Available online at www.sciencedirect.com

SCIENCE DIRECT®

Vision Research 46 (2006) 1057–1069

Vision Research

www.elsevier.com/locate/visres

Visual discrimination of local surface structure: Slant, tilt, and curvedness

J. Farley Norman [a,*], James T. Todd [b], Hideko F. Norman [a], Anna Marie Clayton [a], T. Ryan McBride [a]

[a] *Department of Psychology, Western Kentucky University, Bowling Green, KY 42101-1030, USA*
[b] *Department of Psychology, The Ohio State University, Columbus, OH 43210, USA*

Received 16 June 2005; received in revised form 20 September 2005

Abstract

In four experiments, observers were required to discriminate interval or ordinal differences in slant, tilt, or curvedness between designated probe points on randomly shaped curved surfaces defined by shading, texture, and binocular disparity. The results reveal that discrimination thresholds for judgments of slant or tilt typically range between 4° and 10°; that judgments of one component are unaffected by simultaneous variations in the other; and that the individual thresholds for either the slant or tilt components of orientation are approximately equal to those obtained for judgments of the total orientation difference between two probed regions. Performance was much worse, however, for judgments of curvedness, and these judgments were significantly impaired when there were simultaneous variations in the shape index parameter of curvature.

Keywords: Shape; Orientation; Curvature; Stereo

1. Introduction

One of the most perplexing phenomena in the study of human vision is the ability of observers to perceive the 3-D layout of the environment from patterns of light that project onto the retina. There are many different aspects of optical stimulation, such as shading, texture, motion, and binocular disparity, that are known to provide perceptually salient information about 3-D structure, but an effective computational analysis of this information has proven to be surprisingly elusive. One possible reason for this, we suspect, is that there has been relatively little research to identify the specific aspects of an object's structure that form the primitive components of an observer's perceptual knowledge. After all, to compute shape, it is first necessary to define what "shape" is.

Consider, for example, the shaded image of a smoothly curved surface that is presented in Fig. 1. Clearly, there is sufficient information in this image to produce a compelling impression of 3-D shape, but what precisely do we know about the depicted object that defines the basis of our perceptual representations? When observers are asked to verbally describe the shape of an object, their responses can typically be grouped into three basic categories: (1) They may identify some other object whose shape it resembles, (2) they may estimate the relative proportions of its height, width and depth, or (3) they may describe different parts of the object, such as regions of concavity or convexity. For example, when observers are asked to describe the object depicted in Fig. 1, the most frequent response is that it resembles a tooth. If asked to provide more detail, they typically point with their fingers to identify specific hills and valleys.

Although it is not generally revealed in verbal descriptions of 3-D shape, observers also have knowledge about the surface properties at individual points on an object.

* Corresponding author. Tel.: +1 270 745 2094; fax: +1 270 745 6934.
E-mail address: Farley.Norman@wku.edu (J.F. Norman).

0042-6989/$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2005.09.034

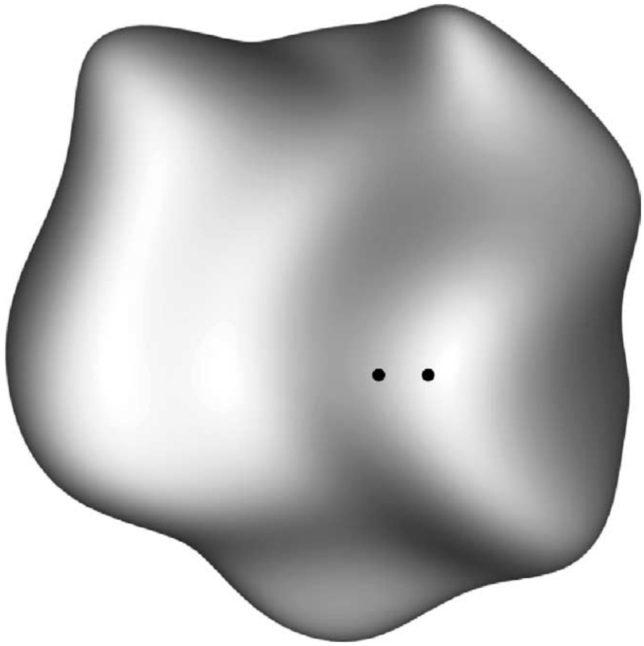

Fig. 1. A shaded image of a smoothly curved surface similar to those used in the present experiments. On a local probe task, observers are required to compare some specified local property (e.g., depth, orientation, or curvature) between two local regions that are marked by small dots.

For example, consider the pair of surface points that are marked in Fig. 1 by small black dots. When examining this figure, observers can easily identify which region appears closest in depth, which region appears more slanted in depth, or which region appears more curved. Indeed, one of the most common methods for representing smooth surfaces in theoretical models of human shape perception involves a map of the surface properties in each local neighborhood—a data structure that we will refer to generically as a local property map.

There are many possible local properties of a surface that could potentially be used for its perceptual representation, and it is useful to consider a few general factors by which they can be distinguished. One of these factors involves varying levels of differential structure. For example, a particularly common method of representing 3-D surface shape is to encode the depth (Z) of each visible point relative to the observer—what is sometimes referred to as a depth map, $Z = f(X, Y)$. A depth map represents the 0th order structure of a surface, but it is also possible to define higher order properties by taking spatial derivatives of this structure in orthogonal directions. The first spatial derivatives of a depth map (f_X,f_Y) define a pattern of surface depth gradients, and is often referred to as an orientation map. Similarly, its second spatial derivatives (f_{XX},f_{XY},f_{YY}) define the pattern of curvature on a surface.

Although a surface depth map is a scalar field, all higher order aspects of differential structure require multiple components to adequately represent each local surface point. There are a number of different coordinate systems in which these components could potentially be parameterized. One possibility for the first order differential structure would be to represent the surface depth gradients in the horizontal and vertical directions (f_X,f_Y), which is sometimes referred to as gradient space. An alternative possibility is to represent surface orientation using a viewer centered spherical coordinate system that is parameterized in terms of *slant* (σ) and *tilt* (τ), where: $\sigma = \arctan\sqrt{f_X^2 + f_Y^2}$ and $\tau = \text{arctan2}(f_X,f_Y)$. To better understand the nature of these parameters, it is useful to consider the possible optical projections of a circular disk at varying orientations relative to the line of sight (see Fig. 2). The optical projection of a circle is always an ellipse. The slant of the circle in 3-D space determines the aspect ratio of its projected ellipse, whereas the tilt component determines the orientation of the ellipse within the image plane. Note that the tilt parameter is undefined when $\sigma = 0$.

