

The perception of surface orientation from multiple sources of optical information

J. FARLEY NORMAN, JAMES T. TODD, and FLIP PHILLIPS
Ohio State University, Columbus, Ohio

An orientation matching task was used to evaluate observers' sensitivity to local surface orientation at designated probe points on randomly shaped 3-D objects that were optically defined by texture, lambertian shading, or specular highlights. These surfaces could be stationary or in motion, and they could be viewed either monocularly or stereoscopically, in all possible combinations. It was found that the deformations of shading and/or highlights (either over time or between the two eyes' views) produced levels of performance similar to those obtained for the optical deformations of textured surfaces. These findings suggest that the human visual system utilizes a much richer array of optical information to support its perception of shape than is typically appreciated.

One of the fundamental issues in the study of human perception concerns how the shapes of objects in the environment are visually specified from the measurable properties of optical stimulation. There are many different aspects of optical structure that are known to provide perceptually salient information about an object's three-dimensional form. Some of these properties—the so-called *pictorial depth cues*—are available within individual static images. These include texture gradients, linear perspective, and patterns of shading. Others are defined by the systematic transformations among a sequence of multiple images, and include the disparity between each eye's view in binocular vision, and the optical deformations that occur when objects are observed in motion.

In the theoretical analysis of motion or binocular disparity, two distinct classes of optical phenomena need to be considered. One involves the optical transformations of identifiable image features, such as surface texture or the vertices of a polyhedron, for which it is possible to establish a point-to-point correspondence over multiple views. The ability to match corresponding features in different images is a necessary condition for most existing computational models for the analysis of 3-D shape from motion or stereo, but there are other types of optical transformations that occur frequently in natural vision, for which this condition cannot be satisfied. These include the optical deformations of occlusion contours and smooth gradients of image shading.

Patterns of shading in an image arise because of systematic changes in local surface orientation. Patches that are oriented perpendicularly to the prevailing direction of illumination reflect the greatest amount of light, while those that are parallel to the direction of illumination re-

flect the least. For matte surfaces that scatter light in all directions, these reflections are governed by Lambert's law, which states that the amount of reflected light is proportional to the cosine of the angle between the surface normal and the direction of illumination. Such surfaces are sometimes referred to in the literature as *lambertian*. For shiny surfaces, the light reflected toward the point of observation is also influenced by the position of the observer, which produces the appearance of specular highlights.

To better appreciate the structure of image shading, it is useful to consider a set of points on a surface that all have the same luminance. For a smoothly curved object with homogeneous reflectance, these points will be aligned along continuous space curves, which we shall refer to as *isophotes* (Koenderink & van Doorn, 1980). The overall pattern of shading at a point of observation is determined by the optical projections of these isophotes, which connect points of equal image intensity.

If an object moves or is viewed from different vantage points, this pattern of image shading will be systematically deformed, but the specific nature of this deformation depends on a number of different factors. Consider, for example, an observer's movements relative to a fixed lambertian surface with a fixed pattern of illumination. Because of the fixed relation between the object and the light source, its isophotes will remain constant, as if they were painted on the surface. The patterns of optical motion in that case are identical to those of identifiable texture elements, and could therefore be analyzed using traditional computational models of structure from motion (e.g., see Horn & Schunck, 1981; Nagel, 1981, 1987). These models are inappropriate, however, when an observer moves relative to a shiny object, or when an object moves relative to its light source. Under these conditions, the isophotes will slide over the object's surface, producing a very different type of optical deformation in its pattern of image shading. The behavior of specular highlights in these situations is particularly complex.

This research was supported by Air Force Office of Scientific Research Grant F49620-93-1-0116 to J.T.T. Correspondence should be addressed to J. F. Norman, Department of Psychology, Ohio State University, 142 Townshend Hall, Columbus, OH, 43210-1222 (e-mail: fnorman@magnus.acs.ohio-state.edu).

They tend to move very rapidly over relatively flat regions of a surface and to cling more stably in regions of high curvature.

During the past decade, there has been a growing amount of theoretical research on the computational analysis of 3-D shape from static patterns of image shading (see Horn & Brooks, 1989, for a review), and there have been numerous psychophysical investigations of how this information is perceptually analyzed by actual human observers (e.g., see Erens, Kappers & Koenderink, 1993a, 1993b; Johnston & Passmore, 1994a, 1994b; Mingolla & Todd, 1986; Todd & Mingolla, 1983; Todd & Reichel, 1989). In general, the available empirical evidence indicates that static shading is primarily used to determine qualitative aspects of surface structure, such as the presence of hills or valleys, but that it is not a particularly powerful source of information for the precise specification of metrical properties, such as relative depths or orientations. There is also evidence to indicate that shading *per se* may only be informative when it occurs in the presence of well-defined occlusion contours (see Ikeuchi & Horn, 1981; Todd & Reichel, 1989).

