

# Visual discrimination of local surface depth and orientation

FRANCENE D. REICHEL

*Brown University, Providence, Rhode Island*

JAMES T. TODD

*Ohio State University, Columbus, Ohio*

and

EMRE YILMAZ

*Harvard University, Cambridge, Massachusetts*

Many theoretical analyses of 3-dimensional form perception assume that visible surfaces in the environment are perceptually represented in terms of local mappings of metric depth and/or orientation. Although this approach is often taken for granted in the study of human vision, there have been relatively few attempts to demonstrate its psychological validity empirically. In an effort to shed new light on this issue, our research has been designed to investigate the accuracy with which observers can discriminate metric depth and orientation intervals on smoothly curved surfaces. Observers were presented with visual images of surfaces defined by shading and/or texture, on which two pairs of points were designated with small dots. In Experiment 1, their task was to identify which pair of points had a greater difference in depth; in Experiment 2 they were required to judge which pair had a greater difference in orientation. The Weber fractions obtained for these tasks were 10 to 100 times greater than those that have been reported for other types of sensory discrimination, indicating that the perception of metric structure from these displays is surprisingly coarse grained.

When we think about the structure of visual space, there is a natural inclination to accept as givens the familiar metrical properties of Euclidean geometry. That is to say, it is sometimes taken for granted that observers can perceive the distance between two points, that they can perceive the angle between two lines, and that they can compare distances in different directions as defined by the Pythagorean theorem. For some researchers, these ideas about the fundamental components of our perceptual experience are so ingrained that they are not even acknowledged as theoretical assumptions.

Until recently, the presumed metrical properties of objects in 3-dimensional (3-D) space has been a central theme of virtually all theoretical discussions of visual form perception. For example, in one of his earliest papers on this topic, Gibson (1950) suggested that the perceived shape of a visible surface is compounded from the more elementary properties of perceived distance and slant at each local region. A very similar approach was adopted many years later by Marr (1982; see also Marr & Nishihara, 1978). Marr coined the term *2<sup>1/2</sup>-D sketch* to describe this type of viewer-centered representation of local surface structure, and he identified its discovery as

one of the most exhilarating events in his career. Ironically, Gibson (1979) concluded at about the same time that thinking about form perception in terms of local depths and orientations was one of his biggest mistakes.

Gibson's change of opinion was based primarily on an extensive series of experiments that he and his colleagues had performed on the perception of optical slant from texture gradients (see Braunstein, 1976, for an excellent review). The results of this research showed clearly that observers' slant judgments for simple planar surfaces tend to be systematically underestimated, and that they can be highly unreliable as well. More recently, similar findings have been obtained for judgments of local metrical structure (e.g., relative depths or orientations) with the use of more complex scenes of smoothly curved surfaces, and for objects depicted with several different types of optical information such as shading (e.g., Erens, Kappers, & Koenderink, 1993; Koenderink, van Doorn, & Kappers, 1992; Mingolla & Todd, 1986), texture (e.g., Todd & Akerstrom, 1987), motion (e.g., Todd & Bressan, 1990; Todd & Norman, 1991), and binocular disparity (e.g., Foley, 1980; McKee, Levi, & Bowne, 1990).

One potential problem in interpreting these results is that there has been relatively little discussion of what should be considered as good (or bad) performance in the perceptual analysis of an object's 3-D structure (cf. Koenderink et al., 1992). Because of the wide variety of stimulus materials and response tasks that have been employed, it is difficult to compare results across different

---

This research was supported in part by AFOSR Grant F49620-93-1-0116 to J.T.T. Correspondence should be addressed to F. D. Reichel, Department of Cognitive and Linguistic Sciences, Brown University, Providence, RI 02912.

experiments. It is partially for this reason, we suspect, that the theoretical implications of this research have been largely ignored. Although there have been several recent attempts—especially in the area of object recognition—to identify alternative forms of perceptual representation that are more qualitative in nature (e.g., Biederman, 1987; Hoffman & Richards, 1984; Koenderink, 1984; Koenderink & van Doorn, 1976, 1980, 1982, 1991; Richards, Koenderink, & Hoffman, 1987; Todd & Bressan, 1990; Todd & Reichel, 1989), it is still a widely accepted assumption that the primary goal of 3-D form perception is to obtain accurate reconstructions of Euclidean metric structure.

In an effort to shed some new light on this issue, the research described in the present article was designed to measure observer sensitivity to the two basic geometric properties explicitly represented by the  $2\frac{1}{2}$ -D sketch—namely, local depth and orientation. Our goal was to measure the Weber fractions for observers' discrimination thresholds of these properties, which would provide a common currency for comparing the results with those obtained for other types of sensory discrimination. Previous investigations of isolated lines viewed stereoscopically (McKee et al., 1990) or in motion (Todd & Bressan, 1990) have found that observers' judgments of relative 3-D length or depth are surprisingly imprecise, producing Weber fractions as high as 30%. The research described in the present article was designed to extend these findings to the visual perception of monocular static images of smoothly curved surfaces depicted with both shading and texture.

