Perception, 1995, volume 24, pages 855-866

The visual discrimination of relative surface orientation

James T Todd¶, J Farley Norman
Department of Psychology (¶ also at The Center for Cognitive Science), The Ohio State University, Columbus, OH 43210, USA
Received 5 November 1994, in revised form 2 March 1995

Abstract. In a series of three experiments, observers' discrimination thresholds for relative surface orientation were measured under full cue conditions, in which surfaces were specified by multiple sources of optical information including shading, texture, motion, and binocular disparity. The results revealed that observers' sensitivity to relative surface orientation varies for different types of structural configuration. Weber fractions were approximately 8% for planar patches connected to form a dihedral angle. They increased to 11% for planar patches that were spatially separated, and jumped to over 26% for patches that were part of a smoothly curved surface.

1 Introduction

In a remarkable paper written almost a half century ago, James Gibson (1950) proposed one of the first accounts of how visible surfaces in the environment could be perceptually represented. Gibson suggested that each local region of a surface can be characterized by a small set of phenomenal primitives, including its depth and orientation relative to the observer, and that the variation of these primitives across different regions could define the layout of surfaces in an entire scene. Gibson was particularly interested in the property of surface orientation. He described how orientation could be perceptually specified by gradients of optical texture (see Purdy 1958 for a more rigorous computational analysis), and he initiated a series of psychophysical experiments to measure observers' sensitivity to this information.

Many years later, a similar hypothesis was proposed in an influential paper by Marr and Nishihara (1978). They also argued that visible surfaces are perceptually represented as a collection of local depths and/or orientations relative to the observer, and that these are the primitive units of our immediate phenomenal awareness of three-dimensional (3-D) form. They adopted 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. However, in order to account for the ability of observers to recognize objects at arbitrary orientations, they also argued that there must be higher-level, object-centered representations of 3-D structure that are invariant over different viewpoints.

It is interesting to note that just before their deaths, Gibson (1979) and Marr (1982) each wrote a book that reflected upon their many accomplishments in the study of visual perception. Marr described the discovery of the $2\frac{1}{2}$-D sketch as "the most exhilarating moment of the whole investigation". Although it is clear that Gibson was equally exhilarated when he had come up with the same basic idea early in his career, he eventually came to believe that this was one of his biggest mistakes.

Gibson changed his mind on this issue as a result of numerous experiments that he and others performed during the 1950s and 1960s on observers' perceptions of slanted planes (see Braunstein 1976 for an excellent review). In a typical paradigm of this period, observers would view a real slanted surface through an aperture and would be asked to adjust a response device such as a palm board to match its orientation in depth. The results of these experiments revealed that adjusted slants tend to be systematically underestimated, and that the variances obtained for multiple judgments of the same stimulus are often quite high. This is hardly what one would expect if local slants were the primary components of our immediate perceptual awareness, as suggested by Marr (1982).

Similar results have also been obtained for local-orientation judgments of more-complicated curved surfaces in a recent series of experiments by Koenderink et al (1992, 1994, 1995; see also Koenderink and van Doorn 1995, Koenderink et al 1996, Norman et al 1995b, Todd et al 1996). The paradigm employed in these studies involves adjusting the 3-D orientation of a circular gauge figure until it appears to rest in the tangent plane at a designated probe point on the surface of an object. The stimuli can be real objects viewed directly in physical space, or computer-generated pictorial displays with various different sources of optical information such as shading, texture, motion, and binocular disparity. From the adjustments obtained over many different probe points it is possible to compute the best-fitting smooth surface that is consistent with the observers' judgments. It is especially interesting to note, however, that this adjusted surface can vary dramatically among different observers, or for a given observer in different viewing conditions (eg when an object is viewed with different directions of illumination). There is also a substantial amount of random variation in the observers' judgments, such that the Weber fraction for adjusted local slant is typically around 25%.

It is important to keep in mind that the ability of observers to match the 3-D orientations of two surface patches does not necessarily require that they have a metrical knowledge of what those orientations are. Nor does it suggest that their underlying representations are scaled in a way that would make it possible to compare the relative orientations of different surface patches. Gibson (1979) made a distinction between optical orientation, which is defined relative to the observer's line of sight, and geographical orientation, which is defined relative to other visible surface regions. When an observer moves from one vantage point to another in an otherwise rigid environment, the optical orientations of all visible surface patches will change, but their relative geographical orientations will remain perfectly invariant. Thus, if a metrical representation of local orientation such as the $2\frac{1}{2}$-D sketch is to account for our perceptual awareness of environmental stability, then it must be sufficiently precise to allow a reasonably accurate computation of relative geographical orientations.

