stereoscopic surfaces. *Perception*, **20**, 789-807.

SAIDPOUR, A., BRAUNSTEIN, M. L., & HOFFMAN, D. D. (1992). Interpolation in structure from motion. *Perception & Psychophysics*, **51**, 105-117.

STEVENS, K. A., & BROOKES, A. (1987). Probing depth in monocular images. *Biological Cybernetics*, **56**, 355-366.

TERZOPOULOS, D. (1982). *Multi-level computation of visual surface representation* (Memo No. 671). Cambridge, MA: Massachusetts Institute of Technology, Artificial Intelligence Laboratory.

TERZOPOULOS, D. (1986). Integrating visual information from multiple sources. In A. P. Pentland (Ed.), *From pixels to predicates: Recent advances in computational vision* (pp. 111-142). Norwood, NJ: Ablex.

TODD, J. T., AKERSTROM, R. A., REICHEL, F. D., & HAYES, W. (1988). Apparent rotation in three-dimensional space: Effects of temporal, spatial, and structural factors. *Perception & Psychophysics*, **43**, 179-188.

TODD, J. T., & BRESSAN, P. (1990). The perception of 3-dimensional affine structure from minimal apparent motion sequences. *Perception & Psychophysics*, **48**, 419-430.

TODD, J. T., & NORMAN, J. F. (1991). The visual perception of smoothly curved surfaces from minimal apparent motion sequences. *Perception & Psychophysics*, **50**, 509-523.

ULLMAN, S. (1976). Filling-in the gaps: The shape of subjective contours and a model for their generation. *Biological Cybernetics*, **25**, 1-6.

ULLMAN, S. (1979). *The interpretation of visual motion*. Cambridge, MA: MIT Press.

UTTAL, W. R., DAVIS, N. S., WELKE, C., & KAKARALA, R. (1988). The reconstruction of static visual forms from sparse dotted samples. *Perception & Psychophysics*, **43**, 223-240.

WALLACH, H., & O'CONNELL, D. N. (1953). The kinetic depth effect. *Journal of Experimental Psychology*, **45**, 205-217.

WEISS, I. (1990). Shape reconstruction on a varying mesh. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, **12**, 345-362.

ZUCKER, S. W., & IVERSON, L. (1987). From orientation selection to optical flow. *Computer Vision, Graphics, & Image Processing*, **37**, 196-220.