

Perception of Local Three-Dimensional Shape

Flip Phillips and James T. Todd
The Ohio State University

The authors present a series of 4 experiments designed to test the ability to perceive local shape information. Observers were presented with various smoothly varying 3-dimensional surfaces where they reported shape index and sign of Gaussian curvature at several probe locations. Results show that observers are poor at making judgments based on these local measures, especially when the region surrounding the local point is restricted or manipulated to make it noncoherent. Shape index judgments required at least 2° of context surrounding the probe location, and performance on sign of Gaussian curvature judgments deteriorated as the contextual information was restricted as well.

In informal conversation we often refer to the configuration of an object's surfaces as its form or shape. A basketball, a chair, a human body, or a cloud, each have a distinct structure that allows us to differentiate between them. Shape is typically distinct from the substance of which the object is made, but often times, the substance itself acts as a determinant of the formal structure of the resulting object. We do not have to necessarily know what substance an object is made of to identify it. Indeed, we can go even further and reliably infer shape from line drawings or photographs of an object. Caricatures of people and objects still look like the things they represent, even though the relationships and configurations of their constituent pieces may vary wildly from the original to the caricature. When we see an illustration of a loaf of bread, we have no trouble understanding that it is a representation of a loaf-sized three-dimensional solid and not an image of a small two-dimensional object, even though the representation may only be a two-inch black on white line drawing.

We often refer to shapes analogically, such as a highway exchange being a "cloverleaf," a person's figure as being "pear shaped," or a golf hole as being a "dog leg." When we do this, we are describing some global, intuitive, and informal characteristic of that object. There is seemingly something atomic or primitive about these types of shapes. We can treat them as a sort of verbal lowest common denominator. However, these primitive descriptions provide us only with nominal distinctions between different shapes.

Given this, what might constitute a good representational strategy for solid shape? One approach that has received a

large amount of attention is based on the occlusion contours of an object or its parts. These approaches rely on the interpretation of the projected two-dimensional sharp edges and outlines of an object, often times depending on their invariance or nonaccidental properties. In one such approach, Biederman's theory of recognition by components (1987), the nonaccidental qualities of the shape profiles and their contained edges allow each to be uniquely identifiable, thus, yielding a small set of primitive parts from which more complex objects can be created. Other similar approaches include Koenderink (1984a), Blum (1973), Hoffman and Richards (1985), and Richards, Koenderink, and Hoffman (1987).

Clearly, there is more to an object's shape than is defined by its edges. The interior of the object also consists of additional information in the form of smooth surface variations or textures. For example, rumpled cloth and natural landscapes defy the occlusion models because they are not easily described by a set of convex occlusion contours. Therefore, the internal smooth structure of objects require some other form of representation.

Gibson (1950) suggested that a local surface region can be described by a small set of phenomenal primitives, such as depth and orientation relative to the observer, and their gradients across the surface and scene. Later, Marr and Nishihara (1978) adopted this idea in their 2½ D sketch. This type of representation was a natural glue for their model because it derives itself from an image that does not contain any information about surfaces behind occlusions, as does the retinal image.

The Marr and Nishihara (1978) model uses depths and orientations relative to some observer as the regional quanta. It is important to note that if the observer or the object under scrutiny moves, an entirely new set of relative depths and orientations results. Using these primitives would result in a rather cumbersome representational strategy because it would require representations for all possible orientations of an object to recognize it.

A more flexible set of primitives would be invariant over the orientation of the object, resulting in a more economical set of representations. One primitive that meets this criteria is the intrinsic curvature of a surface, because it does not

Flip Phillips and James T. Todd, Department of Psychology, Vision Laboratories, The Ohio State University.

This research was supported by Air Force Office Scientific Research Grant F49620-93-1-0116. We thank J. Farley Norman for his contributions at the earlier stages of this research, as well as the other members of the Vision Laboratories for their criticism and advice.

Correspondence concerning this article should be addressed to Flip Phillips, Department of Psychology, Vision Laboratories, Room 314, Lazenby Hall, The Ohio State University, Columbus, Ohio 43210-1222. Electronic mail may be sent via Internet to flip@tv1.psy.ohio-state.edu.

vary relative to the observer. To better understand what we mean by curvature, we can examine some of the concepts of differential geometry and their applicability to our psychophysical problem.

