Perception & Psychophysics
1992, 51 (4), 386-396

The detection of surface curvatures defined by optical motion

J. FARLEY NORMAN
Brandeis University, Waltham, Massachusetts

and

JOSEPH S. LAPPIN
Vanderbilt University, Nashville, Tennessee

The detectability of surface curvatures defined by optical motion was evaluated in three experiments. Observers accurately detected very small amounts of curvature in a direction perpendicular to the direction of rotation, but they were less sensitive to curvatures along the direction of rotation. Variations in either the number of points (between 91 and 9) or the number of views (from 15 to 2) had little or no effect on discrimination accuracy. The results of this study demonstrate impressive visual sensitivity to surface curvature. Several characteristics of this sensitivity to curvature are inconsistent with many computational models for deriving three-dimensional structure from motion.

The results of nearly a century of research have shown that optical motion is a powerful source of information about an object's three-dimensional (3-D) shape. To this day, however, it remains unclear exactly which 2-D image properties are detected and exactly which 3-D structural relations are perceived when observers view structure-from-motion displays. As recently as 1989, Sperling, Landy, Dosher, and Perkins asked : "Does the observer perceive the correct shape in a display? The correct depths? The correct depth order? The correct curvature?" (p. 826). Contemporary research, they argued, has failed to resolve such issues.

Most current computational models operate on a number of discrete points contained within a set of distinct "views." These models typically assume that appropriate correspondences have been established between the same physical object points across the discrete views. Given this assumption, such models recover depth and/or orientation values for each identifiable feature. When the set of all such pointwise depths has been recovered, a primary computational problem has been solved, but other substantial problems remain. What remains is to parse this collection of depths into separate objects and to interpolate smooth surfaces between connected feature points. Object surfaces and properties such as curvature are thought to be secondarily derived from the more elementary features and their depths.

One problematic phenomenon for such models is the finding of directional anisotropies. Rogers and Graham (1983) and Rogers and Cagenello (1989) have documented anisotropies for both stereopsis and motion parallax. Rogers and Graham showed that a depth analogue of the Craik–O'Brien–Cornsweet luminance illusion occurred for surfaces defined by either motion parallax or stereopsis. They found that the illusion was obtained only when the depths varied in the horizontal direction. The illusion did not occur when the depth variations were in the vertical direction. Similarly, Rogers and Cagenello showed that thresholds for discriminating whether stereoscopic cylindrical surfaces were convex or concave depended strongly on the orientation of the cylinders' axis of symmetry: thresholds for vertically oriented cylinders (with horizontal depth variations) were *twice* as high as those for horizontally oriented cylinders (vertical depth variations). Cornilleau-Pérès and Droulez (1989) found a similar anisotropy for cylindrical surfaces defined by motion where the curvatures were much less accurately detected when the cylinders were curved in the direction of rotation than when they were curved in a perpendicular direction.

The existence of such anisotropies challenges models of structure from motion that operate on single points. These models always recover the relative depths of a moving set of points provided the motion satisfies certain constraints; how the surface is oriented relative to the axis of rotation is irrelevant. For example, Ullman's (1979) model requires three distinct views of four noncoplanar points moving rigidly. So long as these conditions are fulfilled, accurate recovery of the 3-D structure will occur.

This research was supported in part by National Research Service Award EY-07007-10 to the first author and NIH Grant EY-05926 to the second author. Many of the analyses accompanying the empirical results contained in this manuscript were aided by Grant 89-0016 from the Air Force Office of Scientific Research, Grant BNS-8908426 from the National Science Foundation, and grants from the Office of Naval Research, and the Air Force Office of Scientific Research to James Todd. We are grateful to James Todd for his helpful comments on an earlier version of this manuscript. Correspondence should be addressed to J. F. Norman, Department of Psychology, Brandeis University, Waltham, MA 02254-9110.

An alternative conception of the processes for detecting structure from motion and from stereopsis is one in which the surface structure is fundamental rather than being derived from a depth map. Several researchers (Koenderink, 1990a, 1990b; Cornilleau-Pérès & Droulez, 1989; Droulez & Cornilleau-Pérès, 1990; Norman, Lappin, & Zucker, in press; Rogers & Cagenello, 1989; Stevens & Brookes, 1988; Sander & Zucker, 1990; Zucker, 1986) have suggested that, for the perception of 3-D shape, the depth values themselves are less important than how they vary over the surface of a 3-D object. For example, is a given surface region flat or curved? If curved, is it a region of positive curvature (a convexity or concavity) or negative (a saddle-like shape)? Does the surface contain discontinuities? If so, where?

The anisotropy of orientation found for both structure from motion and stereopsis would not be expected a priori from a sensitivity to surface structure or curvature if accurate surface curvatures could be recovered in all circumstances. However, the anisotropy might be thought to derive from the ease with which the visual system can detect changes in depth (i.e., derivatives) in various directions relative to the direction of motion. Evidently, curvatures or changes in depth are more easily detected in a direction perpendicular to the direction of motion. Indeed, Braunstein (1977) and Braunstein and Andersen (1984) concluded from their experiments that "variations in the dimension of the axis of rotation and variations in the perpendicular dimension represent separate sources of information about rotating objects" (Braunstein & Andersen, 1984, p. 758).

The present study was part of a larger project to evaluate visual detections and discriminations of surface curvature. The experiments addressed the following questions about the detectability of surface curvature: (1) Does the previously reported anisotropy for detecting 3-D structure from motion apply to discriminations involving shapes other than cylinders and planes? (2) What is the minimal amount of curvature required to discriminate a curved from a planar surface? (3) How does the detectability of curvature depend on the number and distribution of points over the surface? (4) Are two views sufficient for accurate detections of surface curvature, as suggested by Droulez and Cornilleau-Pérès (1990), Koenderink and van Doorn (1991), and Todd and Bressan (1990)? Most alternative computational models (e.g., Hoffman & Bennett, 1985, 1986; Ullman, 1979) need three "distinct" views to recover the 3-D structure of a moving object.

EXPERIMENT 1

This experiment evaluated the relative discriminabilities of four different types of surfaces—sphere, vertical cylinder, horizontal cylinder, and plane. These four surfaces differed in their curvatures in the horizontal and vertical directions. Discrimination accuracies were obtained for all six pairs of these surfaces. The two differently oriented cylinders were used in order to determine if equal curvature magnitudes have similar effects in different orientations. See Figure 1 for a schematic description of the four different surfaces and their position relative to the axis of rotation.