The basic procedure employed in these studies is quite general and can easily be applied to measure observer sensitivity to any metrical property of an object's 3-D structure. On each trial, an observer is presented with a visual image of a smoothly curved surface on which two pairs of points are designated with small dots. The task is to identify which pair of points has a greater difference along some arbitrary dimension (e.g., depth, distance, orientation, curvature, etc.). Comparing performance on this task for different structural properties should make it possible to obtain an objective measure of their relative perceptual salience. In the present series of experiments, we employed this technique to measure the precision with which observers can discriminate intervals of metric depth or orientation on smoothly curved surfaces defined by shading and texture. Depth interval thresholds were measured in Experiment 1, and orientation interval thresholds were measured in Experiment 2.

## EXPERIMENT 1

### Method

**Subjects.** Two of the authors, F.R. and J.T., participated as observers.

**Apparatus.** Stimuli were generated with a Lex-90 graphics system and were displayed on a 19-in. (48.5-cm) color monitor with a spatial resolution of  $640 \times 512$  pixels. Each display consisted of a pair of horizontally aligned surface regions. Each surface was

presented within a  $16.5 \times 12$  cm region of the display screen with a resolution of  $320 \times 250$  pixels. Together, the two surface regions covered approximately one half of the area of the entire display screen and were always positioned centrally. Observers viewed the displays monocularly through a viewing hood at a distance of approximately 76 cm, such that each pixel spanned a visual angle of approximately  $2.4'$ .

**Stimuli.** On every trial, observers were presented with a pair of rectangular surface patches, each of which was taken from one of three smoothly curved surfaces defined by shading, contours, or shading and contours in combination. (A sample pair of experimental surfaces defined by shading and contours can be seen in Figure 1.) Additional variation in surface structure was achieved via presentation of two surfaces that were not bilaterally symmetric in mirror image, resulting in a total of five experimental surfaces. The surfaces were generated using a Cartesian coordinate system ( $x, y, z$ ), where  $x$  and  $y$  were aligned with the horizontal and vertical axes of the display screen and  $z$  was perpendicular to the screen in depth. The depth  $z$  at each point on each of the three experimental surfaces was defined initially via the following equations:

$$z = 100 \cos(\pi(\sqrt{x^2+y^2})/100), \quad (1)$$

$$z = 150 \cos(\pi(\sqrt{(x-150)^2+(y+60)^2})/150) + 400e^{-((x+150)^2+(y+60)^2)/1250\pi} + 200e^{-((x-50)^2+y^2)/800\pi}, \quad (2)$$

and

$$z = -75 \cos(\pi(\sqrt{(x+100)^2+y^2})/125) + 300e^{-((x-200)^2+(y-200)^2)/3200\pi}. \quad (3)$$

They were then rotated by  $40^\circ$  about a horizontal axis through the origin and displayed under parallel projection.

To produce patterns of image shading, the intensity of each picture element was determined with seven-bit precision via the following equation:

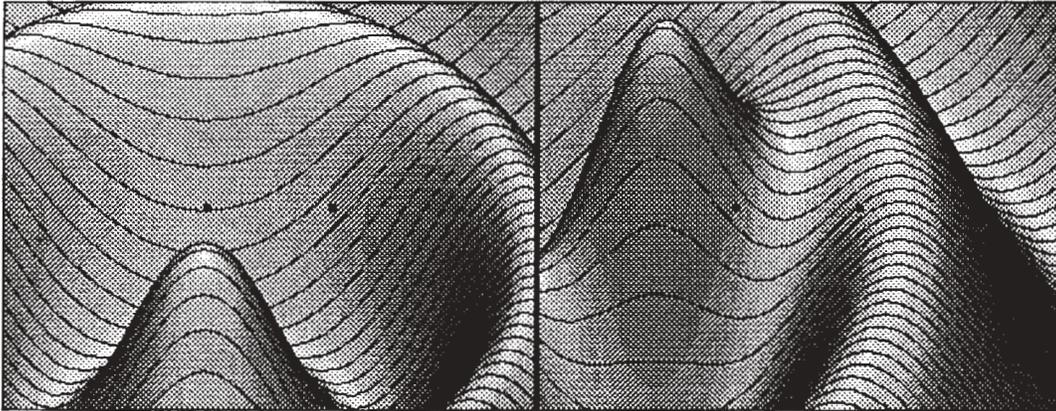
$$I = 128(\mathbf{L} \cdot \mathbf{N}),$$

where  $\mathbf{L}$  is a unit vector in the direction of the light source, and  $\mathbf{N}$  is a unit vector that is perpendicular to the surface at the depicted point (see Todd & Mingolla, 1983). This relation is a reasonable approximation of how light reflects from a pure matte (Lambertian) surface. The pattern of illumination in each display simulated an infinitely distant point light source at a  $20^\circ$  angle directly above the simulated direction of view.

The contours were defined mathematically from Equations 1, 2, and 3 by holding  $y$  constant at a fixed value of  $y_i$ . The pattern of contours was generated by using multiple values of  $y_i$  in 10-pixel increments. These mathematically defined contours were then rotated  $40^\circ$  with respect to the initial coordinate system in the plane of the display screen. They were presented as uniformly black contours on the shaded surfaces (Todd & Reichel, 1990).