For any surface there exists a maximum and minimum normal curvature at any given point, which are called the *principal curvatures* (typically referred to as κ_1 and κ_2) and the corresponding directions of these curvatures are called the *principal directions*.

One special feature of the principal directions is that they are always orthogonal. In Figure 1 we can see the two principal directions labeled as C_1 and C_2 along with a cross section of the patch showing C_1 and its curvature κ_1 .

There are several convenient and intuitive ways to combine these two curvatures into a higher order measure. One potentially useful metric is the *shape operator* of modern differential geometry, also known as the Second Fundamental Form, $\Pi(\bullet, \bullet)$, which measures how a surface bends in space by estimating how the surface normal changes from point to point. Also useful is the surface's *normal curvature*, which is the curvature of a normal section of a surface. It is a real-valued function similar to the shape operator, with the added benefit being that it defines the curvature in a given direction (Gray, 1993; Koenderink, 1989; Lipshutz, 1969). These curvatures range from negative or concave values to positive or convex values. As we will see, we can make this class of measurement more intuitive and perhaps more useful for experimentation purposes.

As Equation 1 shows, the product of the two principal curvatures yields the *Gaussian curvature* at point p ,

$$K_p = \kappa_1 \cdot \kappa_2. \quad (1)$$

The sign of the Gaussian curvature tells us if a surface's principal directions are curving in the same direction (positive, as in a bump or a dimple) or in different directions (negative, like a saddle surface); see Hilbert and Cohn-Vossen, 1952. The Gaussian curvature is also available as the determinant of the matrix yielded by the shape operator (see Figure 2).

Koenderink (1989) has taken these measures one step

further by defining an intuitive, higher order single-valued function that represents the scale invariant local structure along with a companion function, which contains the scale-sensitive information. These are his *shape index* and *curvedness*, respectively. Both scales are one dimensional. Their inputs (principal curvatures) are defined locally and, therefore, so are the resultant operators.

The shape index ranges on the interval $[-1, 1]$ with critical points occurring at the endpoints, ± 0.5 , and 0 (see Figure 3). Note that this yields a potentially nominal scale based on the two principal directions. The critical points conveniently separate the scale into four distinct regions.

For example, for the 0.0 shape index case, one principal direction of curvature is positive and the other negative, yielding a saddle. The two curvatures are said to be anticlastic. The resulting sign of the Gaussian curvature at this point will be negative because the two curvatures have opposite signs. In the +1.0 case, both principal directions have positive curvature yielding a bump. These curvatures are said to be synclastic. The Gaussian curvature in this case is positive because the curvatures are of the same sign. Finally, in the ± 0.5 cases, one principal curvature is either positive or negative, whereas there is zero curvature in the other direction. These curvatures are monoclastic and the resulting Gaussian curvature is zero because one of the curvatures is zero.

The shape index is insensitive to scale, and it only discriminates on the basis of directions of curvatures. Curvedness discriminates on the basis of the magnitudes of these curvatures with its value ranging from zero, for a planar surface, to infinity, for a sharp point. In Figure 4 we show a series of patches in which the shape index is a constant 1.0, whereas the curvedness ranges from 0.0 to 5.0.

We can consolidate our illustration of these two measures into one graph by plotting them as a polar function of the two principal curvatures (see Figure 5).

In this illustration, the shape index increases as a function of the angle from -1.0 in the positive-positive quadrant to 1.0 in the negative-negative quadrant, and the curvedness

