Perception, 1996, volume 25, pages 381 – 398

The discriminability of local surface structure

J Farley Norman, James T Todd
Department of Psychology, The Ohio State University, Columbus, OH 43210, USA;
E-mail: fnorman@magnus.acs.ohio-state.edu
Received 17 September 1995, in revised form 19 December 1995

Abstract. The ability of observers to discriminate depth and orientation differences between separated local regions on object surfaces was examined. The objects were defined by many optical sources of information simultaneously, including shading, texture, motion, and binocular disparity. Despite the full-cue nature of the displays, the observers' performance was relatively poor, with Weber fractions ranging from 10% to 40%. The Weber fractions were considerably lower for discriminations of surface-orientation differences than for similar discriminations of depth differences. The ability of observers to discriminate surface-orientation differences was approximately invariant over the separation of the regions in the projected image. In contrast, the ability to discriminate depth differences was highly influenced by the amount of image separation. This qualitative difference between the perception of depth intervals and surface-orientation differences suggests that knowledge of depths and orientations may be represented separately within the human visual system.

1 Introduction

Imagine that you are seated at your desk. There are many objects visible on its surface: a coffee mug, a paperweight, a stapler, some books, and perhaps an apple or a banana. Using vision, you can easily identify each item, because they have different shapes. How is this possible? One of the primary goals in perceptual research is to discover the precise aspects of shape by which visible objects are perceptually represented and to develop biologically plausible models of how these properties are extracted from patterns of optical stimulation.

One major difficulty in studying how humans perceive three-dimensional (3-D) shape involves selecting the particular attribute to be investigated from the myriad of possibilities. There are a large number of ways in which one could potentially describe the shape of an object (cf Koenderink 1990). Which attributes are the 'primitive' elements of human shape perception? Several possibilities seem intuitively obvious and readily come to mind. First, we have a sense of depth. Some parts of an object appear farther away from us than others. We are also aware of orientation differences—some surface regions face directly towards us, while others do not. Some objects, such as books, have large flat surfaces while others can be quite curved (eg an apple or baseball). Indeed, maps of local depths, orientations, or curvatures have often been proposed as potential human representations of shape. A good example would be the '$2\frac{1}{2}$-D sketch' developed by David Marr (1982).

Representations may exist at differing levels of differential structure. A depth map represents the zero-order structure of a surface, while orientation and curvature maps represent first-order and second-order differential structure, respectively. One could theoretically represent even-higher orders of differential structure as well—for example, changes in curvature across a surface. An additional layer of complexity is introduced when we consider how these levels of differential structure might be encoded by the visual system. For example, depth might be encoded as varying distances to object features in a direction parallel to the line of sight. Alternatively, the depth of the features of an object might be defined in an object-centered way by the distance of each feature from a central axis of symmetry. To completely represent higher-order properties like surface orientation and curvature, a vector quantity with two components is necessary. One possibility for a representation of surface orientation might be in terms of slant and tilt (cf Norman et al 1995) or in terms of depth gradients in the horizontal and vertical directions. Surface curvature could be encoded by the two principal curvatures at a point, or in terms of mean curvature and Gaussian curvature (Hilbert and Cohn-Vossen 1952). Another possibility for encoding curvature is in terms of shape index and curvedness (Koenderink 1990, pages 319-324). Last, our representations of shape could be primarily ordinal. An observer might be able to accurately identify ordering relationships in depth, but be unable to indicate anything precise about specific magnitudes. For example, an observer may know that one part of an object is closer than another, without necessarily knowing how much closer.

It is important to keep in mind that while we have some cognitive awareness of such particular properties as local surface depths, 3-D distances, and orientations, this level of knowledge may not necessarily be precise or reliable enough to account for human perceptual performance in everyday life. In fact, the available evidence indicates that our ability to accurately perceive these properties is rather limited. For example, using structure-from-motion displays, Todd and Bressan (1990) found Weber fractions of 25% to 50% for discrimination tasks involving 3-D lengths and angles. Reichel et al (1995) obtained Weber fractions of 20% to 50% for a depth-interval-discrimination task with surfaces defined by patterns of image shading and contour density. On a similar surface-orientation-discrimination task, Reichel et al found even-larger Weber fractions of 60% or more. McKee et al (1990) found that Weber fractions for discriminations of stereoscopic-disparity differences could be as high as 10% to 20%, while those for horizontal separations in the frontoparallel plane were 2% to 3% under the same experimental conditions. Last, Norman et al (1996) evaluated observers' ability to perceive 3-D lengths when both motion and stereoscopic information was available. Even when both of these sources of information were present, the observers' performance was relatively poor, with Weber fractions as high as 25%.

In all of the previously reviewed studies, the 3-D structure was specified by only one or two optical sources of information. It is possible that the perception of local 3-D structure may be enhanced and closer to veridical in more-full-cue situations where many redundant sources of information are simultaneously present. Therefore, the purpose of the current set of experiments was to probe observers' knowledge of local structure in full-cue situations. In particular, we investigated the perception of depth and orientation differences on smooth surfaces, using objects with naturalistic shapes. These objects were optically defined by motion, shading, texture gradients, and stereoscopic disparities, both horizontal and vertical.

2 Experiment 1

In the first experiment we investigated the ability of observers to compare the magnitudes of depth intervals between two local regions of the surface of an object with that of a standard depth difference. More specifically, observers had to indicate whether a given test interval appeared to be greater than or less than an implicit standard. The two regions of interest were each marked by a small circular spot presented in a plane tangent to the surface of the object. The randomly shaped solid objects were generated by a procedure described in earlier research (Norman et al 1995; Todd and Norman 1995; Koenderink et al, in press). The basic idea is that all of the solid objects were formed by an iterative process which sequentially added a series of sinusoidal distortions in random directions to a sphere. In particular, the initial spherical surface was modulated in depth by a unidimensional sine wave. The surface was then rotated a random amount about all three Cartesian coordinate axes. The surface in its new orientation was again modulated sinusoidally, rotated randomly, etc. This process was repeated ten times. The scaling factor of the modulation controls the amplitude of the bumps, dimples, and ridges of the resulting surface, while the number of iterations controls the apparent 'complexity'. A representative object obtained by following this procedure is shown in figure 1.