There are also different coordinate systems for representing second order differential structure. One could in principle use the second spatial derivatives of the surface depth map (f_{XX},f_{XY},f_{YY}), though it is usually more convenient to parameterize this structure in terms of surface curvature. The *normal curvature* (k) of a surface in any given direction (α) is defined as an angular change in the surface normal per unit arc length. Among all the possible directions at any given surface point, there will always be two principal directions (α_1, α_2): one that has the largest normal curvature (k_1) and another that has the smallest (k_2). As was first discovered by the Swiss mathematician Leonard Euler in around 1760, these principal directions of curvature are always orthogonal to one another. Some alternative

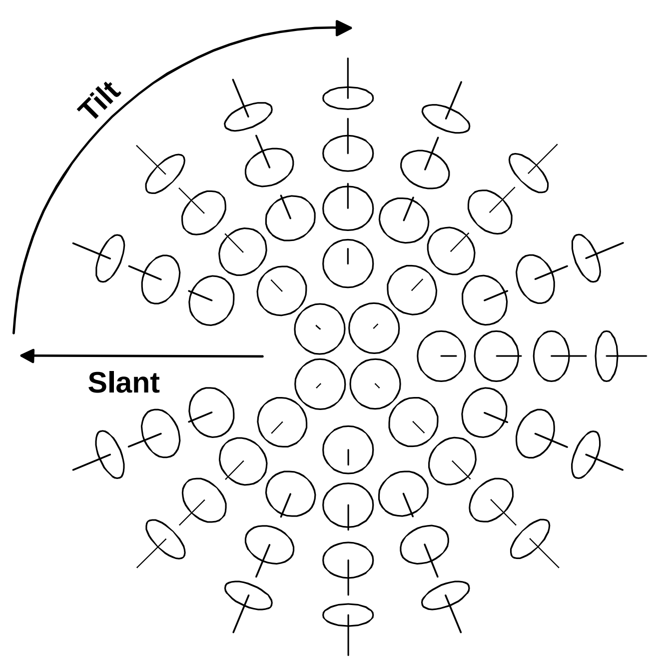

Fig. 2. A set of circular patches arranged on a sphere to illustrate the slant and tilt components of surface orientation. The line at the center of each patch is aligned in the direction of the surface normal. Note that the slant and tilt components form a spherical coordinate system, in which lines of latitude have constant slant, and lines of longitude have constant tilt.

coordinate systems have also been proposed for representing curvatures that are transformations of (k_1,k_2). One possibility is to parameterize the structure in terms of *mean curvature* (H) and *Gaussian curvature* (K), where: $H = (k_1 + k_2)/2$ and $K = k_1k_2$. Another possibility proposed by Koenderink (1990) is to parameterize the structure in terms of *curvedness* (C) and *shape index* (S), where $C = \sqrt{(k_1^2 + k_2^2)/2}$ and $S = -\frac{2}{\pi}\operatorname{atan}(\frac{k_1+k_2}{k_1-k_2})$. The advantage of this latter approach is that shape index is invariant over arbitrary scaling transformations. Thus, a point on a large sphere would have a different curvedness than would a point on a small sphere, but they would both have the same shape index. The variations of 3-D surface structure as a function of these parameters is illustrated graphically in Fig. 3. Note that the variations in shape index produce qualitatively distinct categories such as bumps, cylinders and saddles, whereas the overall magnitude of curvature is determined by the curvedness parameter.

Many researchers have argued that local property maps are a likely form of data structure for the perceptual representation of 3-D shape. This idea was first proposed over 50 years ago by Gibson (1950), and it was also advocated in an influential series of papers by Marr (Marr, 1978, 1980, 1982; Marr & Nishihara, 1978). He coined the term "$2\frac{1}{2}$-D sketch" to describe this type of representation, because it does not incorporate any information about parts of objects that are occluded by others. According to Marr (1982), the fundamental questions for a psychophysical investigation of the $2\frac{1}{2}$-D sketch are: (1) to determine which specific surface properties are perceptually represented (e.g., depth, slant, or curvature); and (2) to determine the precise coordinate system in which those properties are parameterized. The research described in the present article was designed to address these issues.

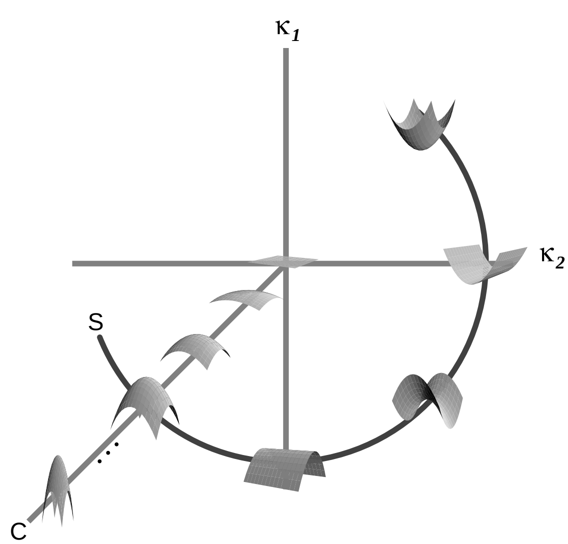

Fig. 3. The effects on surface structure from variations in the curvedness (C) and shape index (S) parameters of curvature, and the relation of these components to the principal curvatures (κ_1 and κ_2)—adopted from Phillips and Todd (1996).

2. The visual perception of local surface structure

Let us now consider what is currently known about the ability of human observers to judge various attributes of local surface structure. There have been many studies reported in the literature in which observers have been required to make judgments about depth, orientation, or curvature at designated probe points on a surface. For purposes of the present discussion, it is useful to separate these studies into two general categories. One popular methodology developed primarily by Koenderink, van Doorn, and Kappers (1992, 1995, 1996) uses judgments of local structure to investigate the perception of global 3-D shape. The basic idea of this approach is to mathematically reconstruct a best-fitting surface that is most consistent with an observer's judgments over an appropriately large sample of probe points. Surface reconstructions from local orientation judgments are particularly interesting in this regard, because it is not mathematically necessary for those judgments to be consistent with the gradient field of a smoothly curved surface. To allow such an interpretation, the overall pattern of an observer's orientation judgments must have a negligible amount of curl. The fact that the data are consistent with this requirement (Koenderink et al., 1992) provides some degree of confidence that the reconstructed surfaces obtained with these procedures may be indicative of a psychologically valid data structure. This confidence is bolstered, moreover, by the high test–retest reliability of these reconstructed surfaces over multiple experimental sessions (Todd, Norman, Koenderink, & Kappers, 1997), and the high correlations among different probe tasks (Koenderink et al., 1996; Koenderink, van Doorn, Kappers, & Todd, 2001). It is also interesting to note, however, that these reconstructed surfaces can be systematically distorted relative to the ground truth, and there can also be large individual differences among different observers (see Koenderink et al., 2001).

A second class of studies reported in the literature has been primarily concerned with the precision with which observers can make judgments of local surface properties (e.g., see Johnston & Passmore, 1994a, 1994b). For example, in one common paradigm, a surface is presented with two small probe points to designate the target regions whose local properties must be compared (e.g., see Fig. 1). This technique has been used to investigate observers' perceptions of both relative depth and relative orientation under a wide variety of conditions (Koenderink et al., 1996, 2001; Norman & Todd, 1996, 1998; Reichel, Todd, & Yilmaz, 1995; Todd & Norman, 1995; Todd & Reichel, 1989). The results of these studies have shown that observers are quite poor at discriminating the depth interval between two probe regions relative to a standard interval. Under full cue conditions, the Weber fractions for such judgments are often in excess of 25% (Norman & Todd, 1996). Other experiments have investigated the ability of observers to judge the difference in orientation between two probe regions. Discrimination thresholds for those judgments are typically in the range of 8° to 11° (Norman & Todd, 1996; Todd & Norman, 1995).