There have been only a few experiments reported in the literature in which observer sensitivity to relative surface orientation has been examined (eg Braunstein et al 1993; Reichel et al 1995; Tittle et al 1995), and most of those have involved displays with limited sources of optical information. The research described in the present paper was designed, therefore, to measure observers' discrimination thresholds for relative surface orientation under full cue conditions with shading, texture, motion, and binocular disparity. Three different tasks were employed in which observers judged the relative orientations of (i) irregular planar surface patches that were connected to form a dihedral angle, (ii) irregular planar patches that were spatially disconnected, and (iii) spatially separated regions on randomly structured smoothly curved surfaces.

2 Experiment 1

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 also used. The stereoscopic half images were presented by means of LCD (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 LCD glasses shuttered synchronously with the display so that each view of the stereo pair was seen only by the appropriate eye. The CRT was refreshed at 120 Hz. Thus each view of a stereoscopic half image was updated at half that, or 60 Hz. The viewing distance was 76 cm, so that the display screen, 1280 pixel wide by 1024 pixel high, subtended 25.2 deg by 20.3 deg.

2.1.2 *Stimulus displays.* On each trial, the display simulated a rotating dihedral angle formed by two differently oriented planes connected at an edge. The optical patterns were generated with the correct polar perspective for each eye based on an interpupillary distance of 6.1 cm. The individual planar patches of the dihedral angles had irregular outlines that varied randomly across trials, and they were covered with a naturalistic pebble-like texture, whose curved contours ensured that point, orientation, and curvature disparities were all present in the binocular optical patterns. A stereogram of a typical configuration is shown in figure 1.

The depicted orientation difference between the two facets of the dihedral angles was varied from trial to trial. On each block, there was a standard orientation difference of either 25°, 35°, or 45°, and the individual test stimuli all had orientation differences that were greater or less than this standard by 4%, 12%, or 20%. For example, when the implicit standard was 35°, the six test orientation differences were 28°, 30.8°, 33.6°, 36.4°, 39.2°, and 42°. In the initial frame of an apparent-motion sequence the depicted object was oriented so that the connecting edge between its two planar facets was parallel to the display screen, though its initial orientation about that edge was varied randomly across trials over a 45° range. The depicted objects oscillated around a vertical axis between −12° and +12° from their initial positions. They rotated in discrete jumps of 2.0° with a 50 ms stimulus onset asynchrony between successive rotary displacements. Thus, each apparent-motion sequence contained a total of thirteen individual frames that were presented at a rate of 20 Hz.

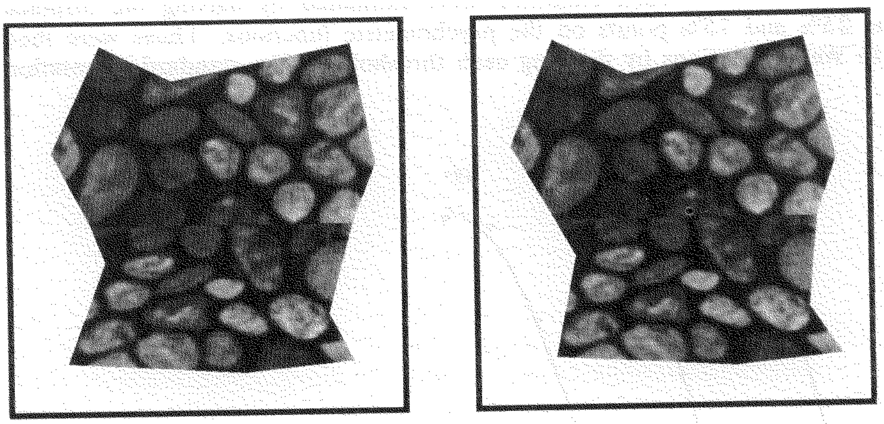

Figure 1. A stereogram of a textured dihedral angle similar to those used in experiment 1.