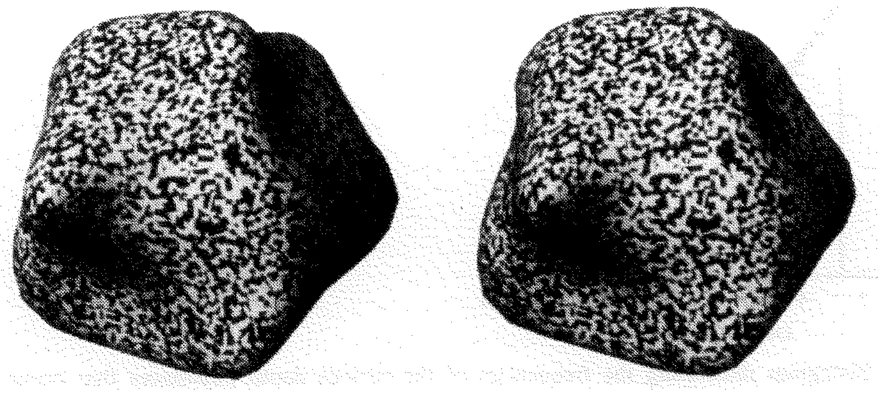

Figure 1. A stereogram of one of the shaded and textured objects that was used in the experiments.

The objects were defined by a polygonal mesh composed of 1280 individual triangles. On the initial generating sphere, the vertices of this mesh were evenly spaced over the surface of the object. Despite the discrete nature of this polygonal mesh, the resulting surface appeared smoothly curved, with no apparent discontinuities in orientation. This was due in part to the ability of our graphics system to interpolate the shading and texture of the surface within the interior of each triangular polygon.

Given that these objects had continuous surfaces, and therefore possessed a wide variety of depth differences, it was not initially obvious what the magnitudes of the depth standards should be. To help us understand the underlying structure of these shapes better, we performed a Monte Carlo simulation, picking arbitrary pairs of vertices from the polygonal mesh on 100000 random objects. For each trial of the simulation, we generated a randomly shaped object, and picked two potentially visible vertices (ie the normals of both vertices faced towards a hypothetical observer). We then calculated the depth difference for each of these 100000 pairs. Preliminary analysis showed that the depth differences ranged from near zero to approximately 14.0 cm. We then divided this range into one hundred discrete bins, and calculated the frequencies of vertex pairs in each one. The resulting frequency histogram is shown in figure 2.

The first thing to notice about this distribution is that the most common depth differences for these objects are small. Most of the pairs of surface regions have depth differences less than 3.0 cm, with a smooth exponential decline at higher and higher depth differences. The solid vertical line marks the median depth difference. Half of the surface pairs had depth differences less than 2.7 cm, while the other half had depth differences greater than 2.7 cm.

The purpose of the first experiment was to evaluate observers' ability to discriminate depth differences between separated local regions. A variety of standards were used to determine whether observers are uniformly sensitive at distinguishing between depth intervals over the entire range of possible differences.

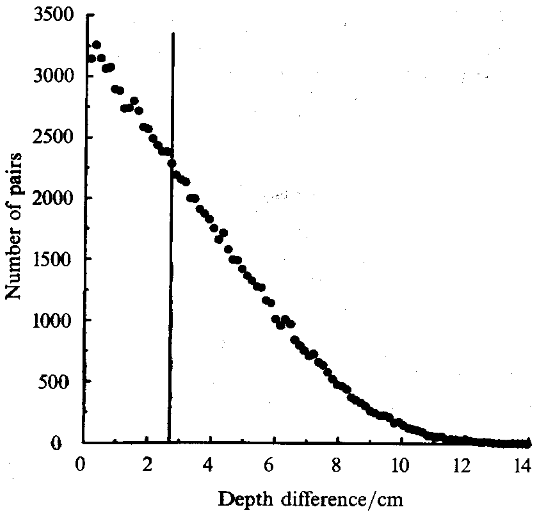

Figure 2. A histogram illustrating the frequencies of the various depth differences that occur between randomly selected pairs of regions on the object surfaces. The median depth difference is indicated by the solid vertical line.

2.1 *Method*

2.1.1 *Apparatus.* The optical patterns were created and displayed on a Silicon Graphics Crimson VGXT workstation with hardware texture-mapping capabilities. Stereoscopic-viewing hardware was used. The stereoscopic half images were presented by means of liquid-crystal shuttered glasses that were synchronized with the refresh rate of the monitor. 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 liquid-crystal 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, or 60 Hz. The viewing distance was 76.0 cm, such that the display screen, 1280 pixels wide by 1024 pixels high, subtended 25.2 deg by 20.3 deg.

2.1.2 *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. A set of 50 objects was generated at random by distorting spheres with an initial radius of 8.0 cm. This transformation was accomplished by 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 Norman et al 1995; Todd and Norman 1995; Koenderink et al, in press).

The objects were presented with motion, stereo, shading, and texture simultaneously. The surface shading was simulated by using a standard computer-graphics reflectance model (see Todd and 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. For the displays in the present experiments, the specular component of the shading model was set to zero, simulating the reflectance properties of a matte surface.

The lambertian component of the shading model was 70%, while the ambient component was 30%.

The shading was consistent with a single simulated 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 applied to the surface by means of a two-dimensional blue-and-white checkerboard pattern. Each polygon was first rotated to a frontoparallel orientation and then mapped onto a random region of the two-dimensional-texture pattern, which ensured that equal areas of the surface contained equal amounts of texture (see Todd and Mingolla 1984).

All of the displays were generated with the appropriate perspective for a 76 cm viewing distance. Each eye's view was computed on the basis of an interpupillary distance of 6.1 cm. The objects oscillated in depth about a vertical axis between -12 deg and $+12$ deg from their 'home' position, with a 2.0 deg angular displacement at each frame transition. Thus, each apparent-motion sequence was composed of a total of thirteen individual frames, which were updated at a rate of 20 Hz. On any given trial, these thirteen frames were oscillated continuously for as long as necessary for the observer to make a perceptual judgment.