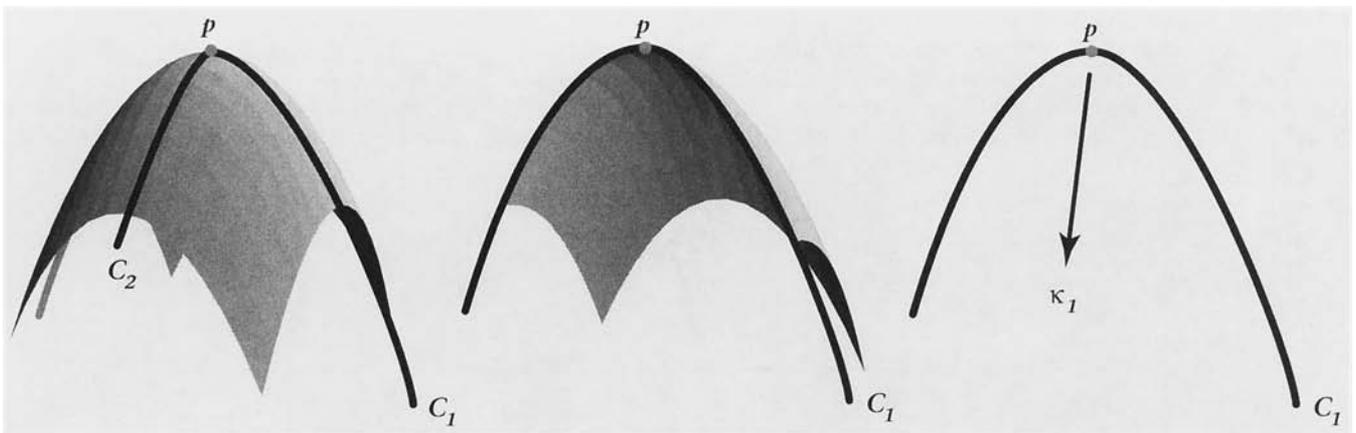


Figure 1. Principal directions and curvatures at point p .

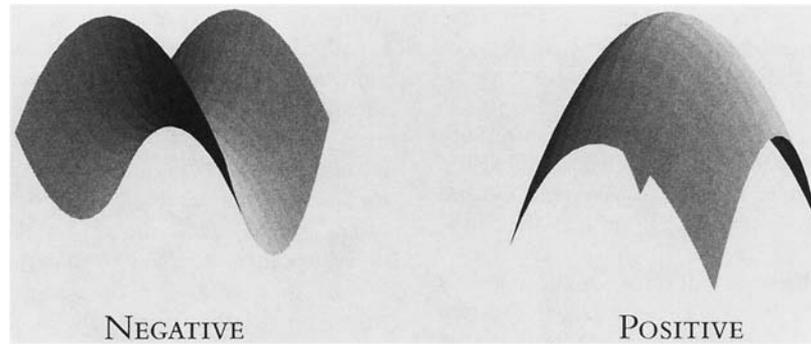


Figure 2. Gaussian curvature.

increases from 0.0 at the origin, increasing toward infinity along the radial axis.

In more formal terms, Koenderink's shape index (1989) is defined as a function of the principal curvatures κ_1 and κ_2 as

$$S(\kappa_1, \kappa_2) = -\frac{2}{\pi} \operatorname{atan} \left(\frac{\kappa_1 + \kappa_2}{\kappa_1 - \kappa_2} \right), \quad (2)$$

where $\kappa_1 > \kappa_2$. The scale sensitive function curvedness is then defined as

$$C(\kappa_1, \kappa_2) = \sqrt{\frac{\kappa_1^2 + \kappa_2^2}{2}}. \quad (3)$$

Regardless of whatever shape is, we know that a large sphere and a small sphere are of the same shape. Even more complicated shapes retain this scale insensitivity, as in the dog-leg example above. This isomorphic quality is true, not only of scale, but of orientation and the local configuration of the objects as well. However, the nominal information provides us with too little information. When we say an object is a ball, we might mean a baseball or a football. Whereas both of these shapes are of the same genus, they are of different shapes. Would it be possible for us to use some sort of measure that preserves our more intuitive idea of shape and also provides us with something more than just a nominal scale?

Previous studies in shape perception have utilized the shape index in several different contexts; shape-from-stereo (de Vries, 1993), shape-from-motion-parallax (van Damme, Oosterhoff, & van de Grind, 1994; van Damme & van de Grind, 1993), and shape-from-shading (Erens, 1993a;

1993b). These studies showed that the shape index and curvedness were appropriate and reliable for use in psychophysical experimentation. In each of these experiments, observers were presented with smooth quadric surfaces of the form $z(x, y) = \frac{1}{2} (\kappa_1 x^2 + \kappa_2 y^2)$. This yields surfaces in which shape index and curvedness varied smoothly over the extent of the object.

In an extension of these experiments, our experiments were designed to measure an observer's ability to judge shape index and Gaussian curvature utilizing more irregular stimuli to better simulate objects in the observer's natural world. Our question becomes: Is it possible for the observers to successfully judge local curvatures, and in so doing, categorize its local shape?