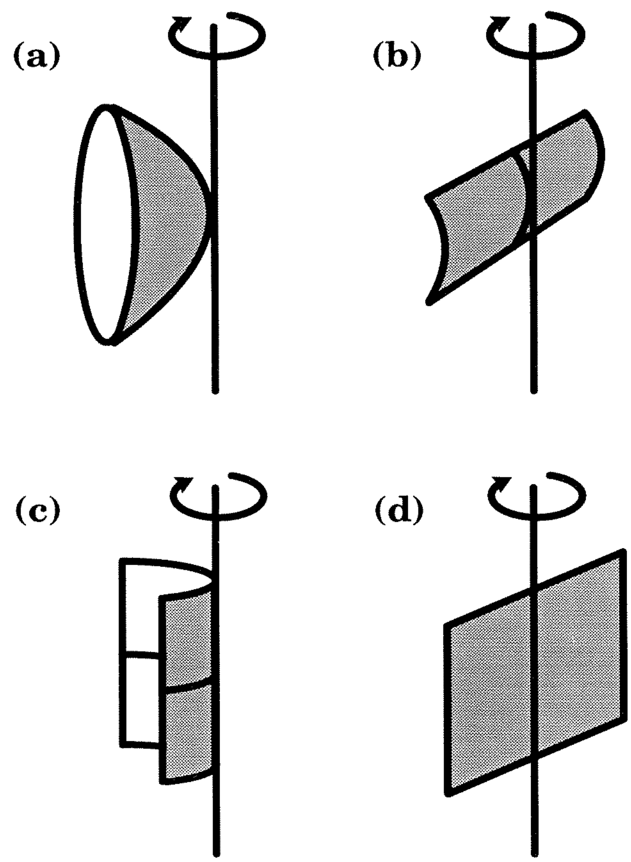

Figure 1. A schematic illustration of the four different types of surfaces used in Experiment 1 and their positions relative to the axis of rotation: (a) spherical surface patch, (b) horizontally oriented cylindrical surface patch, (c) vertically oriented cylindrical surface patch, and (d) planar surface patch. The center of each surface patch fell upon the axis of rotation. As a consequence, image velocities were always zero toward the center of the optical pattern and smoothly increased toward the edge of the pattern. The actual surfaces used in the experiments were defined by the motions of an array of luminous points against a dark background. The points within the optical pattern were arranged into a perturbed hexagonal lattice, as shown in Figure 2.

Method

Stimulus displays. A relatively small surface patch from a large physical object was rotated around a Cartesian vertical axis. The radii of the cylindrical and spherical surface patches were 25.0 cm (¼ m). Curvature along a given direction at a point is defined as the reciprocal of the radius of the best-fitting circle at that point (see Hilbert & Cohn-Vossen, 1952). The surface patches used in this experiment had curvatures of 4.0 m^{-1}.

Ninety-one points were arranged into a rough hexagonal lattice in the frontoparallel plane in order to ensure that the surface be adequately sampled. The height and width of the pattern subtended $2° \times 1.73°$ of visual angle. The average separation between adjacent lattice points subtended 12′. Each point's position was ran-

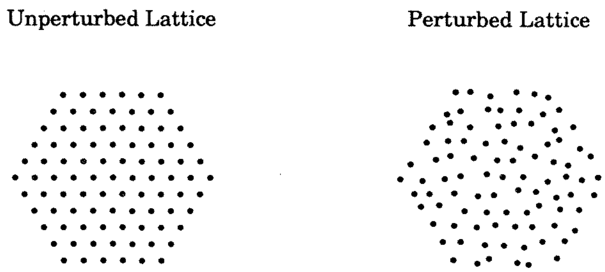

Figure 2. A schematic illustration of the 91-point hexagonal lattice used as the optical pattern for Experiment 1. The left half of the illustration shows the original unperturbed hexagonal lattice. The right half shows an example of a typical pattern used in Experiment 1. As the illustration indicates, significant noise perturbation was added to the positions of the points, breaking the symmetry of the lattice. A different, randomly determined noise perturbation was applied to every optical pattern on every trial.

domly displaced by up to 3′ of visual angle both vertically and horizontally away from a perfect hexagonal lattice to eliminate potential texture-density differences between patches from different 3-D shapes. This 2-D hexagonal lattice was then projected onto the curved 3-D surface. A schematic illustration of both the original unperturbed hexagonal lattice and the noisy lattice used in Experiment 1 are shown in Figure 2.

The apparent-motion sequences consisted of 15 successive views, each separated by a 5° rotation for a total rotation extent of 70°. The rotations were around a Cartesian vertical axis centered at the front surface of the patch. The initial view for every trial consisted of the 2-D hexagonal lattice to which random positional noise had been added. The surface then rotated in one direction for 7 views, reversed direction, rotated for 15 views, reversed direction again, and rotated for 7 views to return to the starting position. The stimulus pattern on each trial consisted of three cycles of this sequence for a total of 90 views. Each view was presented for approximately 28 msec. The total trial duration was approximately 2.5 sec.

The resulting velocity fields produced by these rotating surfaces were qualitatively very similar. A point at the exact center of the optical pattern would fall directly upon the axis of rotation and would not move. Differential velocities across trials existed for all other surface points not falling on the axis of rotation due to the surfaces' varying 3-D shapes.

The stimuli were generated by a Macintosh IIx computer. The 2-D positions of the display's constituent dots, calculated by the computer, were converted into analog voltages by a MacAdios D/A (digital-to-analog) converter. These analog waveforms were then displayed on a Tektronix 608 cathode ray tube (CRT) monitor with P-31 phosphor. The horizontal resolution of the displays was 16 bits, or 65536 positions; vertical resolution was 12 bits, or 4096 positions. The Tektronix monitor was mounted on an optical bench and was viewed by the observer from a distance of 114.6 cm. The positions of the display's constituent points were calculated using the correct perspective projection for the observer's viewing position. Because the depths of the points relative to the projected image plane were small, as compared with the observer's viewing distance, and the surfaces were oscillated in depth around an axis in the image plane rather than rotated through 360°, the effects of perspective were small (the maximum perspective ratio was 1.02; for parallel projection the ratio equals 1.0) (see Braunstein, 1962, for more details on perspective).

The observers viewed the motion sequence monocularly inside a dimly lit room. The optical patterns were viewed through a 10.0-cm-diam (5.0° of visual angle) circular aperture in a black sheet of plastic that occluded the edges of the Tektronix monitor. The observer's head position was constrained during the presentation of each trial.

Psychophysical task. The psychophysical task was to discriminate between two qualitatively different curved surfaces. The two surfaces to be discriminated were different for each of the six experimental conditions.

Within any given session and for any given experimental condition, observers were presented with a randomly ordered sequence of trials containing the two surfaces to be discriminated. On any given trial, a single 3-D surface was presented. The observer's task was to indicate, by pressing either of two keys on a computer keyboard, which of the two surfaces had been presented. Auditory feedback was provided when the observer's response was correct.

Experimental conditions. Six experimental conditions were formed from all pairwise combinations of four basic surfaces: the sphere, the horizontally oriented cylinder, the vertically oriented cylinder, and the plane.

Procedures. Each observer participated in four separate experimental sessions, each consisting of six blocks of 50 trials. A total of 200 trials was therefore obtained for each of the six different discrimination tasks. The order of the six discrimination tasks was determined randomly for each observer and was counterbalanced across the first two experimental sessions. After the first 100 trials had been gathered for each condition, a new random order of conditions was established for the third experimental session. The order of those conditions was counterbalanced across Sessions 3 and 4.

The 3 observers were graduate students in the Department of Psychology at Vanderbilt University. One was the first author. The other 2 observers were naive as to the purpose of the experiment.