2.1.3 *Procedure*. The observers' task on each given trial was to decide whether the orientation difference between the depicted planar surfaces 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 (McKee 1981; Volkmann 1932; Wever and Zener 1928) to obtain psychometric functions and to estimate observers' thresholds for discriminating differences in orientation between the two planes of the dihedral angles. In the traditional method of constant stimuli, the observers are shown both the standard and the 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 (ie fewer stimulus presentations), and has been shown to give similar estimates of observers' discrimination thresholds.

There were six test orientation differences for each of the three implicit standards, for a total of eighteen 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, each of which included one block for each of the three standards (ie there were one hundred trials for each of the eighteen experimental conditions). The three blocks were presented in a random order in the first session, and they were then presented in counterbalanced order in the second.

2.1.4 *Observers.* The displays were presented to three observers, two of whom were authors (JN and JT). All observers had normal or corrected-to-normal vision.

2.2 *Results*

Figure 2 shows the percentage of trials for each test angle that were judged larger than the standard by observer JT in each of the three conditions. These results are also representative of those obtained for the other two observers. The smooth curves in this figure represent the best-fitting cumulative normal distributions, which were computed by using a probit analysis program developed by Foster and Bischof (1991). The thresholds in each condition were estimated by halving the distance between the 25% and 75% points on the psychometric functions. These were then converted to Weber fractions by dividing each threshold by the standard orientation difference.

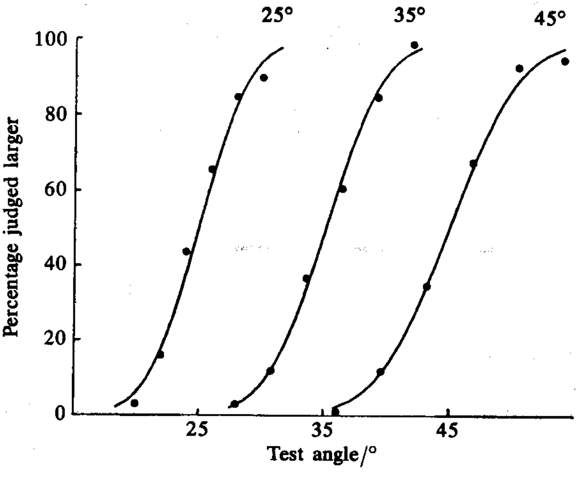

Figure 2. The psychometric functions of observer JT for the three standard angles employed in experiment 2.

The individual results for all three observers are presented in table 1. Over all of the different conditions employed in this experiment, the average Weber fraction for detecting differences among dihedral angles was approximately 8%, though it appears to vary somewhat with the magnitude of the standard. Note, in the table, that the Weber fractions of all three observers decreased by more than 20% as the standard angle was increased from 25° to 45°. A similar trend of increasing sensitivity to relative orientation with the magnitude of the standard has also been reported by Reichel et al (1995) for static monocular displays with shading and texture.

It is interesting to compare these findings with the results obtained in a similar study by Tittle et al (1995). They had observers adjust the angle between two random-dot planar surfaces specified by motion and stereo until they appeared to be orthogonal to one another in 3-D space. Weber fractions were estimated by the standard deviation of multiple judgments in each condition divided by its mean, and the results revealed an overall reliability of approximately 14%. There were two important aspects of the Tittle et al procedure that were most likely responsible for the reduced level of performance relative to that in the present experiment. First, the displays were viewed through a simulated aperture, which may have eliminated potential information from the edges of each surface patch. Second, the convergence angle was varied from trial to trial by using nonius markers, but the distance to the display screen was not, which could have created a potential conflict with accommodation. Thus, it is likely that the results obtained in the present experiment are more representative of what would be expected under full cue conditions in a natural environment.

It should also be kept in mind, on the other hand, that discriminations of dihedral angles may not be representative of observers' judgments of relative orientation for other possible structural configurations. In order to achieve accurate performance on this particular task, the depicted objects need only be specified up to an arbitrary conformal transformation (ie one that preserves local angles). Although a conformal distortion of perceptual space would allow accurate judgments of local angles, it would not in general preserve the relative orientations of visible surface regions that are spatially separated. In an effort to test the relevance of this distinction for human perception, experiment 2 was designed to measure observers' discrimination thresholds for the relative orientation of spatially separated planar patches.

Table 1. The Weber fractions for three observers for discriminating the relative orientations of connected planar patches in experiment 1.