It is important to note that prior research on the perception of local surface properties has focused relatively little attention on how those properties are parameterized in observers' perceptual representations. Although there have been a few studies to investigate the shape index parameter for the representation of local curvature (e.g., de Vries, Kappers, & Koenderink, 1993; Mamassian, Kersten, & Knill, 1996; Phillips & Todd, 1996; van Damme & van de Grind, 1993; van Damme et al., 1994), other aspects of local differential structure have been largely neglected. In an effort to help alleviate this shortcoming, the experiments described in the present article were designed to measure observer sensitivity to several other aspects of local orientation and curvature, including slant, tilt, and curvedness. The stimuli depicted randomly shaped surfaces defined by shading and binocular disparity, with two small probe spheres to indicate the target regions to be compared (e.g., see Fig. 1). For a designated surface property (i.e., slant, tilt, or curvedness), observers were required to judge either interval or ordinal differences between the probe regions, irrespective of other local properties in those regions whose values varied randomly from trial to trial.

3. General methods

3.1. Apparatus

The stimuli were created and displayed on either a Silicon Graphics Crimson VGXT workstation (in Experiments 1–3), or an Apple Power Macintosh Dual-Processor G4 with an ATI Radeon 8500 graphics card (in Experiment 4). The displays were viewed stereoscopically using LCD (liquid crystal display) shuttered glasses such that each stereoscopic half image was presented in temporal alternation at an overall refresh rate of 120 Hz. The computer monitor had a spatial resolution of 1280 × 1024 pixels, and its visible display subtended $25.2° \times 20.3°$ of visual angle. Head movements were restricted using a chin rest.

3.2. Stimuli

The stimuli were similar to those developed originally by Norman, Todd, and Phillips (1995). On each trial, a stimulus was selected from a set of 100 randomly shaped objects that were created by distorting a sphere with a series of 10 sinusoidal perturbations in random orientations. The surface of each object was approximated with a dense triangular mesh, and was textured with a blue and white random check pattern. Texturing was achieved using an interpolation algorithm, in which each polygon was first rotated to a fronto-parallel orientation and then mapped into a random region of texture space. This procedure ensured that equal areas of the surface contained equal amounts of texture. These surfaces were shaded with an additive combination of a 70% diffuse (i.e. Lambertian) component and a 30% constant ambient component. The parameters of the reflectance model were constrained so that the shading of each pixel could not exceed the maximum possible image intensity. The surfaces were illuminated from an upper left direction with a tilt of 45° and a slant of 28° relative to the line of sight. The stereograms presented in Fig. 4 provide two representative examples of the objects created using this procedure.

From the 100 possible stimulus objects, a set of 100,000 vertex pairs were randomly selected to estimate the distributions of the various local surface properties that observers would be asked to judge. At each sampled vertex, we calculated the local surface slant, tilt, and curvedness, and for each pair of vertices we calculated the differences in slant, tilt, and curvedness. Frequency histograms for each of these measures were obtained by sorting the computed values into 100 discrete bins. For example, Fig. 5 shows the sampling distributions that were obtained for curvedness at each vertex and the curvedness differences between pairs of vertices. All of the probe points in the present experiments were selected by randomly sampling the appropriate bin from one of these distributions.

3.3. Procedure

On each trial, observers were presented with one of the possible stimulus objects with two small red probe spheres to designate the target regions whose local properties were to be compared. Each experiment employed one of two basic procedures depending upon whether the required judgment involved an ordinal or an interval relation. For ordinal judgments, observers were required to press a response key to indicate which probe region (left or right) had the largest value for a particular parameter of local surface structure (i.e. slant, tilt, or curvedness). For interval judgments, they were required to judge whether the difference between the probe regions along a designated dimension (i.e., slant, tilt, or curvedness) was greater or less than an implicitly defined standard (see McKee, 1981; Volkmann, 1932; Wever and Zener, 1928). For both types of procedure, observers received 20 trials of practice at the beginning of each block, and they received immediate feedback after each trial in the form of an auditory beep for correct responses. Observers were allowed to view the displays for as long as they desired before making a response.

4. Experiment 1—Judgments of slant

Experiment 1 was designed to investigate the precision with which observers can discriminate differences in the slant component of surface orientation, irrespective of other local aspects of surface structure.

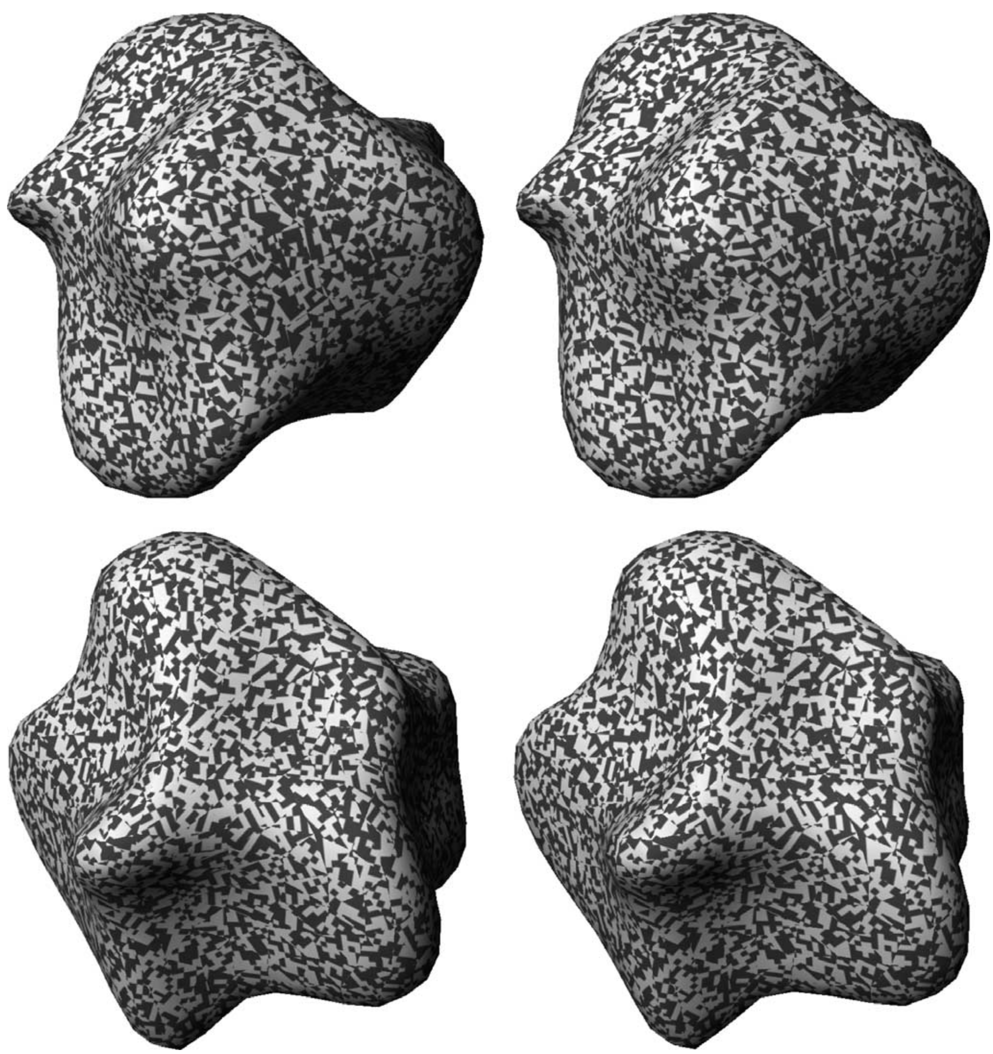

Fig. 4. Stereograms of two example objects that were used in the present experiments. The stereo pairs were designed for crossed free fusion.