The two local regions defining the depth interval to be discriminated on any given trial were highlighted on the surface by small red circular spots (0.2 cm radius) placed tangent to the surface at the appropriate locations. To ensure that the local surface region under each spot was visible and not too close to the occlusion boundary of an object, we chose probe points that had slants less than or equal to 70°. Otherwise, the selection of probe points was done entirely at random, so that some point pairs were close together in the projected image, while others were farther apart.

Since the depth interval between any two surface points changes over time as a 3-D object rotates, the observers made their judgments when the objects were in the middle or home position. To ensure that the judgments were made at the appropriate time, the red probe spots were turned on and off during the apparent-motion sequence. The probe points were turned on when the objects were in their home position, as well as one frame before and after home position. Since the duration of each frame was 50 ms (20 Hz), the total time the probe spots were visible during *each cycle* of the motion sequence was 150 ms. Preliminary pilot observation showed 150 ms to be a sufficient length of time for the observers to estimate the depth interval between separated surface locations. It is important to keep in mind that the observers could view each object and the two local regions of interest for as long as they wanted, until they were ready to make a response.

The depth difference between the two highlighted surface regions was varied from trial to trial. Within any given block of trials, there was a standard depth difference of either 1.35, 2.7, or 4.05 cm. The middle standard of 2.7 cm was the median of the distribution of possible depth differences for these objects (figure 2), while the other two standards were either higher or lower than this standard by 50%. For the 2.7 and 4.05 cm standards, the test depth intervals were either greater or less than these standards by 7%, 21%, 35%, or 49%. For the 1.35 cm standard, the test depth intervals were either greater or less than the standard by 10%, 30%, 50%, or 70%.

2.1.3 *Procedure.* The observers' task on each trial was to decide whether the depth difference between the two probe regions was greater or less than the implicit standard for that block, and they were allowed to view each display for as long as necessary to make this decision. We used the method of single stimuli (Wever and Zener 1928; Volkmann 1932; McKee 1981) to obtain psychometric functions and to estimate observers' thresholds for discriminating differences in depth. In the traditional method of constant stimuli, the observers are shown both the standard and test stimulus on every trial, whereas in the method of single stimuli they are required to judge the test stimulus relative to an implicit standard which is never seen. It has been found that if observers are given feedback and a sufficient number of practice trials, they easily form a mental representation of the standard against which they can compare the test stimuli. This method is more efficient than traditional methods, because it requires fewer stimulus presentations. It has also been shown to provide similar estimates of observers' discrimination thresholds.

There were eight test depth differences for each of the three implicit standards, for a total of twenty-four different experimental conditions. In each block of trials, an observer was presented with all of the different test stimuli for a given standard, presented fifty times each in a random order. Twenty practice trials were shown at the beginning of each block so that observers could learn the magnitude of the standard, and they were provided with immediate feedback after every trial in the form of an auditory beep for correct responses. Observers participated in two experimental sessions for each of the three standards. Each session consisted of one block of four hundred trials plus twenty practice trials. After both sessions for each standard, one hundred trials had been obtained for each of the twenty-four experimental conditions.

2.1.4 *Observers.* The displays were presented to three observers, two of whom were authors (JFN and JTT). HFN was an experienced psychophysical observer, but was unfamiliar with the specific issues involved in the experiment. All observers had normal or corrected-to-normal vision.

2.2 *Results*

The observers' thresholds for all three standards are shown in table 1. The psychometric functions were estimated by finding the cumulative normal distribution that best fit the performance for the eight test depth intervals at each standard. This was done by using a probit-analysis program developed by Foster and Bischof (1991). The thresholds for each standard were estimated by halving the distance between the 25% and 75% points on the psychometric functions. These thresholds were then converted to Weber fractions by dividing by the magnitude of the standard. The Weber fractions range from about 15% to over 40%. They are highest at the smallest (1.35 cm) standard and decrease at the larger standards.

Table 1. Results of experiment 1: Weber fractions expressed as a percentage of the standard.

Observer	Standard depth interval			
	1.35 cm	2.70 cm	4.05 cm	mean
JFN	43.7	26.2	20.7	30.2
HFN	37.8	24.0	17.6	26.5
JTT	41.3	20.6	15.7	25.9
mean	40.9	23.6	18.0	

In general, the observers' performance was surprisingly imprecise, given that Weber fractions for many visual discrimination tasks are 5% or below. For example, the discrimination of the lengths of lines presented in the frontoparallel plane produces Weber fractions of 1%–3% (Fechner 1860/1966, page 181; Norman et al, 1996), while those for velocity discriminations are 5% or less (McKee 1981). For some tasks, human observers are extremely sensitive at discriminating small differences. For example, Weber fractions of 0.3% have been reported for discriminations of auditory pitch (Boring et al 1948). The Weber fractions that we obtained during the present experiment are more than an order of magnitude higher.

The fact that our observers' Weber fractions are high at the low standard and are low at the high standard suggests that they might have an absolute difference threshold rather than the relative one that produces Weber's Law. Figure 3 is a plot of the absolute difference thresholds (in cm) as a function of the magnitude of the standard. The solid lines indicate the predictions of Weber's Law. This was determined by finding the Weber fraction that best fits each observer's difference thresholds by minimizing the least squared differences. The dashed horizontal lines show the predicted performance if the observers had a constant absolute difference threshold. Two of the observers' (JFN and HFN) thresholds fall somewhere in between these predictions, while JTT's thresholds are consistent with the idea of a constant absolute threshold. JTT needed an approximate 0.6 cm difference to discriminate greater than or less than, no matter what the size of the standard depth difference.

Another way to consider the observers' performance in addition to the Weber fractions and absolute difference thresholds is to determine how many of the possible pairs of vertices on the object surfaces have depth differences that are perceptually indistinguishable from the standard depth differences. To illustrate the utility of this concept, consider the following scenario. Imagine that for a given observer we have obtained a depth-interval-difference threshold by using a set of 3-D objects. We then find upon investigation that fully half of the possible pairs of vertices that one might pick on the surface of the objects had depth differences that could not be discriminated from a standard. In a very real sense, this would seem to be poor perceptual performance. If, on the other hand, only 1% of the possible vertex pairs on the object surfaces had depth differences less than an observer's threshold, then the performance would be more impressive. This perspective helps put the magnitude of the difference threshold into a meaningful context by determining how common various depth differences are for a given class of objects. We performed an exhaustive Monte Carlo analysis for each observer and standard, picking 100 000 randomly chosen pairs of vertices on the 50 potato-like objects that we used during the experiment. We then counted how many of those pairs had depth differences that were above and below the observers' difference thresholds. The percentages of pairs below threshold are indicated in table 2. The percentages range from about 15% to 21%. They do not vary appreciably with the magnitude of the standard for observers JFN and HFN. Remember that the Weber fractions for these observers did decline with increasing standard magnitude. In contrast, JTT's percentages of the population of pairs below threshold did decline at the higher standards, following the trend of the Weber fractions.