Standard difference/°	Observer			Mean
	JN	HN	JT	
25	0.099	0.095	0.086	0.093
35	0.084	0.084	0.071	0.080
45	0.065	0.074	0.067	0.069
Mean	0.083	0.084	0.075	0.081

3 Experiment 2

3.1 *Method*

The apparatus and general procedure were identical to those used in experiment 1, with the following exceptions. The centers of the two depicted planar facets were separated horizontally in the image plane by 13 cm (ie 9.7 deg of visual angle), and there was a 4 cm gap between their respective upper and lower edges. Each patch rotated in phase about a separate vertical axis through its center. An example stereogram of a typical display is shown in figure 3. As in experiment 1, there were three possible standard orientation differences, 25°, 35°, and 45°. However, because our pilot research revealed that this task was more difficult, we expanded the range of possible test differences to be greater than or less than the standard by 5%, 15%, and 25%.

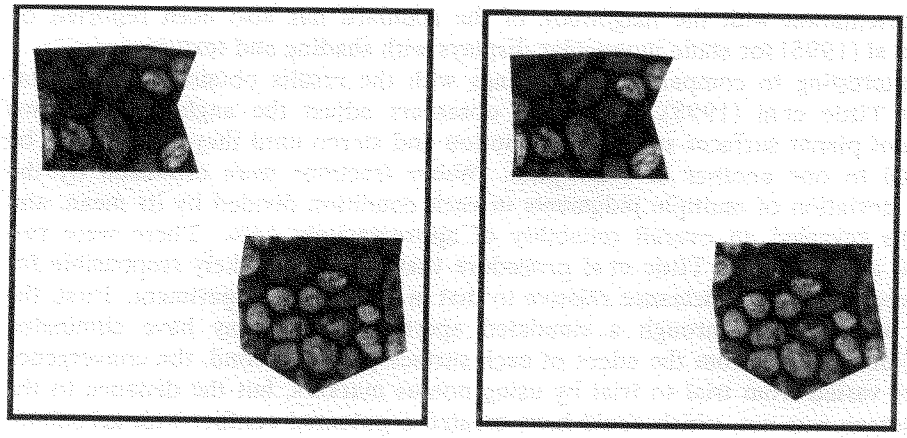

Figure 3. A stereogram of two spatially separated planar patches similar to those used in experiment 2.

3.2 *Results*

As in experiment 1, the best-fitting cumulative normal distributions were computed by using probit analysis, and the thresholds in each condition were estimated by halving the distance between the 25% and 75% points on the resulting psychometric functions. The individual results for all three observers are presented in table 2. Over all of the different conditions employed in this experiment, the average Weber fraction for detecting differences in relative orientation among the separated planar facets was approximately 11%. In comparing these results with those shown in table 1, it is clear that the relative orientation of separated planar facets is more difficult to judge than is the magnitude of a local dihedral angle.

It is important to keep in mind that in most natural scenes there are numerous different surface patches at all possible orientations, but that if two such patches were selected at random, the probability of their forming a connected dihedral angle would be relatively small. It is most likely to be the case, therefore, that the results obtained in the present experiment are more representative of natural vision than are the higher levels of performance for connected dihedral angles obtained in experiment 1.

Another relevant factor that needs to be considered in assessing the generality of these results is that surfaces in the environment are not always restricted to extended

Table 2. The Weber fractions for three observers for discriminating the relative orientations of separated planar patches in experiment 2.

Standard difference/°	Observer			Mean
	JN	HN	JT	
25	0.130	0.168	0.129	0.142
35	0.105	0.130	0.086	0.107
45	0.096	0.091	0.088	0.092
Mean	0.110	0.130	0.101	0.114

planar patches. Although smoothly curved surfaces could potentially be represented as a dense mesh of approximately planar local neighborhoods, as suggested by Marr and Nishihara (1978), there is little evidence to suggest that observers are capable of judging the relative orientations of these neighborhoods with any degree of accuracy. For example, in one recent study by Reichel et al (1995), observers were presented with shaded and textured images of smoothly curved surfaces, and were asked to indicate which of two probe regions was closest in orientation to a designated standard. Although all of the observers agreed that the depicted surfaces were perceptually quite compelling, their ability to judge differences in local orientation was surprisingly poor, as indicated by Weber fractions over 40%.