Pairs of probe points were randomly selected from the experimental surfaces that were all separated by 55 pixels in the image plane. If the depth changed monotonically along a line on the surface connecting the two points, the point pair was categorized as *transitive* (see Todd & Reichel, 1989). Alternatively, if the depth changed nonmonotonically along this line, then the point pair was categorized as *intransitive*. Examples of these two categories are shown in Figure 1. The magnitude of depth difference for a given point pair ranged from 10 to 100 pixels in 2-pixel increments, which allowed us to investigate discrimination accuracy over a wide range of values.

**Procedure.** On every trial of the experiment, two surface patches were presented, each of which was marked with a pair of probe points. On half of the trials, both pairs were transitive, and on the



**Figure 1.** A sample pair of surface patches from Experiment 1 defined by shading and contours. The point pair marked in the left patch is transitive, whereas the one in the right patch is intransitive. The mixture of these categories as shown in the figure is for illustrative purposes only. In the actual experiment, the two point pairs in each display were either both transitive or both intransitive.

remaining trials, they were both intransitive. Observers were instructed to judge which of the two point pairs (left or right) was separated by a greater interval in depth. Responses were recorded by pressing one of two designated keys on a computer keyboard. On each subsequent trial, the difference between the two intervals was adjusted by using an adaptive staircase procedure adapted from Levitt (1970). Starting with an initial step size of 16 pixels, the difference between the two depth intervals was increased by one step after each incorrect response (i.e., making the task easier) and was decreased by one step after each pair of consecutive correct responses (i.e., making the task more difficult). The step size was decreased by half after the 1st, 3rd, and 7th reversals in the direction of this staircase, and the block was terminated after 10 reversals. This procedure converges at a threshold of 70.7%, where the probability of getting two correct responses in succession is .5. The average of the two designated depth intervals in each display was held constant within a given staircase, but was varied across blocks with possible values of 26, 32, 38, 44, 50, 56, and 62 pixels. Observers ran three staircases for each of these average depth intervals, and the threshold in each condition was estimated by the average of all 30 reversal points.

## Results

The data for both observers with all three types of displays (shading, contours, and shading with contours) are shown in Figure 2. For each average depth interval, a Weber fraction was computed via the following equation:

$$WF = 2 * \left( \frac{(|\text{Depth Difference 1} - \text{Depth Difference 2}|)}{(\text{Depth Difference 1} + \text{Depth Difference 2})} \right),$$

where Depth Difference 1 and Depth Difference 2 are the depth intervals separating the right and left point pairs, respectively. Note that the Weber fractions for this task ranged from approximately .2 to .6, which is remarkably high considering the experience of both observers (see Figure 2). In order for them to reliably determine which of two point pairs was separated by a greater interval in depth, the two intervals had to differ by at least 20% relative to their mean. 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 (see, e.g., Boring, Langfeld, & Weld, 1948; Teghtsoonian, 1971). With the exception of smell and taste, most other sensory discriminations produce Weber fractions that are an order of magnitude lower than those obtained in the present experiment for discriminations of relative depth.

As is evident from Figure 2, there were no significant differences in performance among the three different display conditions. The Weber fractions were generally comparable for the shaded and contoured displays, and there was no appreciable facilitation when both sources of information were presented in combination. There was, however, a clear effect of whether the point pairs were transitive or intransitive. For both observers in all three display conditions, the Weber fractions for transitive point pairs were consistently smaller than those obtained for the intransitive point pairs. Similar effects have also been reported with reaction time procedures for judgments of depth order (Todd & Reichel, 1989) and for depth interval discriminations (Reichel, 1991).

Given the high Weber fractions obtained for this task, it does not seem likely that the compelling perceptual appearance of these stimuli as smoothly curved surfaces could have been based on an explicit representation of relative depth in each local region. Experiment 2 was designed, therefore, to examine an alternative possibility that the stable appearance of 3-D form in these displays might be based primarily on a representation of relative orientation.

## EXPERIMENT 2

### Method

**Subjects.** Two of the authors, F.R. and E.Y., participated as observers.

**Apparatus.** Stimuli were generated with a Silicon Graphics Personal Iris 4D-35 workstation and were displayed on a 19-in. (48.5-cm) color monitor with a spatial resolution of  $1,280 \times 1,024$  pixels. Observers viewed the displays monocularly through a view-