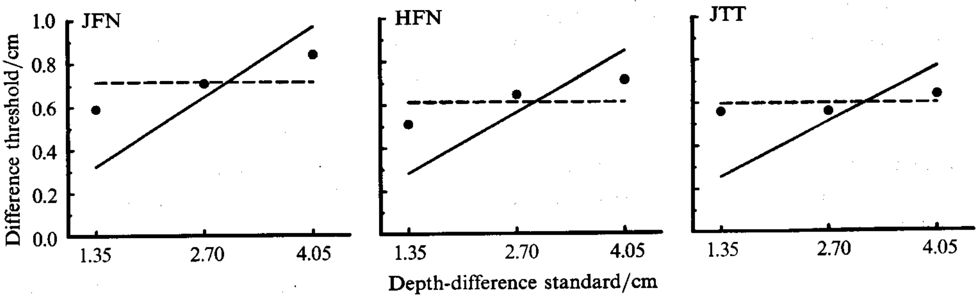

Figure 3. Individual plots showing how each observer's absolute difference thresholds vary as a function of the magnitude of the standard depth difference. The predictions of a constant absolute difference threshold and that of Weber's Law are plotted by the dashed and solid lines, respectively.

Table 2. Results of experiment 1: percentages of the population of depth differences below threshold.

Observer	Standard depth interval			
	1.35 cm	2.70 cm	4.05 cm	mean
JFN	21.4	21.3	19.8	20.8
HFN	18.6	19.8	17.1	18.5
JTT	20.0	16.8	14.7	17.2
mean	20.0	19.3	17.2	

3 Experiment 2

In the previous experiment, the locations of the probe points were picked entirely at random, subject only to the constraint that the slant of each probe region was 70° or less. For smooth, continuous, undulating surfaces like those used in the current experiments, previous research has shown that the local geometric structure of the surface may influence or interfere with observers' ability to make judgments about depth relationships between separated locations (Todd and Reichel 1989; Koenderink and van Doorn 1995; Reichel et al 1995). Todd and Reichel showed observers surfaces defined by image shading and marked two local regions with small circular spots. The observers were asked to indicate which of the two regions was closer in depth. They found that performance was higher and reaction times were shorter when the depth of the intervening surface between the two probe points changed monotonically. On the other hand, when the two regions were separated by a ridge or trough, the observers' ability to judge the order relations in depth deteriorated. This difference in performance occurred despite the fact that there were identical depth differences in both cases. Using a similar ordinal-discrimination task, Koenderink and van Doorn found good performance for nearby surface locations, but poor performance when the points were widely separated in the image. They concluded that the ability to judge depth order was relatively accurate only when the points were on a single 'slope', and not when separated by a depth reversal. Reichel et al found similar deficits in performance on a depth-interval-discrimination task when the depths between surface regions did not change monotonically.

These psychophysical results suggest that observers' perceptions of depth order and depth interval are affected by factors other than the magnitude of the difference in depth between surface regions. However, the surfaces in all three of these earlier studies were optically defined by monocular sources of 3-D information, namely shading and/or texture. These perceptual limitations may not exist when observers are given more optical information about 3-D structure to support their judgments. The purpose of our second experiment was to test this hypothesis. In particular, we investigated whether the amount of spatial separation between the local surface regions defining the test depth intervals affected the observers' ability to discriminate the intervals as being larger or smaller than the standard interval. In some conditions, the two local regions were constrained to be close together in the projected image, while in others the two regions were widely separated.

A second purpose of this experiment was to compare the perception of depth intervals on smooth surfaces with depth intervals between isolated points in 3-D space. It is an interesting fact that human observers can perceive the 3-D structure of a cloud of random points (ie a volume) in laboratory displays, even though such a display would be rare in a natural environment. Also it is common for observers to estimate distances between different objects that are separated in space. At this point in time there is little psychophysical evidence to show whether the perception of depth intervals in these two cases is similar or dissimilar. In both cases (intervening surface or no intervening surface) the binocular disparity or motion parallax between any two given points in space are equivalent. What differs is the context in which these depth intervals appear.

3.1 *Method*

All details of the stimulus displays were identical to those of experiment 1, except for the following differences. In this experiment, only the middle standard depth interval was used, 2.7 cm. There were six basic conditions, formed by the orthogonal combination of two spatial separations in the image between the probe points ('near' and 'far') and three types of optical patterns (solid objects depicted with stereoscopic disparities, texture, shading, and motion; solid objects depicted with stereo, texture, and shading, but without motion; and no solid objects visible, just the probe points, presented stereoscopically against a grid placed in the plane of fixation). This grid subtended 15 deg both horizontally and vertically, and was composed of fifteen evenly spaced vertical bars and fifteen evenly spaced horizontal bars. The 'near' probe points were constrained to have a spatial separation in the image of 4.0 cm or less. The 'far' probe points had a spatial separation of 4.0 cm or more. In the last two types of optical patterns, the probe points were visible continuously, since there was no motion present in the displays and the depicted depth intervals did not change over time.

There were eight test depth intervals for each of the six experimental conditions. For the solid-object-with-motion conditions, the test intervals were greater or less than the standard by 7%, 21%, 35%, and 49%. For the solid objects without motion, the test intervals were greater or less than the standard by 4.5%, 13.5%, 22.5%, and 31.5% for the 'near' image separations, and 7%, 21%, 35%, and 49% for the 'far' image separations. For the probe-point-only conditions, the test intervals were greater or less than the standard by 3.5%, 10.5%, 17.5%, and 24.5% for the 'near' image separations, and 5%, 15%, 25%, and 35% for the 'far' image separations.