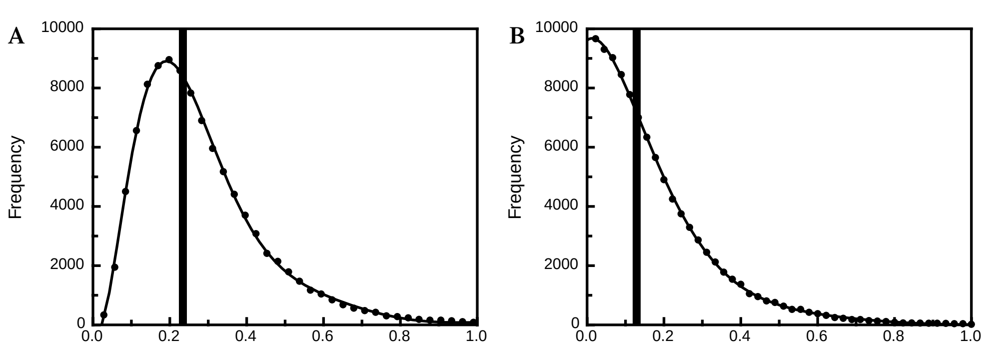

Fig. 5. The distributions of curvedness (A) and curvedness differences (B) for local regions of the objects used in the present studies. The vertical bar in each panel indicates the median of the distribution.

4.1. Method

On each trial, observers were presented with a randomly shaped object with two small red probe spheres to indicate the local regions whose slants were to be compared. As described in the general methods section, two different types of judgments were performed within separate blocks of trials. In the ordinal task, observers were required to judge which of the two designated surface regions (left or right) was more slanted in depth. One of the probe regions always had a standard slant of either 30°, 40°, or 50°, and the slant of the other region differed from the standard by either ±2°, ±6°, ±10°, or ±14°. In the interval task, observers judged whether the difference in slant between the two probe regions was greater or less than an implicit standard that they learned from the immediate feedback after each trial. There were three possible standard slant differences of 20°, 30°, or 40°, and the individual slants all ranged between 0° and 70°. The test slant differences differed from the standard by either ±11, ±33, ±55, or ±77%. For both types of response task, the tilts of the probe regions were varied randomly on each trial. Three of the authors (HN, JN, and JT) participated in the experiment. In each block of trials, the response task and the standard remained constant. Following a series of 20 practice trials, observers made 50 judgments for each of the eight possible test values in a random sequence. They performed two blocks of trials for each of the three possible standards for each response task.

4.2. Results and discussion

The data were analyzed using a probit analysis program developed by Foster and Bischof (1991) to find the cumulative normal distribution that best fit the overall pattern of performance by each observer for each possible combination of standard and response task. The discrimination threshold in each condition was estimated by halving the distance between the 25 and 75% points on the psychometric function. The results of this analysis are presented in Fig. 6, which shows the discrimination thresholds as a function of the standard for all three observers on each of the different response tasks.

There are two important aspects of these data that deserve to be highlighted. First, it should be noted that variations in the magnitudes of the standards had little or no effect on performance for either response task, which is consistent with previous findings on the perception of local surface orientation (Norman & Todd, 1996; Todd & Norman, 1995). A second thing to note in these data is that the thresholds for the ordinal judgments were significantly lower than those obtained for the interval judgments ($t(2) = 72.8$, $p < .001$). This should not be surprising. An accurate perception of ordinal relations can be achieved with a much weaker representation of local slant than would be necessary to discriminate the relative magnitudes of slant intervals (see Todd & Reichel, 1989). It is this latter task, however, that is most indicative of the ability of observers to perceptually represent slant with a consistent metrical scale.

It is especially interesting to compare these results with an earlier study by Todd and Norman (1995), in which the same three observers judged differences in orientation between designated probe regions under conditions that were nearly identical to those used in the present experiment. The average discrimination threshold they obtained was 8.9°. It is important to keep in mind that the orientation intervals observers were required to judge in that study involved simultaneous variations in both slant and tilt, and will therefore be referred to as total orientation differences. Note that this is quite different from the task employed in the present experiment, in which observers were required to base their responses solely on differences in the slant component of surface orientation, while ignoring simultaneous variations in the tilt component. If the perceptual representation of surface orientation is parameterized in terms of slant and tilt, as is commonly assumed, then we would expect that observers' sensitivity to the variations of one parameter should be as good or better than their sensitivity to the total orientation difference between two probe regions. The reason for this is that the computation of a total orientation difference requires an additional operation, in which the variations of individual parameters are metrically combined. The average discrimination threshold

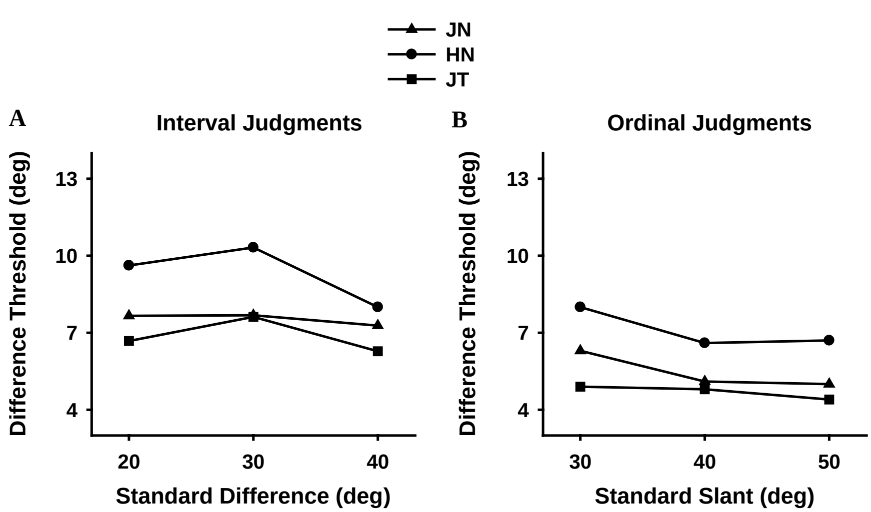

Fig. 6. The thresholds of slant discrimination as a function of the standard for all three observers in Experiment 1. The results obtained with interval judgments are shown in (A), and those obtained with ordinal judgments are shown in (B).

for interval slant judgments in the present experiment was 7.9°, which is slightly smaller than the 8.9° thresholds obtained by Todd and Norman (1995) for judgments of total orientation intervals. Thus, these results are consistent with the hypothesis that slant is a psychologically valid component for the perceptual representation of local surface orientation.

5. Experiment 2—Judgments of tilt

Experiment 2 was designed to investigate the precision with which observers can discriminate differences in the tilt component of surface orientation, irrespective of other local aspects of surface structure.

5.1. Method

On each trial, observers were presented with a randomly shaped object with two small red probe spheres to indicate the local regions whose tilts were to be compared. As in the previous experiment, two different types of judgments were performed within separate blocks of trials. In the ordinal task, observers were required to judge which of the two designated surface regions (left or right) had a more clockwise tilt. One of the probe regions always had a standard tilt of either 0°, 45°, or 90°, relative to the horizontal, and the tilt of the other region differed from the standard by either ±3°, ±9°, ±15°, or ±21°. In the interval task, observers judged whether the difference in tilt between the two probe regions was greater or less than an implicit standard that they learned from the immediate feedback after each trial. There were three possible standard tilt differences of 20°, 30°, or 40°, and the test tilt differences differed from the standard by either ±11, ±33, ±55, and ±77%. For both types of response task, the slants of the probe regions were varied randomly on each trial within a range of possible values between 30° and 70°. Three of the authors (HN, JN, and JT) participated in the experiment. In each block of trials, the response task and the standard remained constant. Following a series of 20 practice trials, observers made 50 judgments for each of the eight possible test values in a random sequence. They performed two blocks of trials for each of the three possible standards for each response task.