In comparison with the large number of studies concerned with static shading, there has been relatively little research on the perceptual analysis of how shading deforms as an object is viewed from multiple vantage points. There have been a few demonstrations that observers can achieve compelling kinetic depth effects or stereopsis from shaded objects with no identifiable features (Bülthoff & Mallot, 1988; Koenderink, Kappers, Todd, Norman, & Phillips, *in press*; Todd, 1985), and that stereoscopically viewed highlights can bias the perception of textured surfaces with ambiguous relief (Blake and Bülthoff, 1990, 1991). However, there are few data to indicate either the way in which dynamic or stereoscopic shading compares with other sources of information for the precise specification of an object's 3-D form or the extent to which the perceptual analyses of

these different types of information might potentially interact with one another.

The research described in the present article was designed to examine these issues for observers' judgments of local surface orientation. This particular task was chosen because surface orientation is often considered as a primitive geometric property for the perceptual representation of 3-D form (e.g., see Marr, 1982), and it is by far the most common representation for the computational analysis of image shading (e.g., see Horn & Brooks, 1989). Observers' judgments were obtained at designated probe points on randomly shaped surfaces with various combinations of texture, lambertian shading, and specular highlights, either moving or stationary, and viewed either monocularly or stereoscopically.

METHOD

Apparatus

The optical patterns were created and displayed on a Silicon Graphics Crimson VGXT workstation with hardware texture-mapping capabilities. For binocular patterns, stereoscopic viewing hardware was used. The stereoscopic half-images were presented using liquid-crystal-display (LCD) shuttered glasses that were synchronized with the monitor's refresh rate. The left and right views of a stereo pair were displayed at the same position on the monitor screen, but they were temporally offset. The left and right lenses of the LCD glasses shuttered synchronously with the display, so that each view of the stereo pair was seen only by the appropriate eye. The CRT was refreshed at 120 Hz; thus, each view of a stereoscopic half-image was updated at half that value (i.e., at 60 Hz). The viewing distance was 76 cm, such that the $1,280 \times 1,024$ -pixel ($w \times h$) display screen subtended $25.2^\circ \times 20.3^\circ$ visual angle.

Stimulus Displays

The stimuli in this study were designed to simulate the optical projections of globally convex smoothly curved surfaces that resembled real-world objects, such as water-worn pebbles or potatoes (see Figure 1 for a representative example). A set of twenty such objects was generated at random by distorting spheres with an initial radius of 8 cm. This transformation was accomplished by




Figure 1. A stereogram of a textured object similar to those used in the present experiment.

adding a series of sinusoidal perturbations on the surface at random orientations. The resulting objects were smoothly curved with no discontinuities, and by keeping track of each successive sinusoidal perturbation, we were able to obtain an analytically defined surface normal at each point (see also Todd & Norman, in press; Koenderink et al., in press).

The objects were presented with shading and/or texture to simulate different types of surface materials: Those depicted with shading alone were defined by 5,120 triangular polygons, while the textured objects were composed of 1,280 polygons. In all cases, the shading and texture were hardware interpolated within the interior of each triangular polygon, so that the depicted surfaces appeared smoothly curved. Surface shading was simulated using a standard computer graphics reflectance model (see Todd & Mingolla, 1983), in which the shading is partitioned into three components: an ambient component that is constant for all surface orientations, a diffuse (Lambertian) component that varies with the cosine of the angle between the surface normal and the direction of illumination, and a specular component that varies as a function of the surface normal, the direction of illumination, and the direction of view. The simulations employed a single-point light source oriented with a slant of 28° and a 45° tilt up and to the left of the observers' line of sight. Texturing was achieved using a 2-D pattern that was designed to resemble granite. Each polygon was first

rotated to a frontoparallel orientation and then mapped onto a random region of the 2-D texture pattern, which ensured that equal areas of the surface contained equal amounts of texture (see Todd & Mingolla, 1984).

There were five different surface-type conditions: (1) a textured surface resembling red granite, whose shading was purely ambient; (2) a smoothly shaded (Lambertian) blue surface, resembling plastic, with a 30% ambient component and a 70% diffuse component; (3) a dark, shiny surface whose shading was purely specular, and which resembled polished obsidian; (4) a textured and shaded granite surface with a 30% ambient component and a 70% diffuse component; and (5) a shiny blue surface with shading and highlights, which had a 30% blue ambient component, a 40% blue diffuse component, and a 30% white specular component. For all surfaces with specular highlights, the exponent of the specular component was 20 (see Todd & Mingolla, 1983, or any computer graphics text for a more complete description of this standard model for the computation of image shading). Figure 2 shows a number of representative examples (in monochrome) of these different surface types.