There were separate blocks of trials for each one of the six base conditions. Each block consisted of fifty trials for every one of the eight test depth intervals plus twenty practice trials. Two blocks were run for each one of the six base conditions, resulting in one hundred trials per test depth interval. The same three observers as in experiment 1 participated in experiment 2.

3.2 *Results*

The Weber fractions for the observers are shown in table 3. Consistent with earlier findings, the simple effect of separation in the projected image had a large impact on the observers' perceptions of the depth intervals between surface regions. For the solid objects, the thresholds for the near condition were about half of those for the far condition (13.5% vs 26%–38%, respectively). It was much easier to make discriminations of depth intervals when the two surface regions were close together in the

Table 3. Results of experiment 2: Weber fractions expressed as a percentage of the standard.

Observer	Condition	Solid object: stereo and motion	Solid object: stereo only	No object: probe points only	mean
JFN	near	13.8	13.7	10.9	12.8
	far	34.2	38.5	15.9	29.5
HFN	near	13.3	12.6	8.6	11.5
	far	26.2	28.5	10.1	21.6
JTT	near	13.7	12.2	9.3	11.7
	far	26.6	20.8	14.1	20.5

projected image. This means that the perception of depth intervals is not as straightforward as performing some calculation based upon the stereoscopic disparities or velocities themselves, since the stereoscopic-disparity differences and velocity differences (in the conditions with motion) were the same in both the near and the far conditions. Although we did not explicitly test for the effects of ordinal transitivity in this experiment, it is likely that the presence or absence of monotonicity in the intervening surface region contributed to the large difference in performance between the near and far conditions. Given the nature of the shapes, the depth changes between the probe points in the near condition were more likely to be monotonic than the depth changes in the far condition. The differential results of these two conditions may be a reflection of the same underlying phenomenon as that studied by Todd and Reichel (1989) and Reichel et al (1995).

It would appear from the results that the motion of the objects contributed little to the observers' perceptions of depth intervals. The performance for the stereo-only conditions is very similar to that of the conditions with stereo and motion. Another interesting finding is that the very best performance for discrimination of depth intervals occurred in the conditions where the probe points were displayed by themselves, without the solid objects. Performance decreased, rather than increased, when the solid objects were present. It would appear that the undulations and variations in surface structure in between the probe points masked the observers' ability to abstract out the depths of the two surface regions. Furthermore, the large effect of spatial separation between the near and far conditions with the solid objects was significantly reduced when only the probe points were present. This was confirmed by a significant interaction in a repeated-measures two-way analysis of variance between near/far and the three types of optical patterns ($F_{2,4} = 8.779$, $p < 0.05$).

4 Experiment 3

In the previous two experiments we investigated observers' ability to discriminate the magnitudes of depth intervals under a wide variety of conditions. The purpose of experiment 3 was to similarly investigate the perception of surface-orientation differences. For several reasons, surface orientation has often been postulated to be a fundamental perceptual variable. For example, the patterns of shading on an object with a rough or matte surface is directly related to how local surface patches are oriented with respect to point light sources like the sun. Texture and velocity gradients in a projected image are also governed by local surface orientations. For these kinds of reasons, surface orientation was chosen as a primary component in Marr's (1982) $2\frac{1}{2}$-D sketch.

We have previously (Todd and Norman 1995) evaluated observers' ability to discriminate the magnitudes of surface-orientation differences both for smoothly curved objects like those in the present experiments and for dihedral-angle surfaces composed of two flat planes. The perceptual performance for the dihedral angles was much higher than that for the smooth surfaces. The Weber fractions for the dihedral angles were about 7%–10%, while those for the smooth surfaces were higher by a factor of about three (20%–35%). We concluded that observers had difficulty in estimating the surface orientations of local regions on smoothly curved objects. However, since the standards used were 25°, 35°, and 45°, only a subset of possible orientation differences was investigated. In general, visible surface regions can differ in orientation by up to 180°, for example at opposite sides of a spherical ball.

In the current experiments we sought to extend our 1995 study by examining observers' performance for a wider variety of standard orientation differences. To facilitate the choice of standards, we again performed a Monte Carlo analysis to learn more about the distributions of surface-orientation differences on these smoothly curved objects, which was directly analogous to that of experiment 1. The results are shown in figure 4.

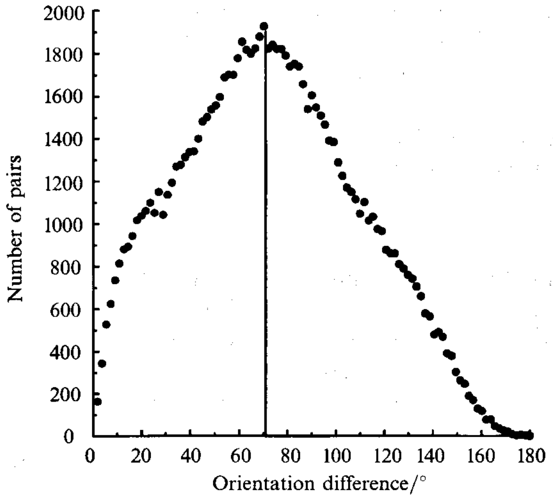

Figure 4. A histogram illustrating the frequencies of the various orientation differences that occur between randomly selected pairs of regions on the object surfaces. The median orientation difference is indicated by the solid vertical line.

This distribution of surface-orientation differences looks quite different from the distribution of depth differences for these objects. The depth distribution was shaped like an exponential decay, while the surface orientation distribution is shaped like an inverted U. The most common surface-orientation difference appears to be approximately 70°. As in experiment 1, we used the median of this distribution as a standard, along with standards that were either higher or lower than this median by 40%. This procedure ensured that a wide variety of surface orientation differences was explored in the current experiment.