One possible factor that could have limited performance in the Reichel et al experiment is that the available information was restricted to shading and texture. In an effort to explore this possibility, experiment 3 was designed to measure observer sensitivity to the relative orientations of small local neighborhoods on smoothly curved surfaces under full cue conditions with available information from shading, texture, motion, and binocular disparity.

4 Experiment 3

4.1 *Method*

The apparatus and procedure were identical to those used in experiments 1 and 2, but the stimulus displays were quite different. The stimuli in this study were designed to simulate the optical projections of globally convex smoothly curved surfaces that resembled real-world objects, such as water-worn pebbles or potatoes—see figure 4 for a representative example. A set of twenty such objects were 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 Koenderink et al 1996; Norman et al 1995b).

The surfaces were covered with a blue-and-white random-check texture and were depicted with Lambertian shading (ie a surface that scatters light equally in all directions with no specular highlights). The simulations employed a single point light source oriented with a slant of 28° and a 45° tilt up and to the left of the observers' line of sight. Slant is defined in this context as the angle between the surface normal

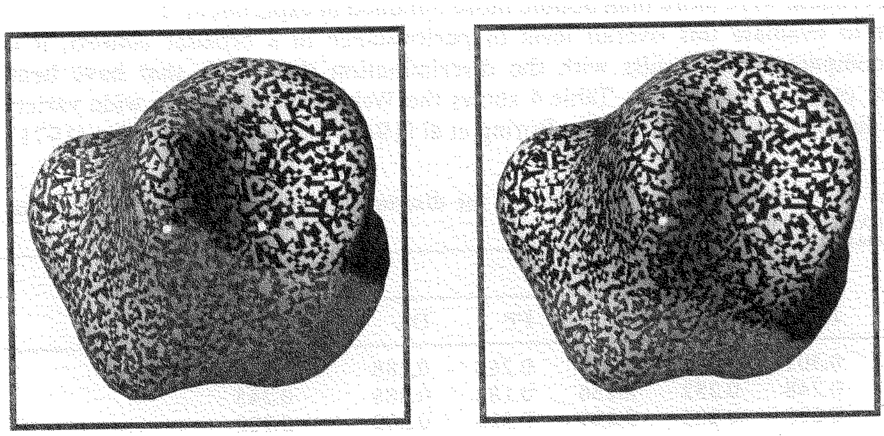

Figure 4. A stereogram of a randomly generated smoothly curved surface similar to those used in experiment 3.

and the line of sight, while tilt is the direction of the surface depth gradient within the frontoparallel plane. In generating the graphics displays, the surfaces were approximated by using a wire-frame mesh composed of 1280 triangular polygons. Texturing was achieved by using a trilinear mip-mapped interpolation algorithm, in which each polygon was first rotated to a frontoparallel 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 (see Todd and Mingolla 1983). The objects were presented stereoscopically and in motion, with exactly the same parameters as those used in experiments 1 and 2.

On any given trial one of the twenty 'potatoes' was shown, and a pair of surface regions was highlighted with red probe points. The surface was oriented so that both of these probe regions had a slant equal to half the test orientation difference. The tilt of one was selected at random over a 360° range, and the tilt of the other was in the opposite direction (ie both probe regions had identical slants, and tilts that differed by 180°). The observers' task was to decide whether the orientation difference between the two separated regions was greater or less than an implicit standard. The orientation difference in this case was defined as the arc cosine of the dot product between the two surface normals at those locations. Once again, there were three standard orientation differences, 25°, 35°, and 45°, and six test orientation differences that were greater than or less than the standard by 12%, 36%, and 60%. All other stimulus parameters, procedures, number of trials per condition, etc were the same as those used in the preceding two experiments. The displays were judged by the same three observers who participated in experiments 1 and 2, and two additional graduate student volunteers who were naive about the overall purpose of the experiment and how the displays were generated.

4.2 *Results*

As in experiments 1 and 2, the best-fitting cumulative normal distributions were computed by using probit analysis, and the thresholds in each condition were estimated by halving the distance between the 25% and 75% points on the resulting psychometric functions. The individual results for all five observers are presented in table 3. Note that the average Weber fraction over all conditions and observers was approximately 26%. By comparing these results with those shown in table 2, it is clear that the relative orientation between small local regions on a smoothly curved surface is much more difficult to judge perceptually than is the relative orientation of spatially separated planar patches. Indeed, the observers' Weber fractions in the present experiment were more than double those obtained in experiment 2.