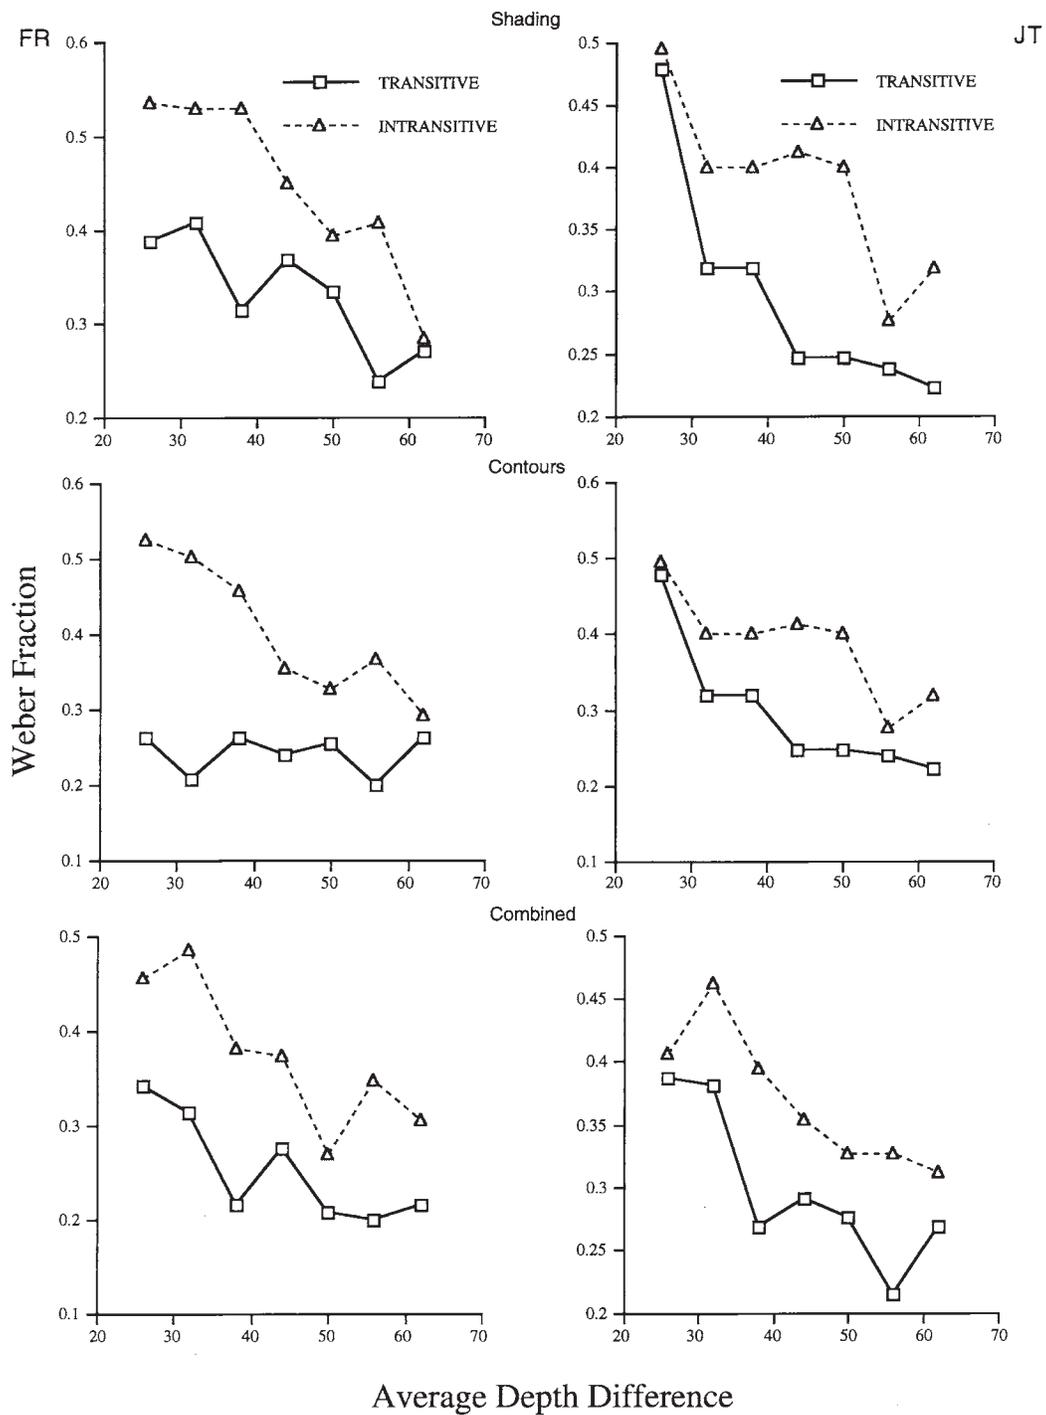
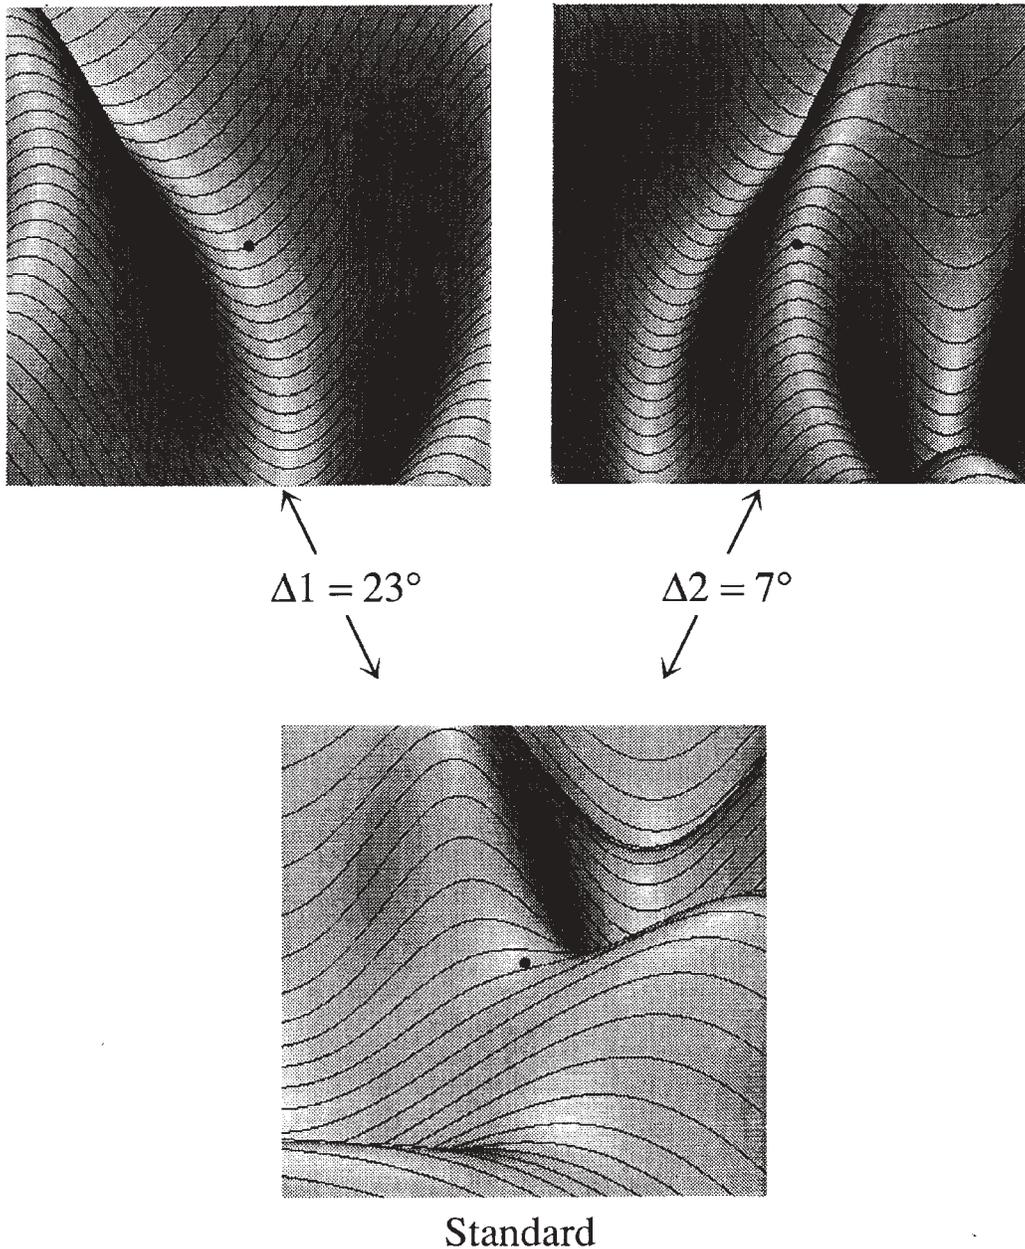