4.1 *Method*

The stimulus displays were identical to those of experiment 1 with the following exceptions. In experiments 1 and 2, the probe points were flat circles that were presented tangent to the object surfaces at the desired locations. These circles in 3-space project to ellipses in the optical projection. This was acceptable for the depth-discrimination experiments, because there is no information in each projected ellipse that could inform the observer about how the depth of that region compares with the depths of other spatially separated regions. However, there is some information about the orientation of a surface region from such projected ellipses. The surface slant can be calculated from the eccentricity of the projected ellipses, if one assumes the projected ellipse is actually a circle in 3-space. The surface tilt can be obtained from the orientation in the image of the minor semiaxis of the projected ellipse. The orientation of the surface could be reconstructed from these slant and tilt values. It would be theoretically possible then to determine the surface-orientation difference between separate regions. This is probably not perceptually possible for the human visual system (as suggested by our earlier results; Todd and Norman 1995), but we wanted to avoid the possibility in the current experiment by using a different type of probe point. Instead, we used small spherical probe points with a 0.125 cm radius. These spherical probe points always project to small circles in the projected image; therefore they serve only to mark the surface regions the observers should attend to, and do not themselves contain any information about the orientation of the underlying local regions.

The orientation difference between the two highlighted surface regions was varied from trial to trial. Within any given block of trials, there was a standard difference of either 42.5°, 70.9°, or 99.2°. For the 42.5° and 70.9° standards, the test differences were either greater or less than these standards by 5%, 15%, 25%, or 35%. For the 99.2° standard, the test orientation differences were either greater or less than the standard by 3.5%, 10.5%, 17.5%, and 24.5%.

Once again, one hundred trials were obtained for the eight test differences for each of the three standards. The three observers were the same as those who had participated in experiments 1 and 2.

4.2 *Results*

The Weber fractions for orientation difference discrimination are shown in table 4. They range from about 10% to 20%, and are about half of those found for the depth-interval discriminations in experiment 1. The observers were able to make finer discriminations of orientation differences than for depth differences. The Weber fractions are highest at the lowest-magnitude standard and are lowest at the highest-magnitude standard. This suggests once again that the observers may have an approximately constant absolute difference threshold. To explore this possibility, in figure 5 the absolute difference thresholds are plotted as a function of the magnitude of the standard. Our earlier results (Todd and Norman 1995) are also indicated, since we used the same task with a different range of standard orientation differences. The best-fitting predictions of Weber's Law and the predictions if the observers had a constant absolute difference threshold are shown. It is visually evident that the pattern of results for each observer does not conform to the predictions of Weber's Law. The hypothesis of a constant difference threshold of around 9° – 10° seems more consistent with the results.

Table 4. Results of experiment 3: Weber fractions expressed as a percentage of the standard.

Observer	Standard orientation difference			
	42.5°	70.9°	99.2°	mean
JFN	22.1	14.5	9.3	15.3
HFN	19.9	18.3	12.8	17.0
JTT	16.8	14.9	11.0	14.2
mean	19.6	15.9	11.0	

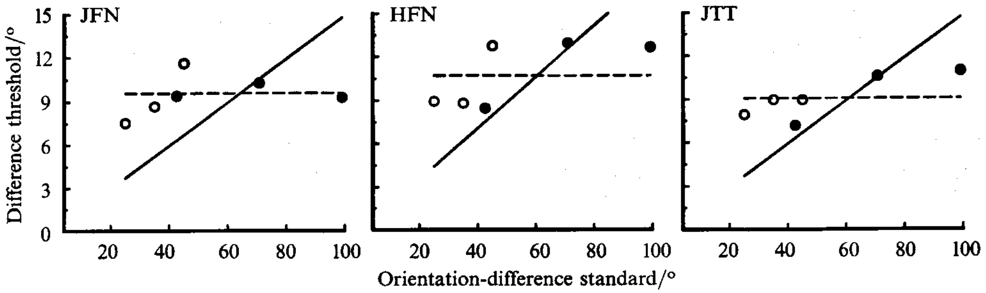

Figure 5. Individual plots showing how each observer's absolute difference thresholds vary as a function of the magnitude of the standard orientation difference. The predictions of a constant absolute difference threshold and that of Weber's Law are plotted by the dashed and solid lines, respectively. The results of experiment 3 are indicated by the filled circles, while the open circles indicate the results obtained by Todd and Norman (1995).

The percentages of the population of surface-orientation differences that were below each observer's thresholds are shown in table 5. These population measures were calculated by using a Monte Carlo analysis identical to that performed in experiment 1. The percentages of the population below threshold range from about 11% to 20%, and are roughly of the same magnitude as the observers' percentages of the population of depth differences that were below threshold (table 2). The wide disparity in performance between the depth-interval-difference and surface-orientation-difference tasks when expressed in Weber fractions is diminished when performance is evaluated in terms of this population measure.

Table 5. Results of experiment 3: percentages of the population of orientation differences below threshold.

Observer	Standard orientation difference			
	42.5°	70.9°	99.2°	mean
JFN	14.8	20.9	14.5	16.7
HFN	13.6	26.2	19.5	19.8
JTT	11.3	21.2	16.3	16.3
mean	13.2	22.8	16.8	

5 Experiment 4

The results of experiment 2 showed that the magnitude of the separation of the probe regions in the image had a large effect upon the observers' ability to accurately perceive the depth intervals. We wondered how variations in the image separation would affect the observers' discriminations of surface-orientation differences. The main purpose of experiment 4 was to investigate this issue. In experiment 2, the perception of the depth intervals was also influenced by the presence of intervening surface structure—performance was enhanced when the solid objects were removed and only the two probe points were visible by themselves. This manipulation is not meaningful in the current experiment, since surface-orientation differences are not defined when the objects themselves are not present.

In the first two experiments, it was necessary for the two probe points on the object surfaces to 'blink' on and off during the conditions with motion, since the test depth intervals had the correct magnitude only when the objects were in the home, or central, position. For orientation-difference discriminations, this practice is not necessary since the relative orientation of two regions does not change as an object rotates. A secondary issue in the current experiment was to compare two motion conditions: with 'blinking' probe points and with continuously visible probe points.

5.1 *Method*

The stimulus displays were identical to those of experiment 3, with the following differences. Half of the conditions in the present experiment included 'blinking' probe points, similar to those of experiment 3. The remaining conditions included probe points which were always present, and they remained attached to the surface regions of interest while the object rotated. We used the middle standard from experiment 3 as the standard orientation difference (70.9°). In contrast to the displays of experiment 3, which employed randomly selected pairs of probe points, we varied the separation of the probe points in the projected image, in a manner identical to that of experiment 2. The eight test orientation differences were either greater or less than the standard by 5%, 15%, 25%, or 35%.