In order to evaluate this overall level of performance in a broader context, it is useful to compare these results with the discrimination thresholds that have been reported for other sensory tasks. Table 4 shows the Weber fractions for a wide variety of sensory discriminations obtained by Boring et al (1948) and by Teghtsoonian (1971).

Table 3. The Weber fractions for five observers for discriminating the relative orientations on smoothly curved surfaces in experiment 3.

Standard difference/°	Observer					Mean
	JN	HN	JT	FP	DL	
25	0.301	0.357	0.317	0.283	0.288	0.309
35	0.248	0.252	0.256	0.183	0.288	0.245
45	0.258	0.283	0.199	0.216	0.202	0.232
Mean	0.269	0.297	0.257	0.227	0.259	0.262

With the exception of smell and taste, most other sensory discriminations produce Weber fractions that are an order of magnitude lower than the results obtained in the present experiment for discriminations of relative orientation.

Why is this task so difficult? Suppose, for example, that observers perform each judgment by first determining the absolute orientation at the designated probe points relative to the line of sight, and then computing a scalar product in order to obtain their relative orientation. One possible source of difficulty with this strategy is that the human visual system might be incapable of computing a scalar product between the surface normals in spatially separated regions, but the results of experiment 2 suggest this is not the case. Observers in that experiment were asked to judge the relative orientations of spatially separated planar patches, and the resulting Weber fractions were less than half of those obtained for curved surfaces in experiment 3.

A more likely source of the difficulty is an inability of the visual system to determine local orientation on curved surfaces—a hypothesis that is also supported by previous research involving local-orientation-adjustment tasks under full cue conditions (Koenderink et al 1996; Norman et al 1995b). These earlier studies have indicated, moreover, that our visual sensitivity to local orientation is anisotropic. Observers are typically quite accurate at adjusting the direction of the surface depth gradient within the frontoparallel plane (ie tilt), but their judgments of slant relative to the line of sight tend to have large standard deviations on the order of 25%.

There are other known anisotropies of perceived 3-D structure that could also have influenced the difficulty of this task relative to the judgments of dihedral angles in experiments 1 and 2. Gillam and her colleagues have shown that the slant of a stereoscopically presented surface about a vertical axis is more difficult to detect than an equivalent slant about a horizontal axis (see Gillam et al 1984, 1988), and there is a similar anisotropy in the detection of surface curvature for both motion and stereo (see Norman and Lappin 1992; Norman and Todd 1995; Rogers and Graham 1983; Tittle et al 1995). It is important to keep in mind that the probed surface regions in the present experiment had a wider range of possible orientations than did the planar patches of experiments 1 and 2. If there were significant anisotropies in the observers' perception of local orientation, then it is likely that the error rates for these judgments would increase with the overall range of orientations presented across trials.

Table 4. The Weber fractions for discriminating various sensory dimensions as reported by Boring et al (1948) and by Teghtsoonian (1971).

Sensory dimension	Boring et al (1948)	Teghtsoonian (1971)
Auditory pitch	0.003	
Electric shock		0.013
Deep pressure	0.013	
Visual brightness	0.016	0.079
Lifted weight	0.019	0.020
2-D line length		0.029
Loudness	0.088	0.048
Smell of rubber	0.104	
Taste of saline	0.200	0.083

5 Discussion

The research described in the present paper has been designed to investigate the abilities of human observers to discriminate the relative orientations of visible surfaces under full cue conditions with potential optical information from shading, texture, motion, and binocular disparity. The results demonstrate that observers' sensitivity can vary dramatically according to the particular type of structural configuration they are asked to judge. For planar patches that are connected to form a dihedral angle, the Weber fraction for discriminating relative orientation is approximately 8%. This increases to 11% if the patches are separated spatially, and it jumps to over 26% if the patches are part of a smoothly curved surface.

It is likely to be the case that there were two distinct sources of error that combined to produce these unusually high Weber fractions—especially for the smoothly curved surfaces employed in experiment 3. Consider, for example, an observer's judgments of perceived relative orientation for a single pair of small local surface regions over an arbitrarily large number of trials. If the observer's perception of 3-D form is systematically distorted relative to the actual structure of the depicted object, then there will be a *constant error* by which the mean of these judgments differs from its correct value. Similarly, there would also be a *random error* due to the variance of the observer's judgments about their central tendency. Although the random variation of display parameters across trials in the design of the present experiments does not make it possible to separate the relative contributions of these two components, they have both been well documented in previous investigations (eg see Koenderink et al 1996; Norman et al 1995a; Tittle et al 1995).