The objects were presented in four possible viewing conditions: (1) monocular, static; (2) monocular, motion; (3) stereoscopic, static; and (4) stereoscopic, motion. All of the displays were generated with the appropriate perspective for a 76-cm viewing dis-

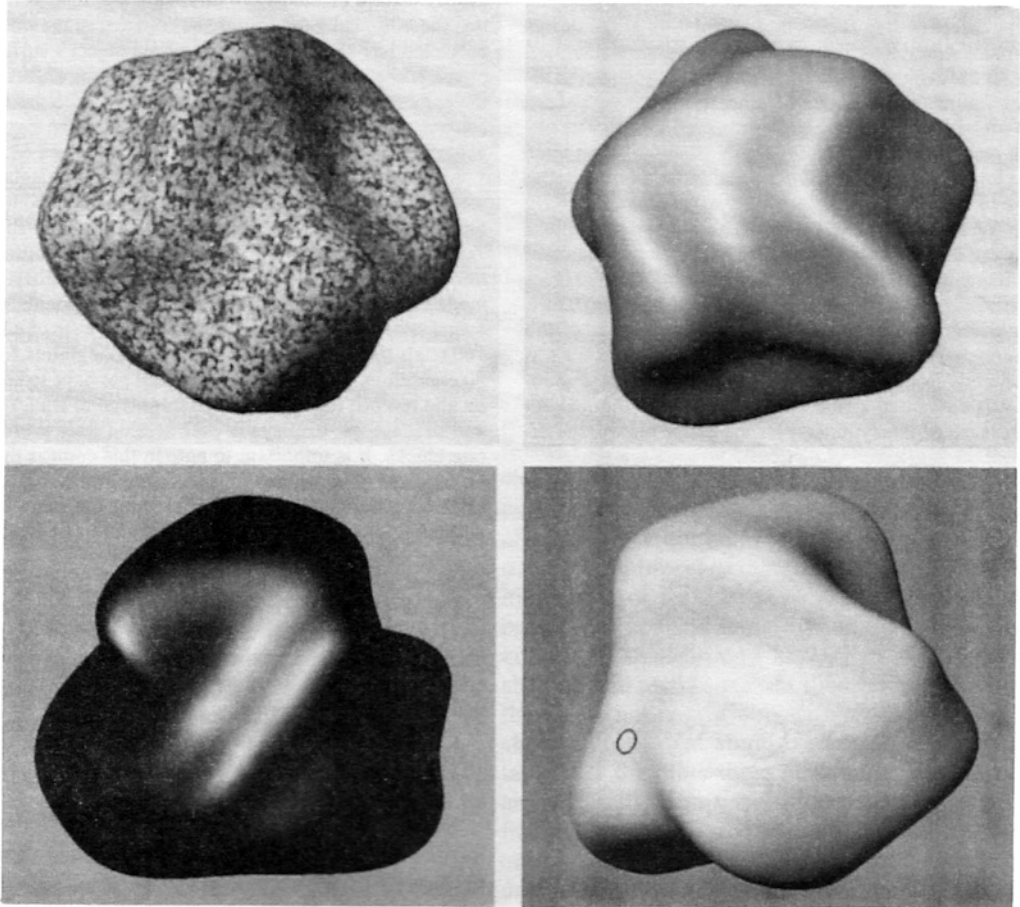


Figure 2. Examples of four surface types used in the present experiment. Counterclockwise from the upper left, the depicted surfaces have (1) texture, (2) pure specular highlights, (3) pure Lambertian shading, and (4) combined Lambertian shading plus highlights. In the actual experimental displays, the Lambertian shading was colored blue. The object depicted in the lower right also shows an elliptical gauge figure similar to those that the subjects were required to adjust.

tance. For the stereoscopic displays, each eye's view was computed on the basis of an interpupillary distance of 6.1 cm. For patterns displayed with motion, the objects oscillated in depth about a vertical axis between -12° and $+12^\circ$ from their home position, with a 2° angular displacement at each frame transition. Thus, each apparent motion sequence was composed of a total of 13 individual frames, which were updated at a rate of 20 Hz.

The perception of the local surface orientation was evaluated at randomly selected probe points on the depicted objects. In this paper, we will describe the orientation of a surface region in terms of its *slant* and *tilt*. Slant refers to the angle of the surface normal relative to the line of sight, and can therefore range from 0° to 90° . Tilt specifies the direction of the surface depth gradient in the image, and ranges from 0° to 360° . To better exemplify these concepts, the optical projections of a 3-D circle with varying amounts of slant and tilt are shown in Figure 3. In our experiment, the values of slant at the designated probe points were restricted to four possible values of 25° , 35° , 45° , and 55° . The tilts were chosen at random on each trial over the full range of 0° – 360° .

Procedure

We used a psychophysical adjustment procedure similar to that developed by Koenderink, van Doorn, and Kappers (1992). The technique is simplest to describe for objects presented in the monocular static conditions. On each trial, a random probe point was selected on one of the 20 possible objects. The object was then rotated appropriately so that the surface orientation at the designated probe point would have one of the four possible slants (i.e., 25° , 35° , 45° , or 55°), and a tilt that was selected at random. An elliptical gauge figure similar to those shown in Figure 3 was centered on the probe point in the object's projected image. The observers' task on each trial was to adjust the shape of the gauge figure in the image plane so that it appeared to be a circle in the tangent plane of the surface at the depicted probe point.