Results

The results (shown in Table 1) show two basic phenomena: (1) a strong *anisotropy* of orientation and (2) discriminations involving surfaces with unidirectional curvature were nearly as accurate as those involving bidirectional curvature. Reliable differences between the six shape-discrimination tasks were indicated by a Friedman rank test for correlated samples using the four sessions for each of the 3 observers as independent replications [$\chi^2_F(5) = 48.9$, $p < .001$].[1]

Two comparisons show an anisotropy of curvature detections. The first comparison indicates that discriminations were much more accurate between the horizontal cylinder and plane than between the vertical cylinder and plane (Wilcoxon matched-pairs signed-ranks test: $T = 0$,

Table 1
Combined Observer Performance for Each of the Six Shape-Discrimination Tasks Used in Experiment 1

	Discrimination Accuracy	
Shape Discrimination Task	% correct	$-\ln \eta$
Sphere vs. Plane	96.5	3.27
Sphere vs. Horizontal cylinder	58.9	0.36
Sphere vs. Vertical cylinder	89.0	2.08
Horizontal vs. Vertical cylinder	87.3	1.94
Horizontal cylinder vs. Plane	91.5	2.36
Vertical cylinder vs. Plane	60.2	0.41

Note—The far right column shows the observers' combined discrimination accuracies in terms of $-\ln \eta$, a measure of discriminal distance developed by Luce (1963).

$n = 12, p < .001$). This finding is analogous to that of Cornilleau-Pérès and Droulez (1989).

The results also show a second anisotropy: discriminations were more accurate between the sphere and vertical cylinder than between the sphere and horizontal cylinder (Wilcoxon signed-ranks test: $T = 0$, $n = 12$, $p < .001$). The earlier results of Cornilleau-Pérès and Droulez documented the existence of the anisotropy involving cylindrical and planar surfaces. The results of this experiment confirm and extend their findings. The current results also show that large perceptual anisotropies exist between other pairs of curved surfaces. The finding of directional anisotropies is not limited to the discrimination between cylinders and planes, but appears to be more general.

Table 1 shows the combined discrimination accuracies for each task in terms of both percent correct and $-\ln \eta$, a measure of discriminability developed by Luce (1963, pp. 113-116, 123-125). This measure is similar to the d' of signal detection theory in that it has many properties

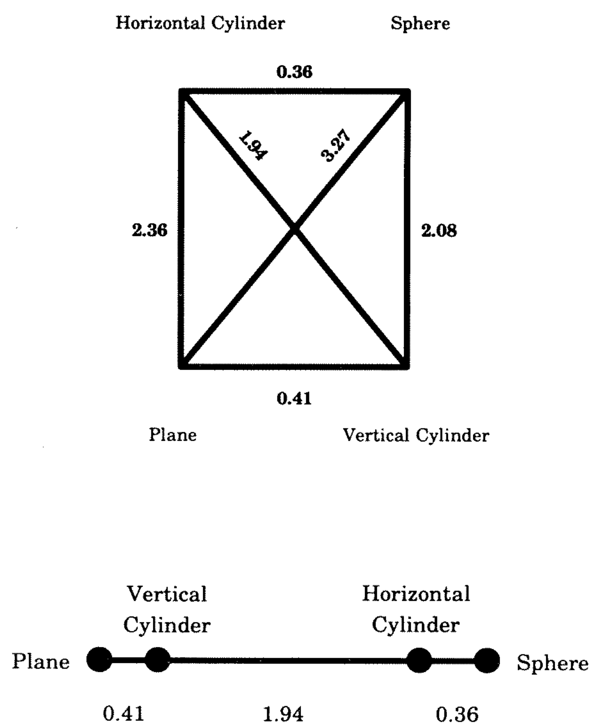

Figure 3. A diagram illustrating the relationships between the observers' combined discrimination accuracies expressed in terms of $-\ln \eta$. In the upper half of the diagram, surfaces that have different curvatures along the direction of rotation are separated horizontally, and surfaces that have different curvatures perpendicular to the direction of rotation are separated vertically. The six line segments connecting the four shapes represent the six discrimination tasks used in Experiment 1. Beside each of these six line segments are plotted that task's discrimination accuracy in terms of $-\ln \eta$. If the length of the line segments were made proportional to the performance in each task, then no spatial configuration could exist. This result indicates that no single measure can completely account for the entire patterns of results. The lower half of the diagram places the four basic surfaces upon a single continuum in a way that is approximately consistent with the pattern of results.

of a distance measure. When performance for the six discrimination tasks is examined using this distance measure (which takes into account the "hit" and "false-alarm" rates), it is evident that horizontal and vertical curvatures are not detected by independent visual mechanisms. Consider the upper half of Figure 3. This figure illustrates the relationships between the different discrimination tasks involving the four different surfaces. Surfaces that differ in their curvatures along the direction of rotation are separated horizontally. Surfaces that differ in their curvatures perpendicular to the direction of rotation are separated vertically. Surfaces that differ in both directions are separated diagonally. The six lines connecting the four surfaces represent the six different discrimination tasks. The discriminability in terms of $-\ln \eta$ is plotted adjacent to each line. If curvatures in and perpendicular to the direction of rotation were detected by independent mechanisms, then various relationships should exist between the discriminabilities of the six pairs of surfaces according to the Pythagorean theorem. If we represent the discriminability of a pair of surfaces as $d(xy)$, where x and y indicate the two surfaces (P = plane, V = vertical cylinder, H = horizontal cylinder, S = sphere), then, in particular, we should have:

$$d(\mathrm{PS})^2 = d(\mathrm{PV})^2 + d(\mathrm{VS})^2,$$
$$d(\mathrm{PS})^2 = d(\mathrm{PH})^2 + d(\mathrm{HS})^2,$$
$$d(\mathrm{VH})^2 = d(\mathrm{PV})^2 + d(\mathrm{PH})^2$$
$$d(\mathrm{VH})^2 = d(\mathrm{VS})^2 + d(\mathrm{HS})^2,$$
$$d(\mathrm{VH}) = d(\mathrm{PS}).$$

Clearly, several of these relationships are violated in the present set of results. Indeed, if the length of the line segments representing the six discrimination tasks is made proportional to the observers' performance, no spatial arrangement is possible. One cannot find a spatial arrangement consistent with the results because the triangle inequality is violated, such that $d(\mathrm{PS}) > d(\mathrm{PV}) + d(\mathrm{VS})$ and $d(\mathrm{PS}) > d(\mathrm{PH}) + d(\mathrm{HS})$.

In the bottom half of Figure 3, we have placed the four shapes along a single dimension. The sphere and the plane occupy opposite ends of the continuum. By placing the vertical cylinder at a distance from the plane appropriate to its discriminability and placing the horizontal cylinder at a distance from the sphere appropriate to its discriminability, we can capture most, but not all, of the remaining observed relationships. In particular, $d(\mathrm{PH})$ approximately equals $d(\mathrm{PV}) + d(\mathrm{VH})$, and $d(\mathrm{VS})$ approximately equals $d(\mathrm{HS}) + d(\mathrm{VH})$. The single relationship not adequately described by such a representation is the discriminability between sphere and plane. The sphere is much more discriminable from the plane than would be predicted by adding $d(\mathrm{PV})$, $d(\mathrm{VH})$, and $d(\mathrm{HS})$. This result raises the possibility that some other property, such as the symmetry of the spherical surface, is perceivable, enhancing its discriminability above that predicted by the combination of the sphere's curvatures in the two orthogonal directions.