This particular design was chosen in an effort to estimate observer sensitivity to relative surface orientation under natural viewing conditions in which surfaces can appear in a wide variety of structural configurations, and can be perceptually specified by several different sources of optical information. There is, however, one potentially important source of variation among visible surfaces in the environment that was artificially constrained in the present experiments, and may therefore have increased the overall level of performance relative to what would otherwise be possible under more natural conditions: whereas real surfaces can appear at arbitrary viewing distances, the surfaces presented in the present experiments were all constrained to a fixed viewing distance of 76 cm.

During the past 100 years, there have been numerous investigations of how observers' perceptions of real objects in a natural environment are influenced by variations in viewing distance (see Norman et al 1995a for a review). The consistent pattern of results from these studies is that perceived intervals in depth become systematically compressed as their distance from the observer is increased. A similar finding has also been reported by Tittle et al (1995) for judgments of relative orientation for computer-simulated dihedral angles depicted with both motion and binocular disparity. It is reasonable to suspect therefore that observers' discrimination thresholds in the present experiments might have been even higher if the simulated viewing distance had varied from trial to trial.

Given the high Weber fractions obtained in these experiments for smoothly curved surfaces at a fixed viewing distance, we believe it is unlikely that local surface orientations could be a primitive component of our immediate phenomenal awareness of 3-D form, as hypothesized by Gibson (1950) and Marr (1982). Let us suppose for the sake of argument that this hypothesis is correct, and that the high Weber fractions were due primarily to random error. We might expect in that case that 25% random perturbations of perceived local orientation would cause smooth surfaces to appear wrinkled, but this is not consistent with the observers' phenomenal impressions that the surfaces appeared perfectly smooth (eg see figure 4). The other possibility to consider is that the high Weber fractions were due to constant error, such that the objects appeared systematically expanded or compressed along the line of sight, as has been reported in previous investigations (eg see Koenderink et al 1996; Norman et al 1995a; Tittle et al 1995). Note in that case that the orientation of an object relative to the axis of compression would have changed continuously as it rotated in depth, which should have been perceived as a nonrigid deformation. All of the observers reported, however, that the depicted motions appeared perfectly rigid.

Two fundamental characteristics of our perceptual awareness of environmental structure are that smooth surfaces appear phenomenally to be smooth, and that rigid objects appear phenomenally to be rigid as they move relative to the observer. Thus we would expect that whatever constitutes the basic representational units of our phenomenal awareness should exhibit these same properties. For a metrical representation of 3-D structure to satisfy these criteria it would have to be both accurate and precise, but the results of the present experiments, together with those from earlier investigations (Koenderink and van Doorn 1995; Koenderink et al 1992, 1994, 1995, 1996; Norman et al 1995b; Todd et al 1996), provide strong evidence that this is not the case for observers' knowledge of local surface orientation.

Another potential aspect of local surface structure identified by Gibson (1950) and Marr (1982) as a potential component of our perceptual representations is the property of relative depth. The available evidence suggests, however, that observers are no more accurate or precise at judgments of relative depth intervals than they are at judging relative orientation. The Weber fractions for discriminating 3-D length intervals are typically around 15% for objects presented at a fixed viewing distance (eg see McKee et al 1990; Norman et al 1995a), and this increases to over 25% when viewing distance is varied across trials (Norman et al 1995a). There are also systematic constant errors in observers' magnitude estimates due to the compression of perceived relative depth with increasing viewing distance (eg see Norman et al 1995; Tittle et al 1995).

How, then, might we account for observers' perceptions of surface smoothness or of rigid motion? From the arguments presented above, we speculate that these properties cannot be derived from a perceptual representation of local surface depth or orientation. The alternative possibility to consider is that they are determined directly from specific patterns of optical stimulation. Surface smoothness, we suspect, is defined by smooth variations of shading, texture, motion, or binocular disparity. The available information for the detection of rigid motion has been described by numerous investigators (eg Koenderink and van Doorn 1991; Todd 1982; Ullman 1977), and there is a growing amount of evidence to suggest that perceived rigidity does not necessarily depend on the invariance of perceived structure over time (Norman et al 1995a; Tittle et al 1995). A particularly compelling example of this phenomenon can easily be experienced while driving along a divided highway. Although the visual environment while driving usually appears quite rigid and stable, if one attends carefully to the hash marks on the road, they will appear to expand dramatically in length as they get closer and closer to the point of observation. Such findings suggest that metrical properties of 3-D structure such as relative depth or orientation may be of only secondary importance to our perceptual awareness of 3-D form.