5.2. Results and discussion

Psychometric functions were again obtained using probit analyses to determine the observers' thresholds in each condition. The results are presented in Fig. 7, which shows the discrimination thresholds as a function of the standard for all three observers on each of the different response tasks. Note that in some respects the overall pattern of performance is similar to the one shown in Fig. 6 for judgments of surface slant. That is to say, the manipulations of the standard had only a minimal effect on observers' discrimination thresholds, and the thresholds for the ordinal judgments were significantly lower than those obtained for the interval judgments ($t(2) = 14.9$, $p < .01$). One important difference, however, is that the thresholds for slant in Experiment 1 were significantly smaller than the thresholds for tilt in Experiment 2 ($t(2) = 5.3$, $p < .05$). It is also interesting to note in this regard that the average threshold for tilt interval judgments was 35% higher than the threshold for total orientation judgments reported by Todd and Norman (1995). If tilt were a basic component in the perceptual representation of local surface orientation, then observers' sensitivity to relative tilt intervals should be as good or better than their sensitivity to total orientation intervals.

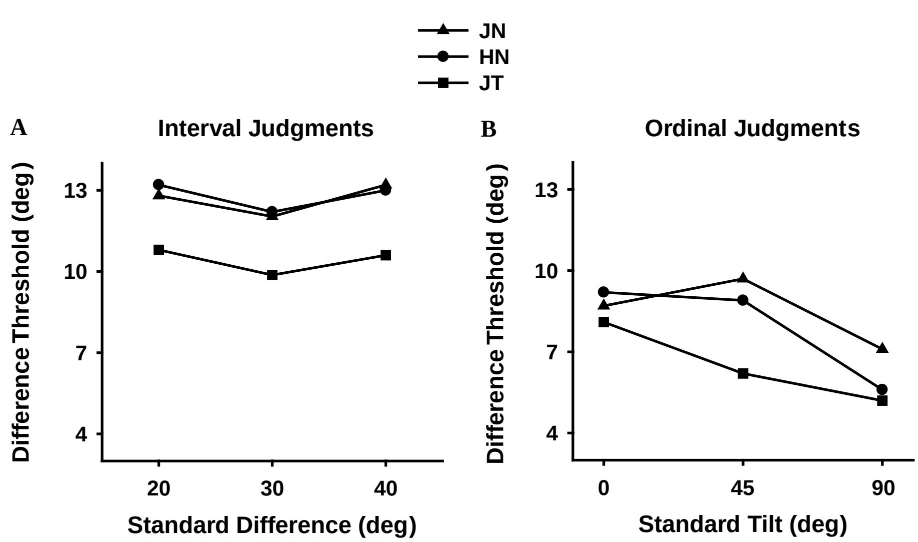

Fig. 7. The thresholds of tilt discrimination as a function of the standard for all three observers in Experiment 2. The results obtained with interval judgments are shown in (A), and those obtained with ordinal judgments are shown in (B). Tilt angles are defined in this context relative to the horizontal direction.

To better understand this seemingly strange result it is useful to represent surface orientations as points on a unit sphere, which is often referred to as the Gauss map (see Koenderink, 1990). Indeed, it is important to keep in mind that slant and tilt are the bases of a spherical coordinate system, in which lines of latitude connect regions with equal surface slant, and the lines of longitude connect regions with equal surface tilt. For example, consider the spherical surface with a slant–tilt coordinate system that is depicted in Fig. 8. Note that there are four local regions on this surface labeled **a–d**. Suppose that two regions have the same tilt but different slants as is depicted in Fig. 8 by the point pairs **ac** and **bd**. The overall orientation difference in that case is equal to the difference in slant between the two probe regions, and is invariant over changes in the shared value of tilt. That is not what occurs, however, in the converse situation where two regions have the same slant but different tilts, as exemplified by the point pairs **ab** and **cd**. If the tilt difference between two regions with identical slant were 45°, then the total orientation difference could range from 0° to 45° depending upon the shared value of slant.

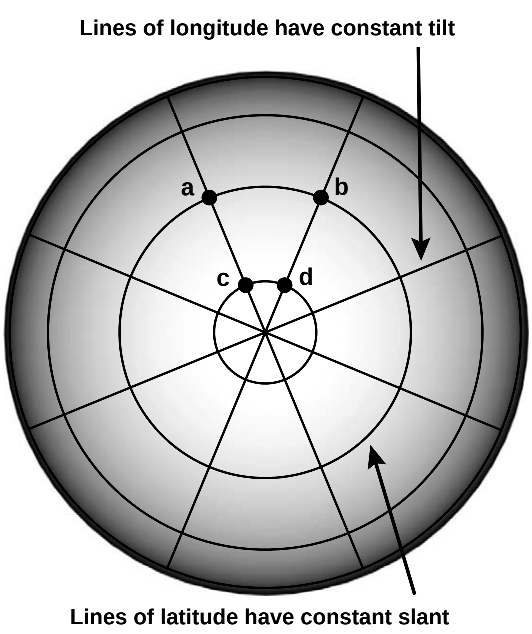

Fig. 8. A spherical representation of local orientation. Any possible orientation in 3-D space can be mapped to a point on the unit sphere called the Gauss map. The slant and tilt components define a spherical coordinate system, in which lines of latitude have constant slant, and lines of longitude have constant tilt. It is important to note that this is a viewer centered coordinate system, and that the point in the center where the lines of longitude intersect represents the surface orientation whose tangent plane is perpendicular to the line of sight. In this particular example, four different orientations labeled **a–d** are marked on the Gauss map by small dots. Note that the tilt angle between **a** and **b** is identical to the tilt angle between **c** and **d**, but that the separation between **a** and **b** in units of spherical arc length is greater than the separation between **c** and **d**. It is the arc length separation that determines the difficulty of tilt discrimination judgments.

In light of this observation, it seems reasonable to question whether the observers in the present experiment were judging differences in tilt angles per se, or some other related measure. There is an alternative way of scaling the difference in tilt between two orientations such that it cannot exceed their total orientation difference. The method is similar to the manner in which one might decompose a diagonal line segment (L) into its component distances (X, Y) in the horizontal and vertical directions (see left panel of Fig. 9). This is done by constructing a right triangle whose component lengths are defined by the Pythagorean theorem: $L^2 = X^2 + Y^2$. An analogous procedure can also be employed for decomposing differences in surface orientation by constructing a spherical right triangle on the Gauss map (see right panel of Fig. 9). If (o) is the total orientation difference between two regions, then the component differences (σ, τ) in the directions of slant and tilt are defined by the spherical Pythagorean theorem: $\cos(o) = \cos(\sigma)\cos(\tau)$. It is important to keep in mind that the arc length of a given tilt interval on the Gauss map increases systematically with the value of slant (see Fig. 8). Thus, if the observers in the present experiment were judging differences in arc length rather than tilt angles per se, then their tilt difference thresholds should be greater for small slants than for large slants. Experiment 3 was designed to test this prediction.