To demonstrate that the pattern of results between the six shape-discrimination tasks was not due to any simple differences in depth or orientation per se, we conducted another set of analyses. These analyses calculated the total depth and surface-orientation differences between all pairs of surfaces. If 3-D surfaces were represented by a depth or orientation map, as many computational analyses suggest, then one might conclude that the most effective way to distinguish between two *different* 3-D surfaces would be to compare the depths or orientations of the two surfaces at all corresponding surface locations. A single measure expressing how different two surfaces are could be obtained by simply summating (i.e., integrating) the depth or orientation differences between corresponding surface locations across the surface. Geometrically, the total depth index would correspond to the volume between the two surfaces.

These total depth and orientation indices were calculated for all pairs of surfaces used in Experiment 1. Identical optical patterns (i.e., no noise perturbation was added to the positions of the points within the hexagonal lattice) were used for both surfaces in each pair. The depth analysis was conducted for each pair of surfaces at both their unrotated and maximally rotated positions. The orientation analysis was performed only for the unrotated position, since this measure is invariant under absolute orientation of the two surfaces relative to the observer. The results of these analyses are shown in Table 2. It is readily apparent that one would not expect a large perceptual anisotropy given the depth or orientation differences that do exist between these surfaces. For example, the depth and orientation differences that exist between vertical cylinder and plane are as large as those that exist between horizontal cylinder and plane. The sphere and horizontal cylinder and the sphere and vertical cylinder pairs also have similar depth and orientation differences.

Relative depths and orientations are *physical* properties of 3-D objects and surfaces. We have seen that these depth and orientation differences cannot explain the significant anisotropies that occur during curvature discriminations. It is possible that some simple optical property unrelated to 3-D shape per se could account for the observed pattern of results. To investigate this possibility, we calculated a measure expressing the total velocity difference between all pairs of optical patterns used during the actual experiment. This measure of total velocity is directly analogous to the total depth and orientation measures described earlier. The velocity difference between points corresponding to the same surface location on the two different 3-D shapes was summated across the area of the optical pattern to give a measure expressing the total velocity difference. A different velocity measure, the maximum velocity difference, was calculated for all pairs of optical patterns. For each pair of optical patterns, the surface location with the largest velocity difference was identified. Both of these velocity measures are shown for all surface pairs in Table 3. Similar to the results of the depth and orientation analyses, it is readily apparent that simple velocity differences are also insufficient to explain the observed anisotropies.

Discussion

The results of this experiment show that strong anisotropies occurred when the two differently curved cylinders were discriminated from flat planar surfaces as well as when they were discriminated from bidirectionally curved

Table 2
Depth and Orientation Analyses Conducted on the Six Pairs of Surfaces Used in Experiment 1

	Total Depth Index (cm)		
Surface Pair	Unrotated Position	Rotated Position	Total Orientation Index (radians)
Sphere vs. Plane	3.65	2.99	4.86
Sphere vs. Horizontal cylinder	1.83	1.50	3.11
Sphere vs. Vertical cylinder	1.83	1.50	3.05
Horizontal vs. Vertical cylinder	2.33	1.91	4.86
Horizontal cylinder vs. Plane	1.83	1.50	3.05
Vertical cylinder vs. Plane	1.83	1.50	3.11

Table 3
Velocity Analyses Conducted on the Six Pairs of Surfaces Used in Experiment 1

	Total Velocity Difference (min/sec)		Maximum Difference in Velocity (min/sec)	
Surface Pair	Unrotated Position	Rotated Position	Unrotated Position	Rotated Position
Sphere vs. Plane	242.8	268.6	6.97	6.12
Sphere vs. Horizontal cylinder	133.7	134.4	6.95	6.12
Sphere vs. Vertical cylinder	122.9	134.4	5.23	4.50
Horizontal vs. Vertical cylinder	167.5	171.6	6.95	6.12
Horizontal cylinder vs. Plane	127.6	134.3	5.23	4.50
Vertical cylinder vs. Plane	147.1	134.3	6.95	6.12

surfaces. These anisotropies should not exist if perceived shape were determined only by depth or orientation differences, because the vertical and horizontal cylinders were similar in terms of their distributions of depths and orientations. Depth and orientation values simply varied in different directions within the two shapes. Informal observations demonstrated that these anisotropies were not related to absolute orientation (horizontal or vertical), but were instead related to the direction of rotation. When the surfaces were rotated about a horizontal axis, instead of a vertical one, the pattern of results reversed. In this case, for example, discriminations between the vertical cylinder and plane were easier than those between the horizontal cylinder and plane.

Another phenomenon noticed by all 3 observers in Experiment 1 was that the planar surface often appeared nonrigid. The plane's 3-D structure and its motion in space seemed ambiguous. The plane sometimes appeared as an expanding and contracting 2-D pattern, rather than as a rigid object undergoing a 3-D rotation. It sometimes appeared to be two planes connected to form a dihedral angle with the apex at the axis of rotation where the angle appeared to continuously vary (i.e., nonrigid, like a hinge) during the presentation of a trial. The vertical cylinder sometimes appeared ambiguous as well. In contrast, the sphere and horizontal cylinder were never subject to alternative interpretations. The presence of curvature in a direction orthogonal to the direction of rotation was necessary for the consistent perception of a 3-D shape rotating rigidly in 3-space.

Our results, involving discriminations between differently curved surfaces, confirm and extend Cornilleau-Pérès and Droulez's (1989) and Rogers and Graham's (1983) earlier findings of orientational anisotropies. Droulez and Cornilleau-Pérès (1990) and Cornilleau-Pérès and Droulez (1989) developed an algorithm that is sensitive to differences in curvature which involves taking second derivatives of the first-order optic-flow field (i.e., the velocity field). They define a measure, the *spin variation*, which calculates the second derivative of the orthogonal component of the velocity field for all possible orientations in the image plane. The spin variation measures the bending of a straight line in the image under motion (other recent models utilizing the bending of lines or collinear points in the image to recover information about the curvature of objects would include those by Lappin, Craft, & Tschantz, 1991, Rogers & Cagenello, 1989, and Weinshall, 1991). The individual spin-variation measures for each orientation are combined (such as mean absolute value) into a single measure whose magnitude is influenced by the curvature of the surface, the orientation of the surface, and how the surface is moving relative to the observer. This algorithm is able to determine whether a given local region is planar, cylindrical, ellipsoidal, or saddle-shaped (a hyperboloid).

Cornilleau-Pérès and Droulez (1989) and Droulez and Cornilleau-Pérès (1990) showed that the spin variation would predict their observed anisotropy involving discriminations between cylindrical and planar surfaces. To evaluate their approach for discriminations between pairs of differently curved, nonplanar surfaces, we constructed optical patterns that were similar to the ones used in Experiment 1 but more amenable to their analysis. The spin variation was calculated for each type of surface (sphere, horizontal cylinder, vertical cylinder, and plane) using parameters (curvature, amount of rotation, etc.) that were the same as those used in Experiment 1. The optical pattern used for this analysis is shown in Figure 4. It was composed of six sets of 11 collinear points. In the first view (left), the points are collinear. The second view (right) shows the projection after a spherical surface patch has been rotated in depth (to make the bending of image lines clearly visible, the curvature of the surface and the magnitude of rotation is higher than that used in the actual simulation). Note how the formerly collinear points have been bent in the second view.