Acknowledgement. This research was supported in part by AFOSR grant number F49620-93-0116.

References

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

Braunstein M L, 1976 *Depth Perception through Motion* (New York: Academic Press)

Braunstein M L, Liter J C, Tittle J S, 1993 "Recovering 3-D shape from perspective translations and orthographic rotations" *Journal of Experimental Psychology: Human Perception and Performance* **19** 598-614

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 visual surfaces" *American Journal of Psychology* **63** 367-384

Gibson J J, 1979 *The Ecological Approach to Visual Perception* (Boston, MA: Houghton Mifflin)

Gillam B, Chambers D, Russo T, 1988 "Postfusional latency in stereoscopic slant perception and the primitives of stereopsis" *Journal of Experimental Psychology: Human Perception and Performance* **36** 559-564

Gillam B, Flagg T, Finlay D, 1984 "Evidence for disparity change as the primary stimulus for stereoscopic processing" *Perception & Psychophysics* **36** 559-564

Koenderink J J, Doorn A J van, 1991 "Affine structure from motion" *Journal of the Optical Society of America A* **8** 377-385

Koenderink J J, Doorn A J van, 1995 "Relief: pictorial and otherwise" *Proceedings of the 5th British Machine Vision Conference* **24** 115-126

Koenderink J J, Doorn A J van, Kappers A M L, 1992 "Surface perception in pictures" *Perception & Psychophysics* **14** 163-175

Koenderink J J, Doorn A J van, Kappers A M L, 1994 "On so-called paradoxical monocular stereoscopy" *Perception* **23** 583-594

Koenderink J J, Doorn A J van, Kappers A M L, 1995 "Depth relief" *Perception* **24** 115-126

Koenderink J J, Kappers A M L, Todd J T, Norman J F, Phillips F, 1996 "Surface range and attitude probing in stereoscopically presented dynamic scenes" *Journal of Experimental Psychology: Human Perception and Performance* in press

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

Marr D, 1982 *Vision* (San Francisco, CA: W H 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, Series B* **200** 269-294

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, 1995 "The perception of 3-dimensional structure from contradictory optical patterns" *Perception & Psychophysics* in press

Norman J F, Todd J T, Perotti V J, Tittle J S, 1995a "The visual perception of 3D length" *Journal of Experimental Psychology: Human Perception and Performance* in press

Norman J F, Todd J T, Phillips F, 1995b "The perception of surface orientation from multiple sources of optical information" *Perception & Psychophysics* **52** 629-636

Purdy W C, 1958 *The Hypothesis of Psychophysical Correspondence in Space Perception* unpublished PhD thesis, Cornell University, Ithaca, NY

Reichel F D, Todd J T, Yilmaz E, 1995 "Visual discrimination of local surface depth and orientation" *Perception & Psychophysics* in press

Rogers B J, Graham M, 1983 "Anisotropies in the perception of three-dimensional surfaces" *Science* **221** 1409-1411

Teghtsoonian R, 1971 "On the exponents in Steven's law and the constant in Ekman's law" *Psychological Review* **78** 71-80

Tittle J S, Todd J T, Perotti V J, Norman J F, 1995 "The systematic distortion of perceived 3-D structure from motion and binocular stereopsis" *Journal of Experimental Psychology: Human Perception and Performance* **21** 663-678

Todd J T, 1982 "Visual information about rigid and nonrigid motion: A geometric analysis" *Journal of Experimental Psychology: Human Perception and Performance* **8** 238-251

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, Koenderink J J, Doorn A J van, Kappers A M L, 1996 "Effects of changing viewing conditions on the perceived structure of smoothly curved surfaces" *Journal of Experimental Psychology: Human Perception and Performance* in press

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

Ullman S, 1977 *The Interpretation of Visual Motion* PhD thesis, Massachusetts Institute of Technology, Boston, MA

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

© 1995 a Pion publication printed in Great Britain