Figure 2. The individual Weber fractions for both observers in all of the different display conditions of Experiment 1.

ing hood at a distance of approximately 43 cm, such that each pixel spanned a visual angle of approximately 2.1°.

**Stimuli.** The displays were composed of rectangular surface patches, each of which was taken from one of three smoothly curved surfaces defined by shading alone or by both shading and contours in combination. These surfaces were generated in exactly the same manner as described in Experiment 1.

**Procedure.** Three different surface patches were presented on each trial in a triangular configuration as shown in Figure 3. Each patch was marked with a single probe point at its center, and observers were required to compare the relative surface orientations in the regions marked by those points. The one on the lower patch was designated as the “standard region,” and those in the two upper patches were designated as “test regions.” Observers were in-



**Figure 3.** A configuration of surface patches with shading and contours similar to those used in Experiment 2. Observers were required to judge which of the two test regions marked by points in the upper patches had a local orientation that was most similar to the surface orientation at a “standard” point marked in the lower patch. In this particular example, the two orientation intervals ( $\Delta 1$  and  $\Delta 2$ ) are  $23^\circ$  and  $7^\circ$ , respectively. The text in this figure is included for illustrative purposes only. In the actual experiment, the patches were presented without any accompanying symbols.

structed to judge which of these test regions (left or right) had a local orientation that was most similar to the orientation of the standard region. Responses were recorded by pressing one of two designated keys on a computer keyboard.

The difference between the orientation intervals of each test relative to the standard was varied from trial to trial by using the same

adaptive staircase procedure as was employed in Experiment 1. The average of these two intervals was held constant within each staircase, but was varied across blocks with possible values of  $15^\circ$ ,  $20^\circ$ ,  $25^\circ$ ,  $30^\circ$ ,  $35^\circ$ ,  $40^\circ$ ,  $45^\circ$ ,  $50^\circ$ , and  $55^\circ$ . Observers ran three staircases for each of these average orientation intervals, and the threshold in each condition was estimated by the average of all 30 reversal points.