The difference between the mean spin variations obtained for the two different shapes used for each pairwise discrimination task was correlated with the observers' actual performance (Pearson r). It was found that the difference in mean spin variation correlated 0.80 when the observers' performance was expressed in terms of percent correct and 0.89 when performance was measured in terms of $-\ln \eta$. It appears that the recent model proposed by Droulez and Cornilleau-Pérès (1990) (and presumably other similar models proposed by Lappin et al., 1991, Rogers & Cagenello, 1989, and Weinshall, 1991) can account quite well for most aspects of our observers' performance. This was true not only for discriminations between curved and noncurved surfaces, as had been shown previously, but also for discriminations involving bidirectionally versus unidirectionally curved surfaces. The results of the $-\ln \eta$ analysis showed, however, that any single measure (such as averaging all of the individual spin-variation measurements taken at each orientation into a single combined measure) is probably insufficient to explain all aspects of curvature detection.

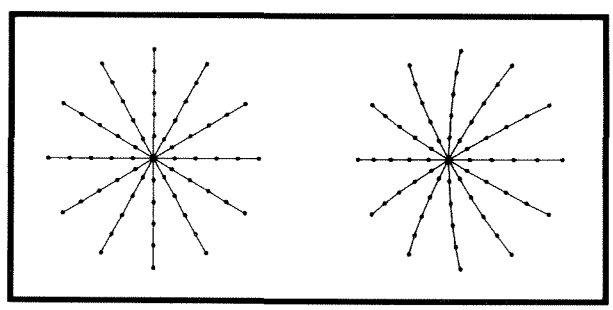

Figure 4. Two views of the optical pattern used to calculate the spin-variation measure developed by Droulez and Cornilleau-Pérès (1990). The spin variation indicates how lines in the 2-D projection bend as the 3-D shape moves relative to the observer. Note how the formerly straight image lines in the left view bend in the subsequent right view due to the object's spherical shape. These images can be stereoscopically free-fused.

EXPERIMENT 2

The preceding experiment documented the existence of perceptual anisotropies whereby curvature perpendicular to the direction of motion was more detectable than curvature along the direction of motion. This second experiment evaluated the *precision* with which observers could discriminate between curved and noncurved surfaces. That is, as the magnitude of the curvature of a spherical surface patch is decreased, at what point does that surface become indistinguishable from a flat surface that lacks curvature? Psychometric functions were obtained showing the observers' discrimination performance as a function of curvature.

This experiment also evaluated how much optical support was necessary for accurate discrimination between differently curved surfaces. The optical patterns used in the previous experiment contained many points. This experiment evaluated the effect of decreases in the number of points on the accuracy of shape discrimination. Is surface curvature less detectable in optical patterns containing few points?

Method

Stimulus displays. The stimulus displays in this experiment were similar to those used in Experiment 1. Differing curvatures and numbers of points were utilized. The three patterns were a 91-point hexagonal lattice, a 19-point hexagonal lattice, and a 9-point diagonal cross. The 91- and 19-point patterns covered the same hexagonal area. Average point spacings were 12′ and 30′ of arc for the 91- and 19-point hexagonal lattices, respectively, and 42′ of arc for the 9-point cross. The cross was composed of two intersecting perpendicular line segments, each containing 5 points. The 9-point cross was chosen after extensive pilot observation, during which it was found that the placement of points became critical when the optical patterns were defined by small numbers of points. Some arrangements of the small number of points appeared ambiguous and did not appear to define a surface. For example, a random 2-D arrangement of 9 points over the $3.46°^2$ area, when back-projected onto the 3-D spherical surface and rotated, appeared to be a set of 3-D vertices connected by invisible line segments, which was deforming rather than undergoing a rigid rotation.

All of the optical patterns used in Experiment 2 were subjected to the same type and magnitude of 2-D positional noise as that used in Experiment 1. All other parameters, such as number of views, angular rotation between views, viewing distance, and so forth, were identical to those used in Experiment 1.

Psychophysical task. The psychophysical task involved discriminations between spherical and planar surface patches. The observer's task was identical to that used in Experiment 1. On any given trial, an observer was asked to indicate whether a spherically curved or flat planar surface had been presented. The surfaces to be discriminated were shown in random order. The observers were again provided with auditory feedback regarding their responses.

Experimental conditions. Nine experimental conditions were formed from the orthogonal combination of *three numbers of points* (91, 19, and 9 points) and *three levels of curvature of the spherical surface* (1.33, 2.0, and 4.0 m^{-1}). These curvatures correspond to spherical surfaces with radii of 75, 50, and 25 cm, respectively.

Procedures. Each observer participated in two separate experimental sessions. Each session consisted of 50 trials for each of the nine conditions. The order of the three different optical patterns (91-point hexagonal lattice, 19-point hexagonal lattice, and 9-point cross) was determined randomly for each observer and was counterbalanced across the two sessions. For each of the three different patterns, observers progressed from high to low curvatures (from 4.0 to 2.0 to 1.33 m^{-1}). Thus, a total of 100 trials were obtained for each of the nine conditions.

Two of the 3 observers had participated in Experiment 1. The 3rd observer was also a graduate student in psychology at Vanderbilt University.

Results

The results (combined over all observers) are shown in Figure 5. The accuracy of discriminations between the spherical and planar surfaces increased linearly with increasing curvatures. Discriminations were less accurate for the 19- and 9-point patterns than for the 91-point pattern. The difference, however, was remarkably small. In fact, Wilcoxon's matched-pairs signed-ranks tests using the two sessions of the 3 observers as six independent replications showed that the 91- and 19-point patterns were not statistically different from one another ($T = 3$, $n = 6$, $p > .05$). The 19- and 9-point patterns were also not different from one another ($T = 9$, $n = 6$, $p > .05$). However, the 91- and 9-point patterns were statistically different from one another ($T = 0$, $n = 6$, $p < .05$).

Discussion

The precision of the observers' discriminations between curved and noncurved surfaces was remarkable. All observers were able to discriminate between a spherical surface patch with a 75-cm radius of curvature and a planar surface patch. Imagine what the curvature of a small area (13.8 cm^2, or $3.46°^2$ visual angle) would be on the surface of a sphere with a diameter of *a meter and a half*. The maximum difference in depth for the spherical surface patch (1.33 m^{-1} curvature, 75 cm radius, 91 points)

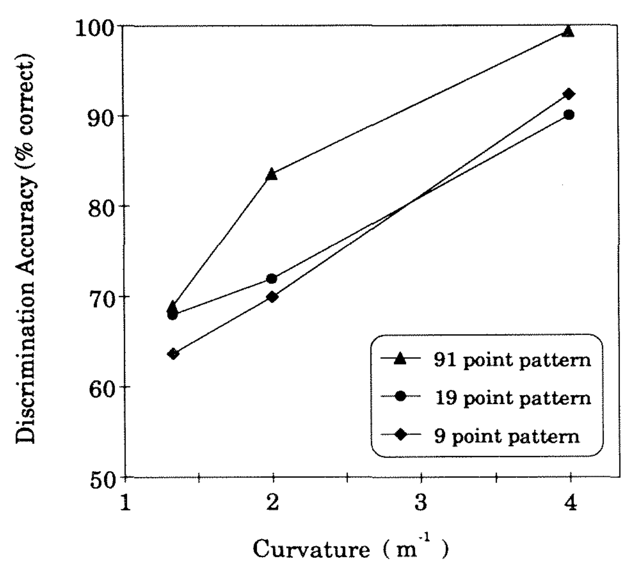

Figure 5. Results of Experiment 2 combined over all observers. This figure plots discrimination accuracy for each optical pattern in terms of percent correct as a function of increasing curvature.

between nearest and farthest points was less than 3/100 of a centimeter at a viewing distance of 114.6 cm. The maximum Weber fraction for velocity between similar points on the least curved spherical surface (75-cm-radius sphere, 91 points) and the plane was only 2.1%.