It is obvious that we need some amount of coherent area surrounding a purely local probe point to estimate the local shape, but how much of it do we actually need? Can we modify the statistical nature of this context and still obtain accurate estimations of the shape? In the four experiments that follow, we attempt to uncover this, along with what other types of information observers are utilizing to judge solid shape.

In these experiments, we investigate the ability of viewers to discriminate two local properties of a surface, Gaussian curvature and shape index, when the context of the local information is restricted or varied through the use of turbulent self-similar surfaces.

Experiment 1

The initial experiment was an attempt to replicate the results of prior studies (de Vries, 1993; Erens, 1993a; van

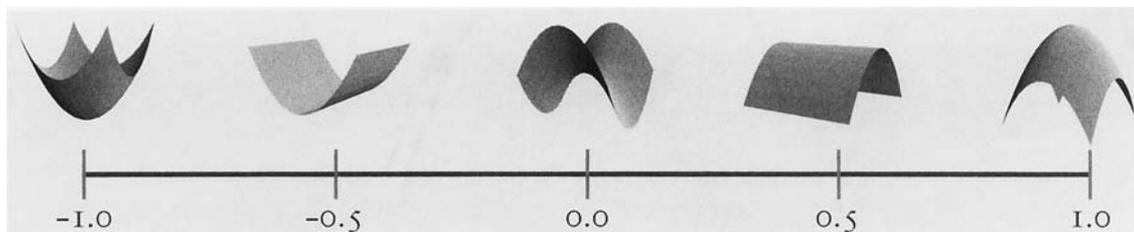


Figure 3. Shape index critical points.

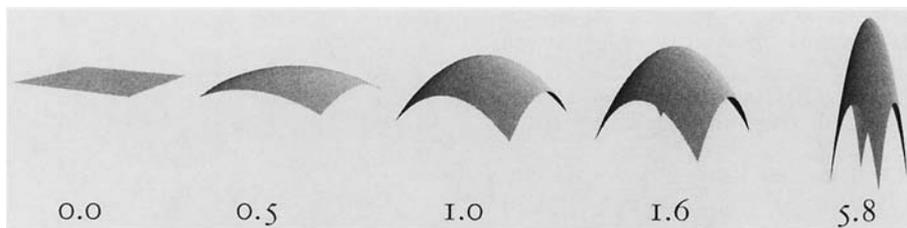


Figure 4. Magnitude of curvedness.

Damme et al., 1994; van Damme & van de Grind, 1993). In addition, we also chose to vary the amount of contextual information available in an attempt to observe its effect on the resulting judgments.

Method

Apparatus. All stimuli were computer-generated surfaces, displayed on a Silicon Graphics Crimson workstation. In all presentations, the surfaces were shown with texture, shading, motion, and in stereo. Viewing distance was a constant 57.4 cm and the entire display screen subtended an angle of 34° horizontally and 28° vertically. The experiments took place in a darkened room with a chin rest used to maintain a constant viewing distance. Responses were provided using the mouse buttons.

Stimuli. The surfaces used as stimuli in this experiment were simple quadrics defined by the Monge formulation,

$$z(x,y) = \frac{1}{2}(\kappa_1x^2 + \kappa_2y^2), \tag{4}$$

where κ_1 and κ_2 are the principal curvatures of the surface under examination (see Figure 6).

These test patches were a circular portion of the surface defined in Equation 4, centered about the origin. To reduce possible contour cues from the edge of the patch, they were smoothly

blended into a flat plane. To simplify the blending, the surface patch defined in Equation 4 can be represented in polar form as shown in Equation 5 (note that r , not θ , corresponds to visual angle). Over the first 2° of the stimulus, the surface is defined entirely by this equation. After 3°, the stimulus is flat and is, z therefore, defined as $z_{flat}(r, \theta) = 0$.

$$z(r,\theta) = \frac{1}{2}(\kappa_1r^2\cos^2\theta + \kappa_2r^2\sin^2\theta). \tag{5}$$

Between 2° and 3° the surface patch was blended with the flat patch. This was accomplished by scaling the surface patch as a function of the radius, r , using a Hermite interpolation shown in Equation 6 (Rogers & Adams, 1990),

$$h(r) = 3r^2 - 2r^3. \tag{6}$$

This function is defined on the range [0,1]. Therefore, the piecewise linear interpolation function shown in Equation 7 is used to bring the radius onto that interval.