The gauge figure was adjusted using a hand-held mouse, whose movements were defined in a polar-coordinate system, such that radial movements altered the elliptical eccentricity of the figure (i.e., the ratio of its major and minor axes) over a range from 0 to 1, and circular movements changed the image orientation of the major and minor axes over a range of 360° . The adjusted tilt was determined directly from the orientation of the minor axis. The adjusted slant was computed as the arc cosine of the elliptical eccentricity. Because of the symmetry of the gauge figure, the di-




Figure 3. A schematic illustration indicating the two components of surface orientation: slant and tilt. Slant refers to the orientation of the figure in depth with respect to the fronto-parallel plane. The tilt component indicates the direction in the image of the surface depth gradient.

rection of adjusted tilt was ambiguous up to a reflection. To eliminate this ambiguity, a set of four circular dials was presented in the corners of the display screen, which pointed in the direction of adjusted tilt. The observers were instructed to examine these dials before recording their responses to ensure that the dials correctly matched the perceived direction of the surface depth gradient. All of the observers reported that with sufficient practice, they could quickly adjust the elliptical gauge figure in the image plane so that it appeared as a circle in three-dimensional space that was attached to the surface of the depicted object. When they were satisfied that the adjustment was correct, they were instructed to press a button on the mouse to proceed to the next trial.

In the stereo conditions, the gauge figure was only presented to the right eye, to eliminate any stereoscopic information from the figure itself. This prevented the observers from potentially performing the task by comparing the disparities of the gauge figure with those of nearby texture elements on the depicted object. All four of the observers who participated in the experiment perceived the gauge figures in the stereoscopic conditions as attached to the surface. In our pilot investigations, however, one other observer had to be excluded because the figures often appeared to him as though they were floating above the textured surfaces.

A similar control was also incorporated in the motion conditions to prevent observers from making their adjustments by minimizing the relative-motion parallax between the moving gauge figure and nearby texture elements. To eliminate any relative-motion parallax, the elliptical gauge figure as adjusted in the image was back projected into the tangent plane of the object's surface at the depicted probe point and rotated rigidly with the object in depth over the apparent motion sequence. The task can be thought of in this context as one of adjusting the shape of a gauge figure in the tangent plane until it appears circular.

All possible combinations of these different display parameters were used, for a total of 80 distinct experimental conditions (5 surface types \times 4 viewing conditions \times 4 possible surface slants). Ten different probe points were selected at random for each condition over the course of the entire experiment, so that each observer made a total of 800 adjustments. These adjustments were recorded over a series of five experimental sessions, each of which involved 160 trials that included two different probe points for each of the 80 conditions. So that the subjects did not have to continually put on and remove the LCD glasses from trial to trial, the monocular and stereoscopic displays in each session were presented in separate blocks. It is important to note in this context that there were no repeated observations for individual probe points, which were selected at random on every trial from the large population of polygon vertices on the 20 possible stimulus objects.

Observers

The displays were presented to four observers, two of whom were two of the authors (JFN and JTT), while the others were experienced psychophysical observers who were naive to the specific details of this particular experiment. All observers had normal or corrected-to-normal vision.

RESULTS

Several different aspects of the results deserve to be highlighted. Let us first consider how the observers' absolute errors in orientation varied across the different experimental conditions. By errors in orientation, we mean the angular difference in degrees between the simulated surface orientation at a depicted probe point and the observers' adjusted orientation. If these orientations are represented as unit surface normals, the absolute error can be calculated as the inverse cosine of their vector dot

product. From these individual measures, an average error was obtained for the 10 adjustments in each of the 80 experimental conditions, which were then combined in various ways to illustrate important differences.

Figure 4 shows the average error for the four different viewing conditions collapsed across slant and surface type for each of the four individual observers. One can see that if either motion or stereo was present, the average error in orientation was approximately 14.5°. However, if neither motion nor stereo was included in the displays, the magnitude of this error increased by 69% to approximately 24.5°. Each observer's adjustments were subjected to a repeated measures analysis of variance. The deficiency of the monocular, static conditions was confirmed by significant two-way interactions between stereo and motion for each observer [JFN: $F(1,39) = 52.1$; JTT: $F(1,39) = 9.0$; VJP: $F(1,39) = 18.5$; DTL: $F(1,39) = 32.5$, all significant with $p < .01$].