6. Experiment 3—Judgments of ordinal tilt at fixed slants

6.1. Method

This experiment was identical to Experiment 2 with the following exceptions: First, observers only made judgments about ordinal tilt. Second, only one tilt standard was employed, which was oriented in the vertical direction. Third, each of the probe regions on any given trial had the same value of slant, which varied across blocks with

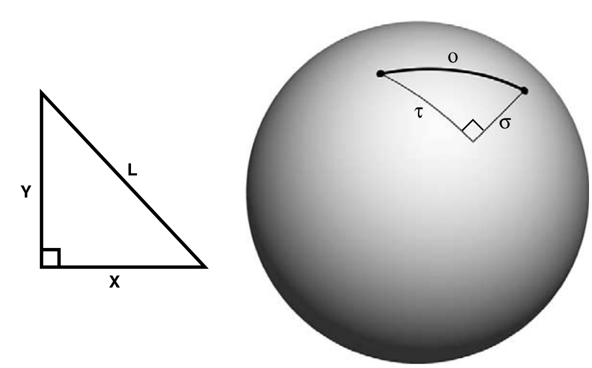

Fig. 9. The decomposition of length into orthogonal components. The left panel shows how a line segment L on a planar surface can be decomposed into horizontal and vertical components (X, Y) by constructing a right triangle, such that $L^2 = X^2 + Y^2$. A similar procedure can also be employed to decompose an orientation interval by constructing a spherical right triangle on the Gauss map. If (o) is the total orientation difference between two regions, then the component differences (σ, τ) in the directions of slant and tilt are defined by the spherical Pythagorean theorem: $\cos(o) = \cos(\sigma)\cos(\tau)$.

possible values of 30°, 45°, and 60°. Finally, the test tilts differed from the standard by ±3.5°, ±10.5°, ±17.5°, or ±24.5° for the 30° slants; by ±2.6°, ±7.8°, ±13.0°, or ±18.2° for the 45° slants; and by ±1.8°, ±5.4°, ±9.0°, or ±12.6° for the 60° slants. Observers performed two blocks of trials for each of the three fixed slants.

6.2. Results and discussion

The left panel of Fig. 10 shows the average tilt discrimination thresholds of all three observers as a function of the fixed slant. Note that the thresholds decrease systematically with increasing slant ($F(2, 4) = 83.3$, $p < .001$), as would be expected if observer sensitivity were based on differences of arc length on the Gauss map rather than tilt angles per se. In an effort to confirm this more clearly, we recalculated the thresholds as arc lengths on a unit sphere instead of tilt angles, and the results of this revised analysis are presented in the right panel of Fig. 10. Note in this case that the arc length thresholds remained approximately constant for all of the different values of fixed slant, thus suggesting that this may be the primary factor that determined the overall level of performance in each condition.

In light of this finding, we also recalculated the tilt discrimination thresholds obtained in Experiment 2 as arc lengths on a unit sphere. This was accomplished by randomly sampling point pairs whose slants were constrained as in Experiment 2 to be between 30° and 70°. For the results obtained for each observer in each condition, a sample of 200 point pairs was chosen all of which had a tilt difference that matched the observer's discrimination threshold. The arc length in the direction of tilt was calculated for each of the 200 point pairs, and the overall average was used to estimate the observers' thresholds in each condition.

Table 1 shows the average tilt arc length discrimination thresholds for the interval and ordinal judgments of each observer in Experiment 2, together with the slant discrimination thresholds from Experiment 1. It is especially important to note in these data that when the thresholds are measured in units of arc length on the Gauss map rather than tilt angles, the results reveal that observers are equally sensitive to orientation differences in the directions of slant and tilt (see also Todd et al., 1997). Moreover, their thresholds for judging individual components of orientation are also slightly less than the total orientation thresholds reported by Todd and Norman (1995) and by Norman and Todd (1996). These findings suggest, therefore, that observers can successfully decompose the directional components of local orientations without significant loss of precision.

There is one other aspect of the results relating to this issue that deserves to be highlighted. In Experiment 3, observers made judgments of ordinal tilt with a vertical tilt standard at fixed values of slant, and the average threshold in units of arc length was 3.8°. In Experiment 2, they also made judgments of ordinal tilt with a vertical tilt standard, but the slants of the probe regions were varied randomly. The average threshold in that case was 4.5°, which is not significantly different from the results obtained in Experiment 3 ($t(2) = 1.79$, $p > .2$). This finding provides especially strong evidence that observers are capable of isolating dif-

Table 1
Observer discrimination thresholds from Experiments 1 and 2 measured in degrees of arc length on a unit sphere

	Interval slant	Interval tilt	Ordinal slant	Ordinal tilt
JN	7.54	9.40	5.47	6.30
HN	9.31	9.57	7.10	5.93
JT	6.86	7.80	4.70	4.90

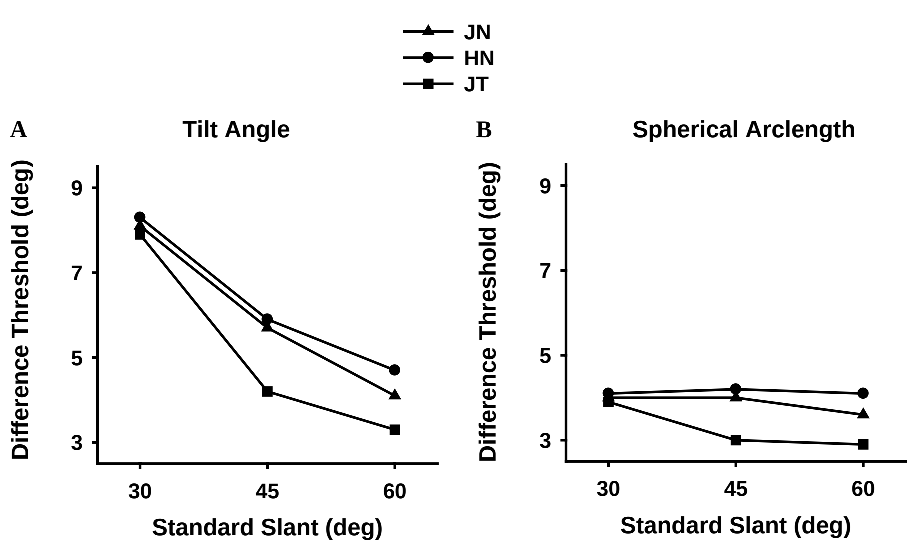

Fig. 10. The thresholds of ordinal tilt discrimination as a function of the standard slant for all three observers in Experiment 3. In (A) these thresholds are measured using the standard definition of tilt—i.e., $\tau = \arctan2(f_X, f_Y)$, and in (B) they are measured as arc lengths on a unit sphere.

ferences in one component of local orientation, while successfully ignoring differences that occur simultaneously in another.

7. Experiment 4—Judgments of curvedness

In contrast to the many studies that have attempted to measure observers' perceptions of local orientation, there has been relatively little research on the perception of local curvature (e.g., de Vries et al., 1993; Mamassian et al., 1996; Phillips & Todd, 1996; van Damme & van de Grind, 1993; van Damme et al., 1994). The majority of these studies have focused on a particular representation of curvature proposed by Koenderink (1990), in which local curvatures are parameterized in terms of curvedness and shape index (see Fig. 3). The advantage of this approach over other possible representations is that the shape index parameter is invariant over arbitrary scaling transformations. Thus, it captures the intuitive notion of most observers that a big sphere and a small sphere both have the same shape. The results of these studies have revealed that observers are most sensitive to differences in the local shape index parameter when the standard shape is a cylinder (de Vries et al., 1993; Phillips & Todd, 1996; van Damme & van de Grind, 1993; van Damme et al., 1994)—as opposed to surfaces that are curved in all directions.