$$t(r,min,max) = \begin{cases} 0 & r < min \\ \frac{r - min}{max - min} & min \leq r \leq max. \\ 1 & r > max \end{cases} \tag{7}$$

The resulting surface equation for the 2° to 3° interpolation range in these experiments is, therefore,

$$z'(r,\theta) = z(r,\theta) \cdot \{1 - h[t(r,2,3)]\}. \tag{8}$$

The stimulus patches were represented computationally as polygonal meshes. To eliminate possible orientation artifacts or cues, the shape patches were arbitrarily rotated within the embedded surface, and the edges of the patch were randomly perturbed for each trial. The probe point was always at the center of the patch, at a constant distance from the viewer, and was indicated by a small blue dot.

A simple, untextured, example surface of the type used in this

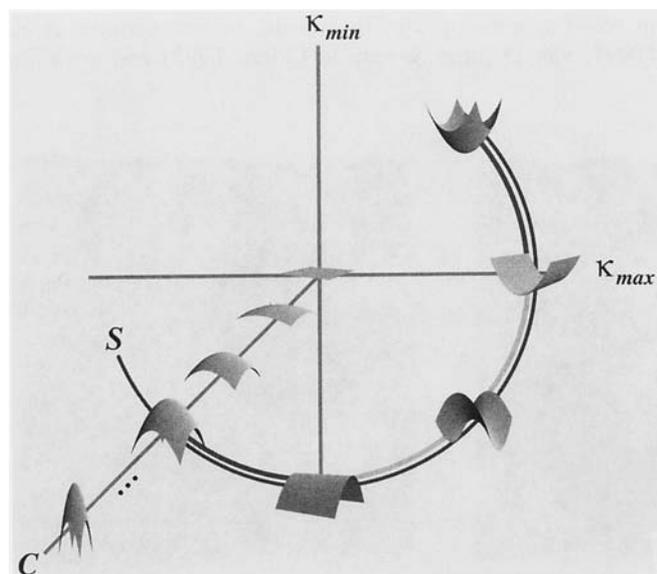


Figure 5. The big picture of shape (S) index and curvedness (C).

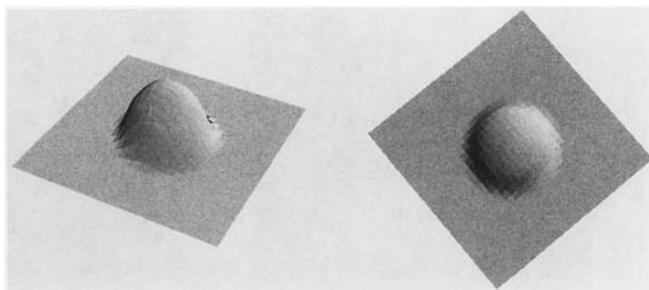


Figure 6. Example surfaces for shape index of 1.0.

experiment is shown in Figure 6, and a stereogram of an actual stimulus with a gravel-like texture is shown in Figure 7. (The actual stimuli had a colored texture.)

Observers. The displays were presented to four observers, all members of the research group that performed the experiments. All observers had normal or corrected-to-normal vision.

Procedure. On each trial, the observer was presented with a pair of three-dimensional surface stimuli. The task on each trial was to indicate the stimulus containing a predefined standard by pressing one of two mouse buttons. Observers were permitted to view the stimuli as long as necessary to make their decision. Feedback was provided in the form of a beep for correct answers.

A 5×3 , shape index by patch radius, within-subjects factorial design was used. The shape index was probed at the five critical locations: -1.0 , -0.5 , 0.0 , 0.5 , and 1.0 (see Figure 3), while curvedness was fixed at a constant 0.5 cm^{-1} . The target patches varied in radius between 1.0, 1.5, and 2.0 cm, subtending visual angles of 2° , 3° , and 4° , respectively. Because the curvedness was held constant, the modifications of the patch radius only changed the amount of the patch seen. It did not change the overall scale of the feature under investigation.

In Figure 8 we see a cross section of a convex bump (with a shape index of 1.0), which shows an example of the amount of the patch that is revealed at each radius. At larger radii, and thus, visual angle, much more of the patch is presented.