Once again, one hundred trials were obtained for each observer at each of the eight test orientation differences. The three observers were the same as those who had participated in experiments 1, 2, and 3.

5.2 *Results*

The Weber fractions for the observers are shown in table 6. In contrast to the large effects of image separation for depth-interval discriminations on object surfaces, there was little effect for similar judgments of surface-orientation differences. In experiment 2, the Weber fractions for the far separation were, on average, more than twice those for the near separation. In the current experiment, there was no difference between near and far for observer JFN, and only very small effects for observers HFN and JTT. There were also only small differences between those conditions with blinking probe points and those conditions where the probe points were continuously visible. Both main effects were not significantly different, as shown by a repeated-measures two-way analysis of variance (image separation, $F_{1,2} = 2.491$, $p > 0.05$; continuously visible probe points vs blinking, $F_{1,2} = 9.924$, $p > 0.05$).

While evaluating the results of this experiment, it is important to keep in mind the nature of the objects that were used. For projections of simple surfaces like a sphere, the amount of surface-orientation difference between regions varies as a direct function of their separation in the image. Regions that are more widely spaced have larger orientation differences. An observer might learn this covariation and perform the discrimination task accurately on the basis of simple image differences and not upon perceived 3-D structure. If this were true, the Weber fractions we found might be artificially low in comparison with those obtained for the depth-interval-discrimination task. However, our surfaces were considerably more complex than that of a sphere, and were specifically chosen because they drastically reduced this covariation between image separation and the magnitude of the surface-orientation difference. For example, we found hundreds of large orientation differences that were close together in the image. Such regions occur on opposite sides of an elongated ridge, or on opposite sides of a sharp peak or valley. In contrast, many objects contained large areas that were relatively flat, so that widely separated regions had small orientation differences. In fact, the complexity of our randomly shaped objects was necessary to test for near vs far effects in the present experiment. We located many pairs of probe regions that had each of the eight test orientation differences for both near and far image separations. For example, we found many far pairs of regions that had orientation differences 35% less than the standard of 70.9° (46.1°). We also found many near pairs of regions that had orientation differences 35% greater than the standard (95.7°). It is not possible to find such pairs of regions on a spherical surface of 8.0 cm radius, because they do not exist. Our stimuli were sufficiently rich to allow us to uncouple and disassociate the relationship between the magnitude of image separation and the

Table 6. Results of experiment 4: Weber fractions expressed as a percentage of the standard.

Observer	Condition	Probe points blink	Probe points continuously visible	mean
JFN	near	13.7	12.6	13.2
	far	13.5	13.5	13.5
HFN	near	16.8	13.7	15.3
	far	19.8	18.8	19.3
JTT	near	13.5	12.5	13.0
	far	15.1	12.9	14.0

amount of surface-orientation difference. It is highly likely, then, that the Weber fractions we obtained are a reasonable estimate of the observers' ability to perceive surface-orientation differences per se, rather than some simple two-dimensional difference in the projected image.

6 General discussion

In general, the performance for discriminations of depth and orientation differences on smooth surfaces was not especially accurate, at least when compared with performances obtained for other psychophysical tasks. The Weber fractions in the present set of experiments varied from about 10% to over 40%. Weber fractions of 1% to 3% are regularly reported in the psychophysical literature, and can be lower than 1% for some tasks, such as the discrimination of auditory pitch or the magnitude of electric shock (Boring et al 1948). The current levels of performance are similar to those found in earlier investigations of perceived 3-D structure involving motion (Todd and Bressan 1990), shading and/or texture (Reichel et al 1995), stereo and motion (Norman et al, 1996), and stereo, motion, texture, and shading (Todd and Norman 1995).

The surfaces used in the present set of experiments were defined by many redundant sources of optical information, including stereoscopic disparities, motion, texture gradients, and shading. The displays were viewed under full-cue conditions, with sufficient lighting so that the observers could see the placement of the monitor and the layout of objects in the room. The observers therefore had knowledge of the viewing distance to the monitor, provided by the convergence angle of the eyes. The observers could also view any given optical pattern for as long as they wanted, until they were ready to make a response. Despite the availability of all this optical information, our observers performed poorly. The results of experiment 2 also indicated that the presence of motion added little to the observers' ability to discriminate the magnitudes of depth intervals. This result is similar to one found by Norman et al (1996), where the perception of 3-D length was no better for conditions including both motion and stereo than for conditions that employed either stereo alone or motion alone.

One intriguing result occurred in experiment 2, where the performance for discriminations of depth intervals actually improved when the object surfaces were removed, so that only the probe points themselves were visible. The perception of depth intervals between separated regions on continuous surfaces seems to behave differently from equivalent judgments between points in 'empty' space. When no surface was present, the observers were able to use the individual binocular disparities of the two probe points themselves to make relatively accurate judgments of the depth interval between them. This level of performance declined when the object surfaces were present. Since the surfaces were textured, there were plenty of individual features around the probe points whose binocular disparities could have been used by the observers to make similarly accurate estimates of depth intervals. Nevertheless, the observers were apparently unable to use these individual feature disparities when the stereoscopic surfaces were visible. Obviously, the visual system is using some other optical information besides individual disparities to perceive the depth interval between spatially separated regions on the surface of an object. In particular, observers may be integrating over the many changes in surface orientation along the shortest path from one probe region to the other. If this is true, one might wonder why the human visual system has two mechanisms to estimate depth intervals, one for surfaces and one for comparing two areas not connected by a continuous intervening surface.

Previous research has also shown that the perception of depth on smooth surfaces may operate differently from the perceived depth of isolated individual features (Todd and Reichel 1989; Koenderink and van Doorn 1995). In particular, it has been suggested in this research that ordinal properties of surfaces, such as the presence or absence of a monotonic change in depth between regions, affect how observers perceive the local surface structure. Todd and Reichel found that when the depth between regions changed monotonically, performance on an ordinal-depth task was high. When a depth reversal occurred between the two local regions (for example, across a peak or trough), the observers' discrimination accuracy declined, while the length of time needed to make a judgment increased by about 50%. Likewise, Koenderink and van Doorn's observers found the same ordinal-depth task easy when the two points were on a single 'slope', but resorted to guessing when the points were widely separated on the surface of the object.