The only previous study we know of that has investigated observer sensitivity to the curvedness parameter of local surface curvature was performed by van Damme and van de Grind (1993). They measured curvedness discrimination thresholds for quadric surface patches defined by motion parallax with fixed values of the shape index parameter. The average Weber fraction they obtained was approximately 20%, which, it should be noted, is an order of magnitude higher than the thresholds obtained for some other visual dimensions, such as brightness or length. Experiment 4 of the present series was intended to extend their study in an effort to determine if judgments of local curvedness are significantly affected by having simultaneous variations in shape index.

7.1. Methods

On each trial, observers were presented with a randomly shaped object with two small red probe spheres to indicate the local regions whose values of curvedness were to be compared. As described in the general methods section, two different types of judgments were performed within separate blocks of trials. In the ordinal task, observers were required to judge which of the two designated surface regions (left or right) had a greater degree of curvedness. One of the probe regions always had a standard curvedness of 0.234 cm^{-1} (see Fig. 11 for an explanation of this measure), which was the median of the distribution of curvedness values for the set of 100 possible objects employed in these experiments (see left panel of Fig. 5). The other region had a curvedness that differed from the standard by either ± 0.04, ± 0.08, ± 0.12, or ± 0.16 cm^{-1}. In the interval task, observers judged whether the difference in curvedness between the two probe regions was greater or less than an implicit standard that they learned from the immediate feedback after each trial. The standard curvedness interval was always 0.129 cm^{-1}, which was the median of the distribution of curvedness differences shown in the right panel of Fig. 5. The test differences differed from the standard by either ± 14, ± 42, ± 70, or ± 98%. For both types of response task, the shape indices of the probe regions were constrained in two different ways: In the random shape conditions, the shape indices of the probe points were selected at random from the distribution of possible values. However, in the similar shape conditions, the probe regions on each trial were constrained so that they had qualitatively similar shapes (i.e., bumps, dimples, or saddles) with shape indices that differed by no more than 0.1. Four of the authors (HN, JN, AC, and TM) participated in the experiment. Following a series of 20 practice trials, observers made 50 judgments for each of the eight possible test values in a random sequence. They performed four blocks of trials for each combination of response task and shape constraint.

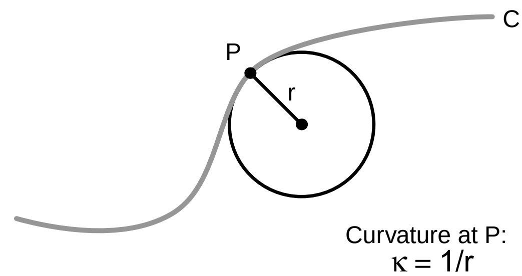

Fig. 11. Radius of curvature. The most common method of measuring the magnitude of curvature at any point (P) on a curve is to specify the radius (r) of the circle that most closely conforms to the shape of the curve at that point. This is also referred to as the osculating circle.

7.2. Results and discussion

The results obtained for each observer in each condition are presented in Table 2 as Weber fractions. Although the values of the standard were not manipulated in the present experiment, the earlier investigation by van Damme, Oosterhoff, and van de Grind (1994) found that ordinal judgments of curvedness satisfied Weber's law over a range of possible standards that included the one used in this study. An analysis of variance for these data revealed that the interval judgments produced significantly higher thresholds than the ordinal judgments ($F(1,3) = 130.68$, $p < .01$), and that the random shape conditions produced significantly higher thresholds than the similar shape conditions ($F(1,3) = 31.23$, $p < .025$). This latter finding is especially interesting because it indicates that observers were unable to isolate differences in curvedness, while ignoring orthogonal differences in the local shape index. Indeed, the curvedness thresholds obtained in the present study with random variations of

Table 2
Observer discrimination thresholds for curvedness in the different conditions of Experiment 4

	Interval thresholds different shape	Interval thresholds same shape	Ordinal thresholds different shape	Ordinal thresholds same shape
JN	71.2	59.9	32.1	23.1
HN	78.9	69.1	32.5	25.2
AC	99.2	81.6	43.2	29.5
TM	99.9	71.8	34.2	24.8

The thresholds are all expressed as a percentage of the standard—i.e., as Weber fractions.

shape index were almost twice as large as those reported by van Damme et al. (1994) when the shape index of the stimuli remained constant on each trial.

All of the observers complained at the conclusion of this study that the interval curvedness judgments were exceedingly difficult, and these subjective reports were clearly confirmed in the data. As is shown in Table 2, none of the observers was able to achieve a Weber fraction lower than 0.6 for interval curvedness judgments either with or without random variations in the local shape index. These findings would seem to suggest that observers are inherently incapable of perceiving the relative curvedness between two local regions with a consistent metrical scale.

8. General discussion

The research described in the present article was designed to investigate the abilities of human observers to perceptually decompose various aspects of local surface structure into multiple components. It is important to keep in mind that all higher order aspects of differential structure, including orientation and curvature, require multiple parameters to be adequately represented. Suppose, for example, that some local property P of a surface were perceptually represented in terms of two parameters (P_α and P_β). There are two fundamental predictions about perceptual performance that would be expected from this type of representation. Prediction 1: If P_α and P_β are represented independently, then judgments of one parameter should be unaffected by simultaneous variations in the other. Prediction 2: Observers' sensitivity to the variations of one parameter should be as good or better than their sensitivity to the total difference in P between two surface points. The reason for this is that the computation of a total difference requires an additional operation, such as a dot product, in which the variations of individual parameters are metrically combined. Let us now examine the extent to which observers' judgments about local orientation and curvature are consistent with these predictions.

8.1. Local orientation

Observers' judgments about slant and tilt, when defined in the traditional manner, do not satisfy either of the predictions described above. The results of Experiment 3 revealed that tilt discrimination thresholds vary systematically with the value of slant (see left panel of Fig. 10), which is incompatible with Prediction 1. Moreover, the results of Experiment 2 revealed that tilt discrimination thresholds are significantly higher than the thresholds for total orientation differences reported by Todd and Norman (1995). This finding is incompatible with Prediction 2. Based on these results, it would appear at first blush that slant and tilt may not be a psychologically valid data structure for the perceptual representation of local orientation.

A very different picture emerges, however, if differences in slant or tilt are rescaled as arc lengths on a unit sphere (see Fig. 9). When measured in that way, the results of Experiments 2 and 3 show that variations in slant have no significant influence on observer sensitivity to differences in tilt arc length (see right panel of Fig. 10), and that discrimination thresholds for intervals of tilt arc length are comparable to those obtained for total orientation intervals. Thus, if differences in slant and tilt are scaled appropriately, the results of these experiments are consistent with both of the key predictions described above.

Although the results of the present experiments show clearly that observers can decompose local orientation differences into components of slant and tilt, it remains to be determined whether those directions are in any way perceptually privileged, or if observers can adopt more flexible representations that are adapted to the required task. An important limitation of a slant–tilt parameterization is that it does not provide a straight forward procedure for computing total orientation differences without first converting the slant and tilt parameters to a different type of data structure. Thus, the preferred representation in virtually all computational applications is to parameterize orientations in terms of orthogonal direction cosines, because that is the most convenient data structure for performing vector addition and multiplication.