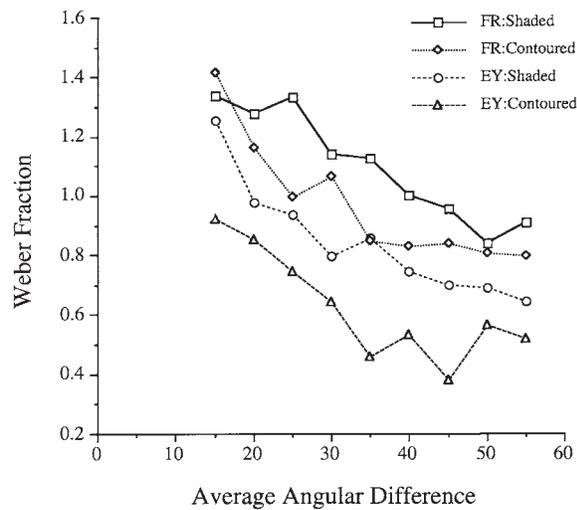


Figure 4. The individual Weber fractions for both observers in all of the different display conditions of Experiment 2.

## Results

The data for both observers with each type of display (shading and shading with contours) are shown in Figure 4. For each average orientation interval, a Weber fraction was computed with the following equation:

$$WF = 2 * \left( \frac{(|\text{Angle 1} - \text{Angle 2}|)}{(\text{Angle 1} + \text{Angle 2})} \right),$$

where Angle 1 is the difference in orientation between the left test region and the standard, and Angle 2 is the difference in orientation between the right test region and the standard. It is clear from the figure that the observers had great difficulty discriminating differences in relative orientation. The Weber fractions for this task ranged from approximately .3 to 1.5, which is roughly twice as large as those obtained in Experiment 1 for discriminations of relative depth.

One other aspect of the results that differs from those of Experiment 1 is that both observers were consistently less accurate when the surfaces were specified by shading alone than when shading and contours were presented in combination. Because we did not run a contour-alone condition, it is not possible to determine from these data whether contours are perceptually more informative than shading about relative surface orientation, or whether the effect could be due to a facilitative interaction that occurs when the two sources of information are presented in combination.

## DISCUSSION

The research described in the present article was designed to evaluate the psychological validity of a particular type of perceptual representation—sometimes referred to as the  $2\frac{1}{2}$ -D sketch (Marr, 1982; Marr & Nishihara, 1978)—which is the desired end product of many existing computational analyses of 3-D structure

from visual images. For smoothly curved surfaces, the  $2\frac{1}{2}$ -D sketch explicitly represents the local depth and orientation of each visible surface point. In considering the precision of this representation for actual human vision, Marr (1982) speculated that local orientations and relative depths are represented quite accurately, but that absolute depths from the point of observation may be represented only roughly.

Experiment 1 of the present series was designed specifically to measure the precision of observers' judgments of relative depth for surfaces defined by shading and/or texture. Two surface patches were presented on each trial, each of which contained a pair of probe points. The observers' task was to indicate which pair of probe points was separated by a greater distance in depth. If observers had accurate knowledge of relative surface depths, as suggested by Marr and Nishihara (1978) and Marr (1982), this task should be quite easy. All that would be necessary is to attend to the relative depth of each pair and to compare them directly to see which one is larger.

Although both observers agreed that the 3-D appearance of the depicted surfaces was extremely compelling, their ability to discriminate between different magnitudes of relative depths was surprisingly imprecise when compared with other types of sensory discriminations. Performance was highest when the point pairs were ordinarily transitive, but even in this most optimal condition, the Weber fractions ranged from 20% to 40%. 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. For judgments of low-level visual, auditory, or haptic properties, such as brightness, 2-D line length, pitch, or lifted weight, observers' Weber fractions are typically less than 5% (see, e.g., Boring et al., 1948; Teghtsoonian, 1971), almost an order of magnitude lower than those obtained in the present experiment.

Why should comparisons of relative depth be so difficult? One possibility is that the available information in our displays was inadequate to specify this particular aspect of surface structure, and that performance might be expected to improve with the addition of more powerful sources of information such as motion or binocular disparity. Although this is a reasonable hypothesis, a growing amount of evidence suggests that judgments of metric depth intervals are inherently difficult regardless of the particular type of optical information by which they are specified. For example, McKee, Levi, and Bowne (1990) obtained Weber fractions of 15% for judged stereoscopic depth intervals at a fixed viewing distance. Norman, Todd, Perotti, and Tittle (in press) obtained similar results when objects at a fixed viewing distance were presented with motion and stereo combined, though the Weber fractions increased to around 25% when viewing distance was varied across trials. There is also confirming evidence to support these findings for judgments of 3-D length from orthographically projected motion and from the variability obtained over repeated observations

on 3-D adjustment tasks under full cue conditions (see, e.g., Norman et al., in press; Tittle, Todd, Perotti, & Norman, 1995).