Furthermore, this level of discrimination performance was maintained when optical patterns containing only 19 and 9 points were used. It appears that few points are needed to accurately recover information about the curvature of surfaces. The performance of the observers in Experiment 2 was quantitatively very similar to that of observers in the earlier experiments by Cornilleau-Pérès and Droulez (1989). Their curved surfaces were defined by the motions of 400 points. It would appear from the present results that many fewer points are sufficient for accurate detections of curvature. Performance did improve with increasing numbers of points, but the improvement was relatively small.

EXPERIMENT 3

The results of Experiment 2 showed that curved surfaces can be discriminated from noncurved surfaces with remarkable precision. Does this result extend to apparent-motion sequences consisting of only two alternating views? Bennett, Hoffman, Nicola, and Prakash (1989), Koenderink and van Doorn (1991), Todd and Bressan (1990), Ullman (1979), and others have shown that three orthographic views are theoretically necessary in order to unambiguously recover *quantitative* information about surface curvature and other metric properties. They have shown that two orthographic views are consistent with the motions of an infinite number of possible 3-D objects.

All of the optical patterns used in the present set of experiments contained the amount of perspective that was appropriate for the observer's actual viewing position. Since these surfaces were nearly planar, were initially oriented parallel to the frontoparallel projection plane, and rotated relatively small amounts around an axis contained within the projection plane, the effects of the perspective on the optical patterns were small. The maximum perspective ratio was 1.02 for both Experiment 1 and Experiment 2. This ratio essentially constitutes a parallel projection. Longuet-Higgins and Prazdny (1980) showed that two views are not ambiguous, and that complete recovery of structure can occur if one has sufficient perspective information. Perspective ratios in psychophysical experiments are often 2.0, 3.0, or higher (Braunstein, 1962, 1977; Doner, Lappin, & Perfetto, 1984). It is therefore unlikely that the very small effects of perspective contained within the present optical patterns could be detected by human observers.

Theoretical analyses developed by Droulez and Cornilleau-Pérès (1990), Koenderink and van Doorn (1991), Lappin et al. (1991), and Todd and Bressan (1990) show that while two orthographic views are quantitatively ambiguous, they do provide qualitative information about distinctly different *types* of surfaces. In particular, discriminations between curved and noncurved surfaces should be possible. The Todd and Bressan (1990) and Koenderink and van Doorn (1991) analyses show that two orthographic views allow for the determination of an object's 3-D shape up to an affine stretching transformation. In the Todd and Bressan analysis, all 3-D objects that can be made congruent by an affine stretching transformation along an observer's line of sight are indistinguishable from the information provided by two successive views. The spherical and planar surface patches used in the previous experiments differ not only in terms of metric curvature, but also in terms of their affine structure. A spherical surface patch cannot be made planar by an affine stretching transformation along any axis. Therefore, these two 3-D shapes could be discriminated by a visual system sensitive only to affine structure.

The Droulez and Cornilleau-Pérès (1990) and Lappin et al. (1991) analyses both recover information about the presence or absence of curvature within local regions. They can also differentiate between qualitatively different types of curved regions where the surface is locally flat, cylindrical, ellipsoidal, or hyperbolic. The Lappin et al. analysis shows that points within different types of curved regions have different projected patterns of differential motion. The Droulez and Cornilleau-Pérès analysis is derived from second-order derivatives of the first-order velocity field. While their spin-variation measure is null for planar regions, and increases with increasing curvatures, it is also affected by the orientation of the surface as well as the relative motion between surface and observer. Therefore, while much qualitative information about surface curvature is available from a two-view motion sequence, that information is incomplete.

These recent analyses suggest that accurate discriminations of curved from noncurved surfaces should be possible with only two views. Earlier experiments involving discriminations of surface curvature, such as those reported by Cornilleau-Pérès and Droulez (1989), have used many more views in the apparent-motion sequence (the apparent-motion sequences used by Cornilleau-Pérès & Droulez contained 46 distinct views). It remains to be determined whether accurate discriminations of curved versus noncurved surfaces can be made using only the information provided by two alternating views. If discrimination between these two qualitatively different surfaces is possible, how does observer sensitivity compare with that obtained with longer motion sequences?

Method

Stimulus displays. The stimulus displays in this experiment were similar to those in Experiments 1 and 2. The spherical and planar surface patches were defined by 91 points arranged in a perturbed hexagonal lattice. The apparent-motion sequence consisted of two views of the 3-D surface separated by a 15° angular rotation. The duration of each view was 250 msec. Each trial lasted for 2.5 sec and consisted of five repetitions of the two-view motion sequence.

All other stimulus parameters were identical to those used previously in Experiments 1 and 2.

Psychophysical task. The psychophysical task was identical to that used in Experiment 2. The observers were required to discriminate between spherical and planar surface patches. The spherical and planar surfaces were presented randomly within a block of trials. On any given trial, observers indicated, by pressing either of two keys on a computer keyboard, which of the two surfaces had been presented. Feedback was provided if the observer's response was correct.

Experimental conditions. Three conditions evaluated the discriminability of differently curved spherical surface patches (the curvatures were 1.67, 2.5, and 5.0 m^{-1}, corresponding to radii of curvatures of 60, 40, and 20 cm, respectively) versus noncurved planar patches.

Procedures. Each observer participated in two separate experimental sessions. Each session consisted of 50 trials for each of the three conditions. The order of the conditions progressed from high to low curvatures (from 5.0 to 2.5 to 1.67 m^{-1}). A total of 100 trials were thus obtained for each of the three conditions.

The 3 observers were the same as those in Experiment 2.

Results

The results for the individual observers are shown in Figure 6. Results for the 3 observers were unusually consistent. Discrimination accuracies between the spherical and planar surface patches increased from near-threshold levels for curvatures of 1.67 m^{-1} to nearly perfect performance for curvatures of 5.0 m^{-1}.

Discussion

Similar to the results of previous experiments, the precision of the observers' discriminations between differently curved surfaces was remarkable. The maximum difference in displacement between corresponding points on the spherical and planar surface patches was 47.1″ of arc for the 5.0-m^{-1} curvature condition, 23.5″ of arc for the 2.5-m^{-1} curvature, and only 15.7″ of arc for the lowest curvature condition, 1.67 m^{-1}. All 3 observers were able to discriminate between the spherical and planar surfaces with the lowest curvatures at a 70% accuracy level.