Figure 5 provides a more detailed breakdown of this effect for the five different surface types. One can readily see that a considerable improvement in performance occurred for each of the five surface types as stereo and/or motion was added to the optical patterns. Consider, for example, the extent of this improvement in the texture-only displays. When static monocular texture gradients were the only available sources of information, the average absolute error was approximately 25°, but this was reduced to only 12° when the displays were presented in stereo or in motion. This large improvement of 53% was presumably due to the projected motions or binocular disparities of identifiable feature points, which have been well established as perceptually salient sources of information about the 3-D structure of objects in space. What is theoretically important about the present results, however, is that there were similar improvements in performance with the addition of motion or stereo even for those displays that contained no identifiable feature points.




Figure 4. The average errors in adjusted orientation for each individual observer for the four combinations of motion and stereo collapsed over the different surface types. The motion/stereo conditions are labeled as follows: with stereo (S), without stereo (nS), with motion (M), and without motion (nM).




Figure 5. The average errors in adjusted orientation for the five different surface types, depicted with or without stereo or motion. The different surface types are labeled as follows: lambertian shading (S), texture, with no shading (T), specular highlight shading (H), lambertian shading combined with texture (S+T), and lambertian shading combined with specular highlights (S+H).

For the displays composed of pure lambertian shading, pure specular highlights, or shading and highlights in combination, the addition of motion and/or stereo produced improvements of 37%, 31%, and 36%, respectively. These findings provide strong evidence that the optical deformations of these shading and highlight fields, perhaps in conjunction with the deforming boundary contour (see Norman & Todd, 1994), can provide useful information for the perceptual analysis of 3-D form. This is also supported by the phenomenological impressions of the observers, all of whom reported that the moving or stereoscopic displays with shading or highlights appeared just as compelling as those with texture.

To test for the presence of any systematic anisotropies in the observers' judgments of local orientation, we performed an additional analysis to compare errors in their slant and tilt components. Figure 6 shows a scatterplot of adjusted tilt versus simulated tilt for Observer J.F.N. in all of the different viewing conditions in which individual surface properties (i.e., lambertian shading, specular highlights, or texture) were presented in isolation. To preserve space and legibility, the pattern of results for combined shading and highlights and for combined shading and texture are not illustrated in the figure, but they were not appreciably different from those that are shown. These data are also representative of those obtained for the other three observers. Note that the plots are essentially linear, with slopes near 1.0, especially for the nine conditions that utilized stereo and/or motion. In an effort to quantify these general observations, an analysis of linear regression was performed on the 40 different adjustments (10 probe points at four distinct slants) for each observer with each of the 20 combinations of surface type and viewing condition. The correlation coefficient (r^2) values computed from this analysis are shown in Table 1. As is evident from the table, the simulated tilt accounted for about 95% of the total variance

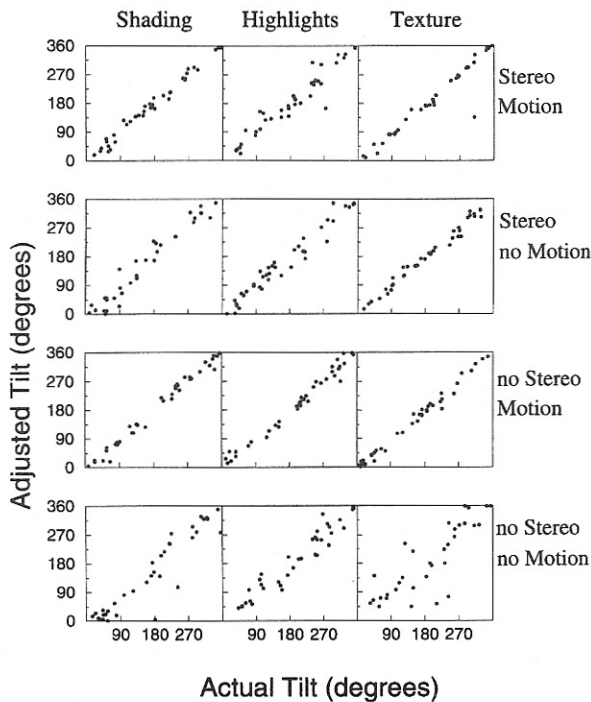


Figure 6. Scatterplots of the tilt component of adjusted surface orientation for Observer J.F.N. as a function of simulated tilt.

in the observers' adjustments when the displays were presented in stereo or in motion, although this dropped somewhat, to around 85%, in the monocular static conditions.

In contrast to their relatively accurate adjustments of surface tilt, the observers' judgments of slant were much less precise. Figure 7 shows scatterplots of the slant data for Observer J.F.N. in the same 12 conditions as are shown in Figure 6, together with their best-fitting regression lines. These results are again representative of the other conditions and observers. It is clear in this figure that there was only a moderate correlation between the simulated and adjusted slants. While noisy, these data are similar to the tilt data, in that the relationship between simulated and adjusted slant is much stronger in conditions with either stereo or motion than it is in the monocular static conditions. The individual r^2 values for all observers in all conditions are shown in Table 2, in which it can be seen that the r^2 values are mostly clustered around 50%, indicating that the simulated slants accounted for only about half of the total variance in the observers' judgments.