In view of the imprecision of observers' judgments of relative depth, Experiment 2 was designed to investigate whether judgments might be more accurate for some other aspect of local surface structure. It is important to keep in mind, when one is considering this issue, that most current computational analyses for determining shape from shading or contours are designed to reconstruct the local orientation of each visible surface region (see, e.g., Horn & Brooks, 1989, for a recent review), rather than a depth map. Thus, to the extent that these analyses are relevant to human perception, it might be reasonable to expect that local orientation could be a more salient aspect of surface structure for these particular sources of optical information. The results of Experiment 2 revealed, however, that observers are even less precise at judging relative orientation in these displays than they are at judging relative depth—that is, the Weber fractions ranged from 40% to 140%.

One important factor that needs to be considered in comparing performance between Experiments 1 and 2 is that depth is a scalar quantity whereas orientation is a vector. Consider, for example, two possible perceptual representations—one that encodes the absolute depth of each visible surface region relative to the point of observation, and another that encodes its local orientation relative to the line of sight. In order to determine the relative depth between two probed locations, all that would be necessary is a simple subtraction of their respective absolute depths, but to determine the difference between two local orientations, it would be necessary to compute their scalar product. It is likely to be the case, therefore, that there are two distinct sources of error in observers' judgments of relative orientation, one that arises from the representation of absolute orientation in each local region, and another that arises from computing the difference between them.

Although it is not possible to separate these components in the design of the present experiment, several recent studies employing matching tasks can be used to estimate the precision of perceived orientation at individual probe regions on a smoothly curved surface. This technique has been developed primarily by Koenderink, van Doorn, and Kappers (1992, 1994, 1995; see also Koenderink, Kappers, Todd, Norman, & Phillips, in press; Koenderink & van Doorn, 1995; Norman, Todd, & Phillips, 1995; Todd, Koenderink, van Doorn, & Kappers, in press). The observer's task in these studies was to adjust the 3-D orientation of a circular gauge figure until it appeared to rest in the tangent plane at a designated probe point on an object's surface. In such a task, 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. One interesting finding that has emerged from these experiments is that observers' judgments of local orientation at

a designated probe point are no less precise than their judgments of local depth using a stereoscopic matching procedure (see, e.g., Koenderink et al., in press). This finding suggests that the differences obtained in the present experiments between judgments of depth and orientation intervals may indeed be due to the difficulty of computing a scalar product from the individual surface normals. An interesting finding from the gauge figure experiments of Koenderink and his colleagues is that observer sensitivity to local orientation is anisotropic with respect to the components of slant and tilt. Slant is defined in this context as the angle between the surface normal and the line of sight, and tilt is the direction of the surface depth gradient within the frontoparallel plane. Observers are typically quite accurate at adjusting the direction of tilt. Their estimates of slant, on the other hand, tend to have large standard deviations (on the order of 25%) and large constant errors that can vary dramatically among different observers or for a given observer in different viewing conditions (e.g., when an object is viewed with different directions of illumination).

Although most of our discussion thus far has focused on local properties such as depth or orientation that are defined at a given point, other aspects of the data suggest that the observers' judgments were also affected by more global aspects of surface structure. In Experiment 1, for example, the discrimination accuracy was significantly improved when the intervening surface regions along each of the designated intervals varied monotonically in depth, in such a way that they were ordinally transitive (see Todd & Reichel, 1989). A similar result has also been reported by Reichel (1991), who showed that reaction times for discriminating ordinally transitive depth intervals are significantly faster than those obtained for comparable intervals that are ordinally intransitive. Such findings provide strong evidence that the process of determining the magnitude of perceived depth between two designated points involves some form of interpolation over the intervening surface regions.

There are many possible forms of representation that could potentially be used to describe an object's perceived shape. Marr (1982) speculated that our immediate conscious awareness of environmental structure is based primarily on an explicit representation of local depths and orientations, which he referred to as the  $2\frac{1}{2}$ -D sketch, but a growing body of evidence suggests that our perceptual knowledge of these properties is surprisingly coarse grained. It is important to keep in mind when one is evaluating these results that many other types of perceptual judgments do not require an accurate or precise knowledge of the metrical aspects of an object's structure. Observers typically do quite well on these tasks. For example, Reichel (1991) has shown that judgments of depth order are performed much faster and with greater accuracy than are comparable judgments about the magnitude of a depth interval. High levels of performance are also common for tasks involving basic-level object recognition (see, e.g., Biederman, 1987). Such findings suggest to us that the perceptual appear-

ance of objects is more likely to be based on some form of nonmetrical representation, such as one that defines 3-D structure as a mapping of local depth order relations (see, e.g., Todd & Reichel, 1989), or as a connected set of qualitatively distinct parts (see, e.g., Biederman, 1987; Hoffman & Richards, 1984; Richards, Koenderink, & Hoffman, 1987).