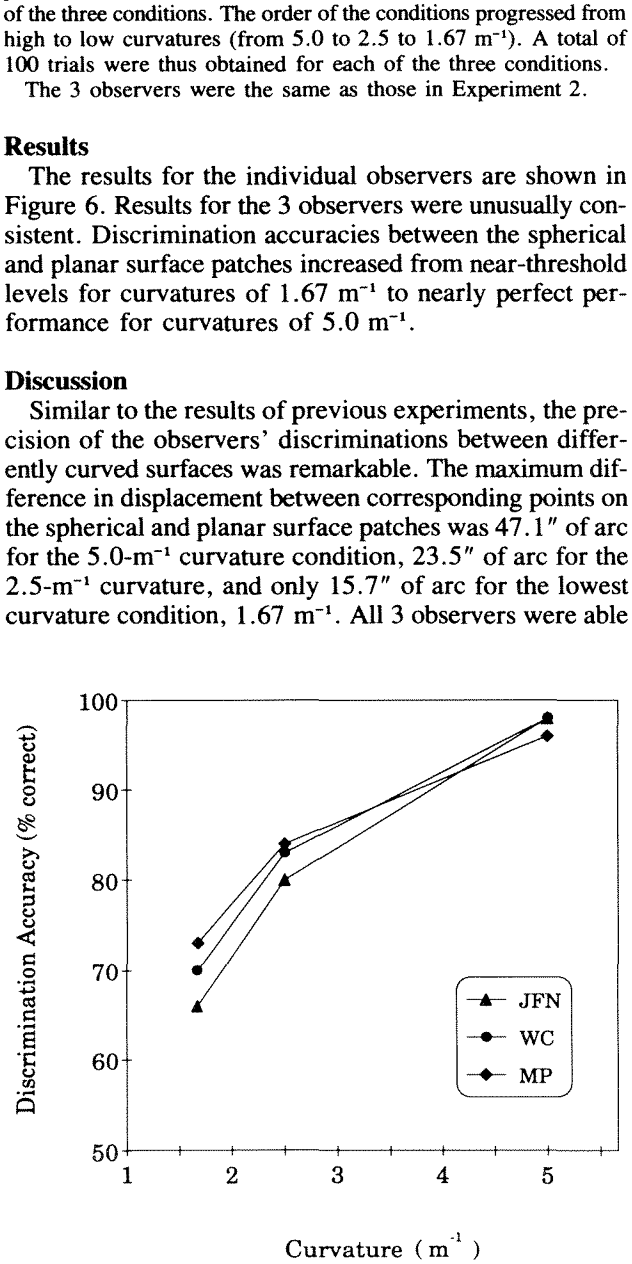

Figure 6. Results of Experiment 3 plotted separately for each observer. This figure plots each observer's discrimination accuracy (percent correct) as a function of increasing curvature.

Nakayama (1981) presented observers with random-dot patterns in which both common and differential image motion occurred. Upper and lower half-fields of his patterns moved with similar velocities, but in opposite directions. Common motion was also added as an independent vector component. In conditions similar to those in the present experiment, with a 2′ of arc common-motion component, he found thresholds for detection of differential motion *within a single pattern* to be around 5″ of arc. Performance for the current experiment involved detection of differential motion *across different trials* on the order of 16″ of arc. These results show that discrimination of curved versus noncurved surfaces can be executed with remarkable precision, given only two alternating views, and is almost as accurate as the detection of differential 2-D motion.

Evidence for the determination of exactly which aspects and relationships within the different optical patterns enabled the precise discrimination between these two different 3-D shapes, however, is insufficient. The Koenderink and van Doorn (1991) and Todd and Bressan (1990) analyses can recover affine but not euclidean structure from two views. The Droulez and Cornilleau-Pérès (1990) spin-variation measure recovers information about surface curvature from the velocity field (i.e., two views), but the measure is affected by other variables as well as surface properties. Longuet-Higgins and Prazdny (1980) showed that one could determine euclidean structure from the optic-flow field (i.e., two views) if one had a sufficiently strong perspective and could take first and second derivatives over the flow field.

Currently, little information exists to evaluate whether human observers can take advantage of the information provided by strong perspective to recover euclidean structure from two views as Longuet-Higgins and Prazdny have suggested. It is impossible to recover euclidean structure from two orthographic views. Since our experimental displays contained very small amounts (maximum perspective ratio in this experiment was 1.009) of perspective, it seems very unlikely that this perspective information was visible, although this possibility cannot be definitely ruled out. The high degree of correlation between the results of Experiment 1 and the spin-variation measures of Droulez and Cornilleau-Pérès suggests that their analysis is relevant to the performance of our observers.

CONCLUSIONS

The preceding set of experiments have shown that surface curvatures portrayed by the kinetic depth effect can be detected with remarkable accuracy. The small differences in curvature that were discriminated in these experiments involved very small differences in the velocity fields of the 2-D optical images. Indeed, the implicit discriminations in velocity implied by these curvature discriminations appear to be at the limits of the visual abilities documented in experiments specifically concerned with the detectability of spatial displacement and velocity.

Extensive velocity variations within optical patterns by themselves were not always sufficient to produce reliable percepts of rotating 3-D surfaces. The noncurved planar surfaces used in these experiments were often found to be perceptually degenerate. These planar surfaces were not always perceived as completely rigid objects rotating in 3-D space. This ambiguity is remarkable, because well-defined velocities existed for all of the points (91 in Experiments 1 and 2) within the optical patterns across an entire 15-view apparent-motion sequence. Variations in velocity within optical patterns have long been thought to indicate variations in depth (Helmholtz, 1925). Despite well-defined optical patterns with internal variations in velocity covarying with the depths of the planes' constituent points, these optical patterns were not always perceived as coherently rotating 3-D surfaces. Many of the computational models that recover 3-D structure from motion would be able to accurately detect rigid rotations over a 15-view motion sequence of a planar arrangement of 91 points. This perceptual ambiguity did not exist for surfaces that were curved in a direction orthogonal to the direction of rotation.

In line with the stereoscopic results of Uttal and his colleagues (Uttal, 1987; Uttal, Davis, Welke, & Kakarala, 1988) the discriminability of differently curved surfaces defined by motion did not require dense optical patterns containing many points. Accurate discriminations of curved from noncurved surfaces (Experiment 2) were possible even with optical patterns containing as few as 9 points. However, the placement of these few points was critical. With purely random 2-D positioning of the 9 points, observers' percepts were of a bending, jointed space curve where the points appeared to be connected by "invisible" line segments. The configuration of 9 points used in Experiment 2 was a cross composed of two intersecting perpendicular lines, each containing 5 points. This configuration reliably elicited the perception of a rotating surface. Presumably, the collinear arrangements of the points enabled the determination of curvature. Three points, at least approximately collinear, would be the minimum number necessary for a determination of curvature in any given direction. Two collinear lines in different directions would be necessary to unambiguously specify a surface.