DISCUSSION

The results of this experiment clearly demonstrate that the deformations of shading and highlights are informative sources of optical information that are used by the human visual system to support its perception of shape. The key finding to support this conclusion is that motion and stereo improved performance in the observers' judgments of local orientation even when the

displays contained no identifiable feature points, so that the only available information was in the optical deformations (over time or in binocularly disparate views) in the occlusion contours of an object and its smooth pattern of image shading. Indeed, the overall level of performance for shaded objects in motion or stereo was only slightly worse than the performance obtained with textured objects (see Figure 5). It is important to keep in mind that most existing theoretical models for the computational analysis of shape from motion or stereo are only applicable to the projected displacements of identifiable feature points (see, however, Blake & Bülthoff, 1990, 1991; Cipolla & Blake, 1990; Giblin & Weiss, 1987; Koenderink & van Doorn, 1980; and Pentland, 1991, for some notable exceptions). If the optical deformations of smooth occlusion contours or image shading were analyzed in this manner, there would in general be no mathematically possible rigid interpretation. Thus, the results of the present experiment and the related findings of Bülthoff and Mallot (1988), Koenderink et al. (in press), Norman and Todd (1994) and Todd (1985) provide strong evidence that the basic concepts of optical "motion" and binocular "disparity" should be viewed in a more general and less restrictive manner.

The current finding that surface tilt is perceived more accurately than slant replicates similar findings by Koenderink and his colleagues using the same gauge figure adjustment task for a variety of different stimuli, including real objects and both monocular and stereoscopic

Table 1
Correlation Coefficient (r^2) Values From the Linear Regression Analysis of the Tilt Adjustments for all Four Observers

Tilt Adjustment	Observer			
	J.F.N.	J.T.T.	V.J.P.	D.T.L.
Stereo				
Motion				
Shading	.983	.945	.956	.851
Texture	.926	.966	.973	.985
Highlights	.916	.978	.925	.862
Shading & texture	.984	.979	.982	.979
Shading & highlights	.983	.966	.973	.973
No motion				
Shading	.940	.963	.874	.907
Texture	.983	.959	.977	.983
Highlights	.956	.930	.880	.892
Shading & texture	.980	.988	.965	.993
Shading & highlights	.983	.956	.825	.959
No Stereo				
Motion				
Shading	.983	.977	.905	.903
Texture	.987	.973	.981	.964
Highlights	.934	.981	.971	.739
Shading & texture	.971	.980	.973	.977
Shading & highlights	.987	.978	.972	.962
No motion				
Shading	.856	.896	.843	.614
Texture	.761	.867	.829	.683
Highlights	.874	.918	.845	.618
Shading & texture	.870	.945	.884	.774
Shading & highlights	.900	.949	.777	.910

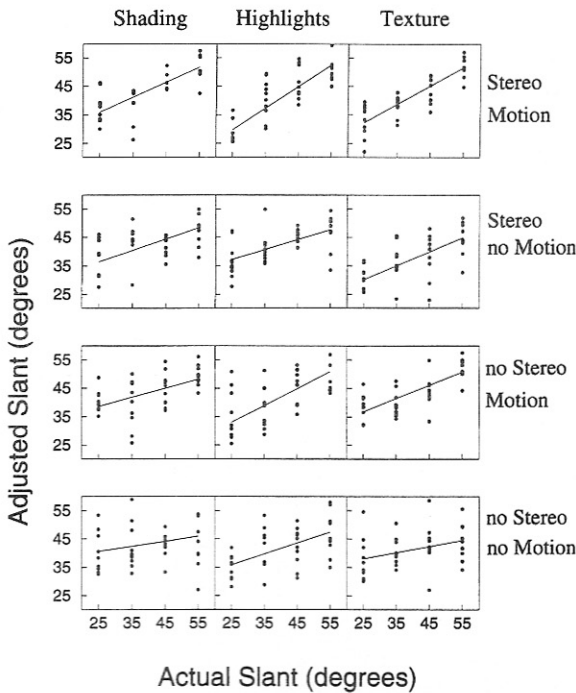


Figure 7. Scatterplots of the slant component of adjusted orientation for Observer J.F.N. as a function of simulated slant. The solid lines indicate the best-fitting linear regression line.

photographs (Koenderink & van Doorn, in press; Koenderink, van Doorn, & Kappers, 1992, 1994, in press; Koenderink et al., in press; Todd, Koenderink, van Doorn, & Kappers, in press). In all of these earlier studies, observers made repeated adjustments to individual probe points, so that it was possible to obtain a pure measure of their test-retest reliability. The general result for all of the viewing situations investigated to date is that the tilt component of observers' judgments has much smaller variance than does the slant component.