Could similar considerations be important for the representation of surface orientation in human perception? In order to address that question, it would be interesting in future research to investigate the abilities of observers to judge the individual components of local orientation differences in the horizontal and vertical directions, which can also be scaled into units of arc length on the Gauss map (see Fig. 9). If the thresholds for these judgments are similar to those obtained for judgments of slant and tilt, it would provide strong evidence that the perceptual representation of local orientation does not depend on a single privileged coordinate system.

9. Local curvature

The results of Experiments 1–3 suggest that the slant and tilt parameters of orientation are represented independently, such that judgments of one parameter are unaffected by simultaneous variations in the other. That does not seem to be the case however for judgments of local curvature. The results of Experiment 4 reveal that observers are unable to isolate differences in the curvedness parameter while ignoring simultaneous differences in the local shape index. Moreover, when observers are required to discriminate whether the curvedness difference between two regions is greater or less than a nonzero standard value, the task is basically impossible. The Weber fractions in that case are typically in excess of 70%, thus suggesting that curvedness may not be a perceptually relevant parameter for the representation of 3-D surface shape.

One aspect of curvedness judgments that may make them so difficult is that observers must combine information about surface curvature along multiple directions. If integrating information over different directions is a significant source of error in judgments of curvedness, it would explain why performance is improved when the regions to be compared have the same value of shape index. Accurate discrimination can be achieved in that case by only judging the magnitude of curvature along a single principal direction. This finding may suggest therefore that judging the magnitude of curvature in just one direction could be significantly easier than judging measures of curvature that combine information over multiple directions.

However, there are other sources of evidence to indicate that observers are able to combine information over multiple directions to make accurate judgments of the shape index parameter. For example, Perotti, Todd, Lappin, and Phillips (1998) used a global adjustment task to measure apparent 3-D shapes of quadric surface patches that were rotating in depth. The results revealed that judged shape index was almost perfectly correlated with the ground truth, but that the correlations with curvedness accounted for only 50% of the variance in observers' judgments. Accurate judgments of the shape index parameter have also been reported by de Vries et al. (1993) and van Damme et al. (1994), who found that observers' discrimination thresholds for moving or stereoscopic quadric surface patches could be as small as 1/100 of the total range of possible shape indices.

When evaluating these results on the perception of local shape, it is important to point out that the stimulus displays in all of these studies depicted large homogeneous surface patches that spanned at least 6.5° of visual angle. Phillips and Todd (1996) have shown that if the patch size is decreased to 2° or 3°, then the shape index discrimination thresholds are increased by an order of magnitude over what has been reported for larger displays. Performance is even worse when observers are required to make shape index judgments at designated probe points on randomly shaped surfaces. The just noticeable difference in that case can be as large as one fourth of the total range of possible shape indices (Phillips & Todd, 1996). These findings suggest that the accurate judgments of shape index reported in earlier studies may be based on the perception of global surface features, such as hills and valleys, rather than a map of the individual curvature parameters in each local neighborhood of a surface.

Acknowledgments

We thank Joseph Lappin for the detailed and helpful comments regarding a previous version of the manuscript. We also thank Flip Phillips for the assistance in preparing Figs. 3 and 11.

References

de Vries, S. C., Kappers, A. M., & Koenderink, J. J. (1993). Shape from stereo: A systematic approach using quadratic surfaces. *Perception & Psychophysics, 53*, 71–80.

Foster, D. H., & Bischof, W. F. (1991). Thresholds from psychometric functions: Superiority of bootstrap to incremental and probit variance estimators. *Psychological Bulletin, 109*, 152–159.

Gibson, J. J. (1950). *The perception of the visual world*. Boston, MA: Houghton Mifflin.

Johnston, A., & Passmore, P. J. (1994a). Shape from shading I: Surface curvature and orientation. *Perception, 23*, 169–189.

Johnston, A., & Passmore, P. J. (1994b). Independent encoding of surface orientation and surface curvature. *Vision Research, 34*, 3005–3012.

Koenderink, J. J. (1990). *Solid shape*. Cambridge, MA: MIT Press.

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. (1995). Depth relief. *Perception, 24*, 115–126.

Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. L. (1996). Pictorial surface attitude and local depth comparisons. *Perception & Psychophysics, 58*, 163–173.

Koenderink, J. J., van Doorn, A. J., Kappers, A. M. L., & Todd, J. T. (2001). Ambiguity and the "Mental Eye" in pictorial relief. *Perception, 30*, 431–448.

Mamassian, P., Kersten, D., & Knill, D. C. (1996). Categorical local shape perception. *Perception, 25*, 95–107.

Marr, D. (1978). Representing visual information. *Lectures on Mathematics in the Life Sciences, 10*, 101–180.

Marr, D. (1980). Visual information processing: The structure and creation of visual representations. *Proceedings of the Royal Society of London, B, 290*, 199–218.

Marr, D. (1982). *Vision*. San Francisco: Freeman.

Marr, D., & Nishihara, H. K. (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London, B, 200*, 269–294.

McKee, S. P. (1981). A local mechanism for differential velocity detection. *Vision Research, 21*, 491–500.

Norman, J. F., & Todd, J. T. (1996). The discriminability of local surface structure. *Perception, 25*, 381–398.

Norman, J. F., & Todd, J. T. (1998). Stereoscopic discrimination of interval and ordinal depth relations on smooth surfaces and in empty space. *Perception, 27*, 257–272.

Norman, J. F., Todd, J. T., & Phillips, F. (1995). The perception of surface orientation from multiple sources of optical information. *Perception & Psychophysics, 57*, 629–636.

Perotti, V. J., Todd, J. T., Lappin, J. S., & Phillips, F. (1998). The perception of surface curvature from optical motion. *Perception & Psychophysics, 60*, 377–388.

Phillips, F., & Todd, J. T. (1996). Perception of local three-dimensional shape. *Journal of Experimental Psychology: Human Perception & Performance, 22*, 930–944.

Reichel, F. D., Todd, J. T., & Yilmaz, E. (1995). Visual discrimination of local surface depth and orientation. *Perception & Psychophysics, 57*, 1233–1240.

Todd, J. T., & Norman, J. F. (1995). The visual discrimination of relative surface orientation. *Perception, 24*, 855–866.

Todd, J. T., Norman, J. F., Koenderink, J. J., & Kappers, A. M. L. (1997). Effects of texture, illumination, and surface reflectance on stereoscopic shape perception. *Perception, 26*, 806–822.

Todd, J. T., & Reichel, F. D. (1989). Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psychological Review, 96*, 643–657.

van Damme, W. J. M., & van de Grind, W. A. (1993). Active vision and the identification of three-dimensional shape. *Vision Research, 33*, 1581–1587.

van Damme, W. J. M., Oosterhoff, F. H., & van de Grind, W. A. (1994). Discrimination of 3-D shape and 3-D curvature from motion in active vision. *Perception & Psychophysics, 55*, 340–349.

Volkmann, J. (1932). The method of single stimuli. *American Journal of Psychology, 44*, 808–809.

Wever, E. G., & Zener, K. E. (1928). The method of absolute judgment in psychophysics. *Psychological Review, 35*, 466–493.