The present experiments have important implications for some of the existing models which recover various aspects of 3-D structure from optical motion. The results of Experiment 3 showed that two views were sufficient to very accurately discriminate between curved and noncurved surfaces. All observers reported that the appearance of the curved surfaces was compelling. The surfaces looked curved. The finding that two views were sufficient to perceive the curvature of rotating surfaces conflicts with some theoretical models which require three views to recover the 3-D structure of a moving object. Usually, these models recover the depths and/or orientations of all visible feature points on an object's surface. To perceive the curvature of a surface based upon such a representation would require the detection of changes in depth or orientation across the layout of the depth or orientation map. If one cannot obtain a depth or orientation map from a two-view motion sequence, and observers can perceive curved surfaces defined by the motion between two views, the perception of curved surfaces must be based on other properties besides depths and orientations.

One of the most well-known algorithms that recover 3-D structure from motion was developed by Ullman (1979). Ullman's structure-from-motion theorem requires at least three distinct views (containing the projections of four noncoplanar points) in order to accurately recover 3-D structure. Human perception of 3-D shape from motion does not seem to require establishment of correspondences beyond two adjacent views. Other models that require more than two views include those of Webb and Aggarwal (1981) and Hoffman and Bennett (1986). These models also assume a fixed axis of rotation. This additional assumption enables these models to recover aspects of 3-D structure that Ullman's algorithm cannot. Given three (or more) views, the models of Webb and Aggarwall (1981) and Hoffman and Bennett (1986) can recover the 3-D structure of planar arrangements of points. The planar surfaces used in Experiments 1 and 2 rotated around a fixed axis for more than three views. According to these models, the noncurved planar surfaces should not have been perceptually ambiguous—the information necessary to recover their structure was available. The fact that the observers often perceived the noncurved surfaces as nonrigid deforming 3-D structures, rather than rigid planar configurations, would seem to indicate that the human visual system cannot take advantage of all of the information present within longer motion sequences.

To summarize, these psychophysical experiments have served to refine existing knowledge about the perception of 3-D shape from motion. They clarify the conditions under which accurate detections of curvature are possible. They have revealed that the human visual system is extremely sensitive to the presence or absence of curvature, comparable to the hyperacuity found under optimal circumstances.

This sensitivity to small differences in curvature is reduced when surfaces are curved only along the direction of rotation. The poor performance of observers in tasks requiring them to discriminate between 3-D shapes with differing curvatures along the direction of rotation cannot be explained in terms of the shapes' depth, orientation, or simple projective velocity differences.

While demonstrating the precision and robustness (i.e., few points, two views) of human observers' detections of surface curvature, these experiments also demonstrate that a particular property of the first-order optic-flow field, namely the spin variation developed by Droulez and Cornilleau-Pérès (1990), correlates highly with the responses of human observers performing discriminations of curvature. Further experimentation will be necessary to refine our knowledge of how the human visual system acquires information about the 3-D shapes of environmental surfaces and objects.

REFERENCES

BENNETT, B. M., HOFFMAN, D. D., NICOLA, J. E., & PRAKASH, C. (1989). Structure from two orthographic views of rigid motion. *Journal of the Optical Society of America A*, **6**, 1052-1069.

BRAUNSTEIN, M. L. (1962). Depth perception in rotating dot patterns: Effects of numerosity and perspective. *Journal of Experimental Psychology*, **64**, 415-420.

BRAUNSTEIN, M. L. (1977). Perceived direction of rotation of simulated three-dimensional patterns. *Perception & Psychophysics*, **21**, 553-557.

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.

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

DONER, J., LAPPIN, J. S., & PERFETTO, G. (1984). Detection of three-dimensional structure in moving optical patterns. *Journal of Experimental Psychology: Human Perception & Performance*, **10**, 1-11.

DROULEZ, J., & CORNILLEAU-PÉRÈS, V. (1990). Visual perception of surface curvature: The spin variation and its physiological implications. *Biological Cybernetics*, **62**, 211-224.

HELMHOLTZ, H. L. F. VON (1925). *A treatise on physiological optics* (Vol. 3, 3rd ed., J. P. C. Southall, Trans.). New York: Optical Society of America.

HILBERT, D., & COHN-VOSSEN, S. (1952). *Geometry and the imagination*. New York: Chelsea.

HOFFMAN, D. D., & BENNETT, B. M. (1985). Inferring the relative three-dimensional positions of two moving points. *Journal of the Optical Society of America A*, **2**, 350-353.

HOFFMAN, D. D., & BENNETT, B. M. (1986). The computation of structure from fixed-axis motion: Rigid structures. *Biological Cybernetics*, **54**, 71-83.

HOWELL, D. C. (1982). *Statistical methods for psychology*. Boston: Duxbury Press.

KOENDERINK, J. J. (1990b). Some theoretical aspects of optic flow. In R. Warren & A. H. Wertheim (Eds.), *The perception and control of self-motion* (pp. 53-68). Hillsdale, NJ: Erlbaum.

KOENDERINK, J. J., & VAN DOORN, A. J. (1991). Affine structure from motion. *Journal of the Optical Society of America A*, **8**, 377-385.

LAPPIN, J. S., CRAFT, W. D., & TSCHANTZ, S. T. (1991). Perceiving shape from motion. *Investigative Ophthalmology & Visual Science* (Suppl.), **32**, 1277.

LONGUET-HIGGINS, H. C., & PRAZDNY, K. (1980). The interpretation of a moving retinal image. *Proceedings of the Royal Society of London B*, **208**, 385-397.

LUCE, R. D. (1963). Detection and recognition. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 1, pp. 103-189). New York: Wiley.

NAKAYAMA, K. (1981). Differential motion hyperacuity under conditions of common image motion. *Vision Research*, **21**, 1475-1482.

NORMAN, J. F., LAPPIN, J. S., & ZUCKER, S. W. (in press). The discriminability of smooth stereoscopic surfaces. *Perception*.

ROGERS, B., & CAGENELLO, R. (1989). Disparity curvature and the perception of three-dimensional surfaces. *Nature*, **339**, 135-137.

ROGERS, B. J., & GRAHAM, M. E. (1983). Anisotropies in the perception of three-dimensional surfaces. *Science*, **221**, 1409-1411.

SANDER, P. T., & ZUCKER, S. W. (1990). Inferring surface trace and differential structure from 3-D images. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, **12**, 833-854.

SPERLING, G., LANDY, M. S., DOSHER, B. A., & PERKINS, M. E. (1989). Kinetic depth effect and identification of shape. *Journal of Experimental Psychology: Human Perception & Performance*, **15**, 826-840.

STEVENS, K. A., & BROOKES, A. (1988). Integrating stereopsis with monocular interpretations of planar surfaces. *Vision Research*, **28**, 371-386.

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

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

UTTAL, W. R. (1987). *The perception of dotted forms*. Hillsdale, NJ: Erlbaum.

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.

WEBB, J. A., & AGGARWAL, J. K. (1981). Visually interpreting the motion of objects in space. *Computer*, **14(8)**, 40-46.

WEINSHALL, D. (1991). Direct computation of qualitative 3-D shape and motion invariants. *IEEE Pattern & Machine Intelligence*, **13**, 1236-1240.

ZUCKER, S. W. (1986). The computational connection in vision: Early orientation selection. *Behavior Research Methods, Instruments, & Computers*, **18**, 608-617.

NOTE

1. The nonparametric statistics presented in this manuscript were calculated following the methods contained in Howell (1982).

(Manuscript received June 10, 1991;
revision accepted for publication December 3, 1991.)