In addition to the random errors in repeated observations of individual probe points, there can also be large systematic constant errors. In the earlier studies of Koenderink et al. (1992, 1994, in press), Koenderink & van Doorn (in press), Koenderink et al. (in press), and Todd et al. (in press), each stimulus was sampled at numerous different probe points, so that it was possible to reconstruct the best-fitting smooth surface consistent with the overall pattern of an observer's adjustments. When any given stimulus display is judged by an observer over multiple occasions, the 3-D structure of this reconstructed surface remains quite stable. Its structure can vary dramatically, however, among different observers or for a given observer in different viewing conditions (e.g., when an object is viewed with different directions of illumination). In general, these differences are related by an overall scaling of the perceived surface relief in depth, which affects slant but not tilt. This same general pattern of perceptual distortion has also been confirmed using a variety of other procedures, and with

several different types of optical information, both individually and in combination (e.g., see Norman, Todd, Perotti, & Tittle, in press; Tittle, Todd, Perotti, & Norman, in press; Todd & Norman, 1991).

The methodology employed in the present experiment differs significantly from that of previous studies using the gauge figure adjustment task, in that a unique random stimulus was presented on each trial. The disadvantage of this approach is that it does not provide sufficient data for any given stimulus display to adequately reconstruct a complete 3-D representation of the observers' perceptions. The advantage, however, is that with sparse sampling of the individual objects, performance could be compared over a much broader range of stimulus conditions than would be possible if hundreds of trials were required for each display. We believe that these two approaches can complement one another, and that they are best used in combination. For example, a detailed reconstruction of observers' perceptions of a moving stereoscopic smooth surface with lambertian shading and specular highlights has been reported by Koenderink et al. (in press). Their findings demonstrate that observers can reliably perceive an object's 3-D structure from motion and stereo in the absence of any identifiable features. The present experiment enhances this result by comparing observers' performance in this condition with what would otherwise be possible with alternative combinations of optical information, both with and without identifiable texture elements.

Table 2
Correlation Coefficient (r^2) Values From the Linear Regression Analysis of the Slant Adjustments for all Four Observers

Slant Adjustment	Observer			
	J.F.N.	J.T.T.	V.J.P.	D.T.L.
Stereo				
Motion				
Shading	.583	.433	.373	.348
Texture	.708	.714	.599	.558
Highlights	.652	.495	.356	.517
Shading & texture	.581	.591	.643	.696
Shading & highlights	.521	.718	.329	.452
No motion				
Shading	.296	.417	.057	.133
Texture	.441	.585	.668	.787
Highlights	.366	.508	.202	.424
Shading & texture	.571	.787	.675	.716
Shading & highlights	.178	.517	.101	.093
No Stereo				
Motion				
Shading	.281	.433	.287	.197
Texture	.518	.529	.707	.457
Highlights	.404	.442	.293	.079
Shading & texture	.752	.556	.506	.379
Shading & highlights	.491	.579	.198	.016
No motion				
Shading	.051	.284	.010	.041
Texture	.122	.332	.423	.149
Highlights	.225	.308	.122	.070
Shading & texture	.464	.472	.330	.114
Shading & highlights	.015	.146	.108	.030

At present, there are few theoretical analyses to suggest how deformations of shading and highlight fields could be used to generate a useful representation of 3-D shape. The work of Koenderink and van Doorn (1980) is a notable exception. Their analysis demonstrated that saddle points in a lambertian shaded image where isoluminance contours cross at an X correspond to parabolic points on an object's surface. Over a sequence of views as an object is observed in motion, these saddle points could gradually trace out the parabolic lines on a surface separating regions of positive and negative gaussian curvature. It is important to keep in mind, however, that many of our conditions contained deformations of highlights, either by themselves or in combination with lambertian shading. There is no analysis yet available to suggest how 3-D structural information could be obtained from these types of deforming images. Thus, it remains for future theory to identify and reveal the particular aspects of these patterns that provide information about 3-D shape for human vision.