The findings of Todd and Reichel (1989) and Koenderink and van Doorn (1995) demonstrate the importance of ordinal transitivity for our perceptions of depth on surfaces. The results of experiment 2, that the magnitude of image separation greatly affects observers' ability to discriminate the magnitudes of depth intervals, may be a reflection of the same underlying phenomenon. In experiment 4, however, there was little or no effect of image separation on the observers' ability to discriminate surface-orientation differences. One possibility for this qualitative difference is that there might be multiple representations of 3-D form. If there were only one primary representation of form, from which knowledge of all other 3-D properties were derived, then one would expect that if there was a lack of some information in the primary representation, there would be similar deficits in the knowledge of all subsequent properties derived from that representation. For example, if the primary representation was composed of depth values and the magnitudes of apparent depth differences depended not only upon the actual depths, but also upon image separation, it would be difficult to understand how the perception of surface orientation or curvature would be unaffected by image separation. On the other hand, it seems unlikely that the perception of depth intervals is always derived from a local map of orientation or curvature values, since visible surfaces do not need to be present in order for us to perceive the depths of isolated points in space, as shown by the results of experiment 2.

The current results raise the possibility that there may be separate representations which encode the depths and surface orientations of smooth surfaces. Our findings are consistent with those of Johnston and Passmore (1994), who found similar qualitative dissociations between the perception of surface orientation and surface curvature from optical patterns containing stereo, shading, and texture. They found that as the slant of a point light source was varied, the observers' curvature and orientation thresholds were affected differently. In particular, as the slant of the light source was reduced towards 0 (aligned with the observers' line of sight) the curvature thresholds increased, while those for orientation decreased. In addition, when texture was added to the shaded surfaces, the orientation thresholds decreased but the curvature thresholds increased. From these results Johnston and Passmore concluded that our perceptions of surface curvature are not derived from an already existing representation of local surface orientation but are formed in parallel from the retinal input. In this vein, Droulez and Cornilleau-Pérès (1990) have shown how it is possible to obtain information about surface curvature directly from binocular-disparity fields and from the velocity fields that result when a 3-D object moves relative to an observer.

The body of research that has been accumulating over the past 30 years has shown that an observer's performance depends greatly on the particular 3-D task they are asked to perform. For example, if we are asked to discriminate whether a given surface is curved or flat, we can do so with amazing accuracy (Norman and Lappin 1992). We can also indicate the depth order of two points with stereoscopic disparities as low as 3 s arc (Westheimer and McKee 1978, 1979, 1980). At the same time, we cannot reliably discriminate between two different 3-D lengths oriented in different directions in space (Norman et al, 1996) or accurately perceive the depth interval between surface regions that are widely separated in the projected image (experiment 2; also Koenderink and van Doorn 1995). This pattern suggests that we, as humans, may have multiple representations of 3-D form that are not general, but are best suited for performing certain, specific tasks. In general, we perform poorly on those metrical tasks that require us to distinguish between various magnitudes of a given property like depth interval, 3-D length, or surface-orientation difference. In contrast, performance for nonmetrical tasks that only require an order judgment, such as nearer vs farther, shorter vs longer, can often be more accurate and occur with shorter response latencies (eg Todd and Reichel 1989). It is important to keep in mind that we may have knowledge of which we are not consciously aware and cannot use to make subjective perceptual judgments, but which is available to the motor system. This is suggested by the results of Loomis et al (1992), who found significant perceptual distortions in a depth-to-width-matching task, but no apparent distortions in depth when observers were asked to walk blindfolded towards a previously viewed target. This result suggests that motor activity based upon visual information may behave according to different rules. It will be the task of future research to fully resolve this possibility.

Acknowledgements. This research was supported by Air Force Office of Scientific Research grant F49620-93-1-0116 to James Todd.

References

Boring E G, Langfeld H S, Weld H P, 1948 *Foundations of Psychology* (New York: John Wiley)

Droulez J, Cornilleau-Pérès V, 1990 "Visual perception of surface curvature: The spin variation and its physiological implications" *Biological Cybernetics* **62** 211 – 224

Fechner G T, 1860/1966 *Elements of Psychophysics* Volume I (translation by H Adler) (New York: Holt, Rinehart and Winston) (Original work published 1860)

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

Hilbert D, Cohn-Vossen S, 1952 *Geometry and the Imagination* (New York: Chelsea)

Johnston A, Passmore P J, 1994 "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, Doorn A J van, 1995 "Relief: pictorial and otherwise" *Image and Vision Computing* **13** 321 – 334

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 and Performance*

Loomis J M, Da Silva J A, Fujita N, Fukusima S S, 1992 "Visual space perception and visually directed action" *Journal of Experimental Psychology: Human Perception and Performance* **18** 906 – 921

Marr D, 1982 *Vision* (San Francisco, CA: W H Freeman)

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

McKee S P, Levi D M, Bowne S F, 1990 "The imprecision of stereopsis" *Vision Research* **30** 1763 – 1779

Norman J F, Lappin J S, 1992 "The detection of surface curvatures defined by optical motion" *Perception & Psychophysics* **51** 386 – 396

Norman J F, Todd J T, Perotti V J, Tittle J S (1996) "The visual perception of three-dimensional length" *Journal of Experimental Psychology: Human Perception and Performance* **22** 173 – 186

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

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, Bressan P, 1990 "The perception of 3-dimensional affine structure from minimal apparent motion sequences" *Perception & Psychophysics* **48** 419 – 430

Todd J T, Mingolla E, 1983 "Perception of surface curvature and direction of illumination from patterns of shading" *Journal of Experimental Psychology: Human Perception and 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 and Performance* **10** 734 – 739

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

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

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

Westheimer G, McKee S P, 1978 "Stereoscopic acuity for moving retinal images" *Journal of the Optical Society of America* **68** 450 – 455

Westheimer G, McKee S P, 1979 "What prior uniocular processing is necessary for stereopsis?" *Investigative Ophthalmology and Visual Science* **18** 614 – 621

Westheimer G, McKee S P, 1980 "Stereogram design for testing local stereopsis" *Investigative Ophthalmology and Visual Science* **19** 802 – 809

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

© 1996 a Pion publication printed in Great Britain