#### REFERENCES

- BIEDERMAN, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, **94**, 115-147.
- BORING, E. G., LANGFELD, H. S., & WELD, H. P. (1948). *Foundations of psychology*. New York: Wiley.
- BRAUNSTEIN, M. L. (1976). *Depth perception through motion*. New York: Academic Press.
- ERENS, R. G. F., KAPPERS, A. M. L., & KOENDERINK, J. J. (1993). Estimating local shape from shading in the presence of local shading. *Perception & Psychophysics*, **54**, 334-342.
- FOLEY, J. (1980). Binocular distance perception. *Psychological Review*, **87**, 411-434.
- 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: Houghton Mifflin.
- HOFFMAN, D. D., & RICHARDS, W. A. (1984). Parts of recognition. *Cognition*, **18**, 65-96.
- HORN, B. K. P., & BROOKS, M. J. (1989). *Shape from shading*. Cambridge, MA: MIT Press.
- KOENDERINK, J. J. (1984). What does the occluding contour tell us about solid shape? *Perception*, **13**, 321-330.
- KOENDERINK, J. J., KAPPERS, A. M. L., TODD, J. T., NORMAN, J. F., & PHILLIPS, F. (in press). Surface range and attitude probing in stereoscopically presented dynamic scenes. *Journal of Experimental Psychology: Human Perception & Performance*.
- KOENDERINK, J. J., & VAN DOORN, A. J. (1976). The singularities of the visual mapping. *Biological Cybernetics*, **24**, 51-59.
- KOENDERINK, J. J., & VAN DOORN, A. J. (1980). Photometric invariants related to solid shape. *Optica Acta*, **27**, 981-986.
- KOENDERINK, J. J., & VAN DOORN, A. J. (1982). The shape of smooth objects and the way contours end. *Perception*, **11**, 129-137.
- KOENDERINK, J. J., & VAN DOORN, A. J. (1991). Affine structure from motion. *Journal of the Optical Society of America A*, **8**, 377-385.
- KOENDERINK, J. J., & VAN DOORN, A. J. (1995). Relief: Pictorial and otherwise. *Image & Vision Computing*, **13**, 321-334.
- KOENDERINK, J. J., VAN DOORN, A. J., & KAPPERS, A. M. L. (1992). Surface perception in pictures. *Perception & Psychophysics*, **52**, 487-496.
- KOENDERINK, J. J., VAN DOORN, A. J., & KAPPERS, A. M. L. (1994). On so-called "paradoxical monocular stereoscopy." *Perception*, **23**, 583-594.
- KOENDERINK, J. J., VAN DOORN, A. J., & KAPPERS, A. M. L. (1995). Depth relief. *Perception*, **24**, 115-126.
- LEVITT, H. (1970). Transformed up-down methods in psychoacoustics. *Journal of the Acoustical Society of America*, **49**, 467-477.
- MARR, D. (1982). *Vision*. San Francisco: 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.
- MCKEE, S. P., LEVI, D. M., & BOWNE, S. F. (1990). The imprecision of stereopsis. *Vision Research*, **30**, 1763-1779.
- MINGOLLA, E., & TODD, J. T. (1986). Perception of solid shape from shading. *Biological Cybernetics*, **53**, 137-151.
- NORMAN, J. F., TODD, J. T., PEROTTI, V. J., & TITTLE, J. S. (in press). The visual perception of 3D length. *Journal of Experimental Psychology: Human Perception & Performance*.
- 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. (1991). *Visual knowledge of smoothly curved surfaces*. Unpublished doctoral dissertation, Brandeis University, Waltham, MA.
- RICHARDS, W. A., KOENDERINK, J. J., & HOFFMAN, D. D. (1987). Inferring three-dimensional shapes from two-dimensional silhouettes. *Journal of the Optical Society of America A*, **4**, 1168-1175.
- 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 & Performance*, **21**, 663-678.
- TODD, J. T., & AKERSTROM, R. A. (1987). Perception of 3-dimensional form from patterns of optical texture. *Journal of Experimental Psychology: Human Perception & Performance*, **11**, 689-710.
- TODD, J. T., & BRESSAN, P. (1990). The perception of three-dimensional affine structure from minimal apparent motion sequences. *Perception & Psychophysics*, **48**, 419-430.
- TODD, J. T., KOENDERINK, J. J., VAN DOORN, A. J., & KAPPERS, A. M. L. (in press). Effects of changing viewing conditions on the perceived structure of smoothly curved surfaces. *Journal of Experimental Psychology: Human Perception & Performance*.
- TODD, J. T., & MINGOLLA, E. (1983). Perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental Psychology: Human Perception & Performance*, **9**, 583-595.
- TODD, J. T., & NORMAN, J. F. (1991). The visual perception of smoothly curved surfaces from minimal apparent motion sequences. *Perception & Psychophysics*, **50**, 509-523.
- TODD, J. T., & REICHEL, F. D. (1989). Ordinal structure in the perception and cognition of smoothly curved surfaces. *Psychological Review*, **96**, 643-657.
- TODD, J. T., & REICHEL, F. D. (1990). The visual perception of smoothly curved surfaces from double projected contour patterns. *Journal of Experimental Psychology: Human Perception & Performance*, **16**, 665-674.

(Manuscript received August 17, 1994;  
revision accepted for publication April 7, 1995.)