REFERENCES

- BLAKE, A., & BÜLTHOFF, H. (1990). Does the brain know the physics of specular reflection? *Nature*, **343**, 165-168.
- BLAKE, A., & BÜLTHOFF, H. (1991). Shape from specularities: Computation and psychophysics. *Philosophical Transactions of the Royal Society of London: Series B*, **331**, 237-252.
- BÜLTHOFF, H. H., & MALLOT, H. A. (1988). Integration of depth modules: Stereo and shading. *Journal of the Optical Society of America A*, **5**, 1749-1758.
- CIPOLLA, R., & BLAKE, A. (1990). The dynamic analysis of apparent contours. In *Proceedings of the Third International Conference on Computer Vision* (pp. 616-623). Los Alamitos, CA: IEEE Computer Society Press.
- ERENS, R. G. F., KAPPERS, A. M. L., & KOENDERINK, J. J. (1993a). Estimating local shape from shading in the presence of global shading. *Perception & Psychophysics*, **54**, 334-342.
- ERENS, R. G. F., KAPPERS, A. M. L., & KOENDERINK, J. J. (1993b). Perception of local shape from shading. *Perception & Psychophysics*, **54**, 145-156.
- GIBLIN, P., & WEISS, R. (1987). Reconstruction of surfaces from profiles. In *Proceedings of the IEEE First International Conference on Computer Vision* (pp. 136-144). IEEE Computer Society Press.
- HORN, B. K. P., & BROOKS, M. J. (1989). *Shape from shading*. Cambridge, MA: MIT Press.
- HORN, B. K. P., & SCHUNCK, B. G. (1981). Determining optical flow. *Artificial Intelligence*, **17**, 185-203.
- IKEUCHI, K., & HORN, B. K. P. (1981). Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, **17**, 141-184.
- JOHNSTON, A., & PASSMORE, P. J. (1994a). Shape from shading: I. Surface curvature and orientation. *Perception*, **23**, 169-190.
- JOHNSTON, A., & PASSMORE, P. J. (1994b). Shape from shading: II. Geodesic bisection and alignment. *Perception*, **23**, 191-200.
- KOENDERINK, J. J., KAPPERS, A. M. L., TODD, J. T., NORMAN, J. F., & PHILLIPS, F. (in press). Surface range and attitude probing in stereoscopically presented dynamic scenes. *Journal of Experimental Psychology: Human Perception & Performance*.
- KOENDERINK, J. J., & VAN DOORN, A. J. (1980). Photometric invariants related to solid shape. *Optica Acta*, **27**, 981-996.
- KOENDERINK, J. J., & VAN DOORN, A. J. (in press). Relief: Pictorial and otherwise. *Proceedings of the 5th British Machine Vision Conference*.
- KOENDERINK, J. J., VAN DOORN, A. J., & KAPPERS, A. M. L. (1992). Surface perception in pictures. *Perception & Psychophysics*, **52**, 487-496.
- KOENDERINK, J. J., VAN DOORN, A. J., & KAPPERS, A. M. L. (1994). On so-called "paradoxical monocular stereoscopy." *Perception*, **23**, 583-594.
- KOENDERINK, J. J., VAN DOORN, A. J., & KAPPERS, A. M. L. (in press). Depth relief. *Perception*.
- MARR, D. (1982). *Vision*. San Francisco: W. H. Freeman.
- MINGOLLA, E., & TODD, J. T. (1986). Perception of solid shape from shading. *Biological Cybernetics*, **53**, 137-151.
- NAGEL, H.-H. (1981). On the derivation of 3D rigid point configurations from image sequences. In *Proceedings of the IEEE Conference on Pattern Recognition and Image Processing* (pp. 103-108). New York: IEEE Computer Society Press.
- NAGEL, H.-H. (1987). On the estimation of optical flow: Relations between different approaches and some new results. *Artificial Intelligence*, **33**, 299-234.
- NORMAN, J. F., & TODD, J. T. (1994). Perception of rigid motion in depth from the optical deformations of shadows and occlusion boundaries. *Journal of Experimental Psychology: Human Perception & Performance*, **20**, 343-356.
- NORMAN, J. F., TODD, J. T., PEROTTI, V. J., & TITTLE, J. S. (in press). The visual perception of 3D length. *Journal of Experimental Psychology: Human Perception & Performance*.
- PENTLAND, A. (1991). Photometric motion. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, **13**, 879-890.
- TITTLE, J. S., TODD, J. T., PEROTTI, V. J., & NORMAN, J. F. (in press). The systematic distortion of perceived 3-D structure from motion and binocular stereopsis. *Journal of Experimental Psychology: Human Perception & Performance*.
- TODD, J. T. (1985). Perception of structure from motion: Is projective correspondence of moving elements a necessary condition? *Journal of Experimental Psychology: Human Perception & Performance*, **11**, 689-710.
- TODD, J. T., KOENDERINK, J. J., VAN DOORN, A. J., & KAPPERS, A. M. L. (in press). Effects of changing viewing conditions on the perceived structure of smoothly curved surfaces. *Journal of Experimental Psychology: Human Perception & Performance*.
- TODD, J. T., & MINGOLLA, E. (1983). The perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental Psychology: Human Perception & Performance*, **9**, 583-595.
- TODD, J. T., & MINGOLLA, E. (1984). The simulation of curved surfaces from patterns of optical texture. *Journal of Experimental Psychology: Human Perception & Performance*, **10**, 734-739.
- TODD, J. T., & NORMAN, J. F. (1991). The visual perception of smoothly curved surfaces from minimal apparent motion sequences. *Perception & Psychophysics*, **50**, 509-523.
- TODD, J. T., & NORMAN, J. F. (in press). The visual discrimination of relative surface orientation. *Perception*.
- TODD, J. T., & REICHEL, F. D. (1989). Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psychological Review*, **96**, 643-657.