Mathematica (Wolfram, 1994) was used to determine inverse solutions for the principal curvatures from the associated curvedness and shape index. These solutions were then utilized to calculate the probe patches for any given curvedness or shape index values.

All stimuli were presented initially in a frontoparallel orientation such that the normal at the probe point pointed in the direction of the line of sight. Each trial consisted of the side-by-side presentation of two stimuli, 15 cm apart center to center, each rotating about its own central vertical axis through an angular extent of $\pm 12^\circ$. In each trial, one of the two stimuli was the target patch, whereas the other was a test patch whose shape index varied from trial to trial. The ordering of the two patches was varied randomly from trial to trial. An adaptive staircase was used that would yield a 75% discrimination threshold. A correct response caused the difference between the shape indices of the two patches to be decreased by one step of a predefined step size, whereas incorrect responses caused the difference to be increased by three steps. Ten reversals of direction were averaged to yield the threshold. At the

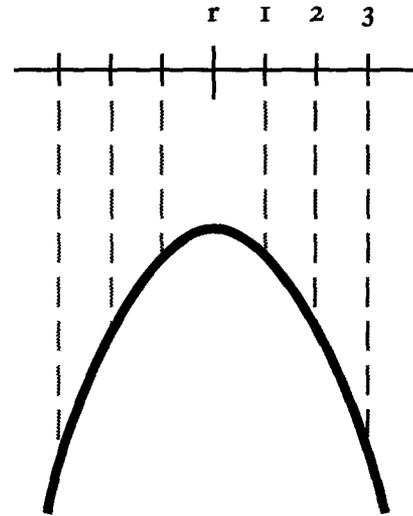


Figure 8. Amount of patch visible as a function of radius.

first, third, and seventh reversal, the step sizes used above were halved. Four runs of 25 shapes at each level were performed by each observer, yielding 100 measurements of each shape.

All observers were familiar with the shape index scale and were made aware of the patch radius and target shape index at the start of each staircase. Before each staircase, the target shape was presented in both positions of the display at random orientations so that the observers could become comfortable with its representation. Observers were allowed to view this display as long as they wished before they continued with the experiment.

Results and Discussion

The individual results for all four observers are shown in Figure 9. Note that for patch sizes of 3° and 4° , the observers showed a higher level of performance than at the 2° case at all shape indices. The results for the 3° and 4° cases are in basic agreement with the results of van Damme et al. (1994; van Damme & van de Grind, 1993) and de Vries

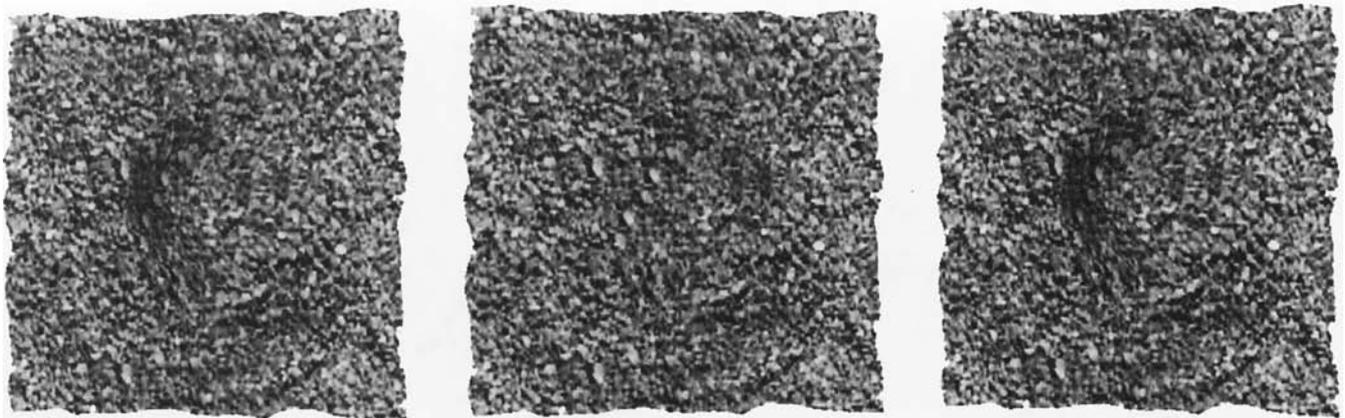


Figure 7. Example stimulus (left pair are crossed, right are divergent).

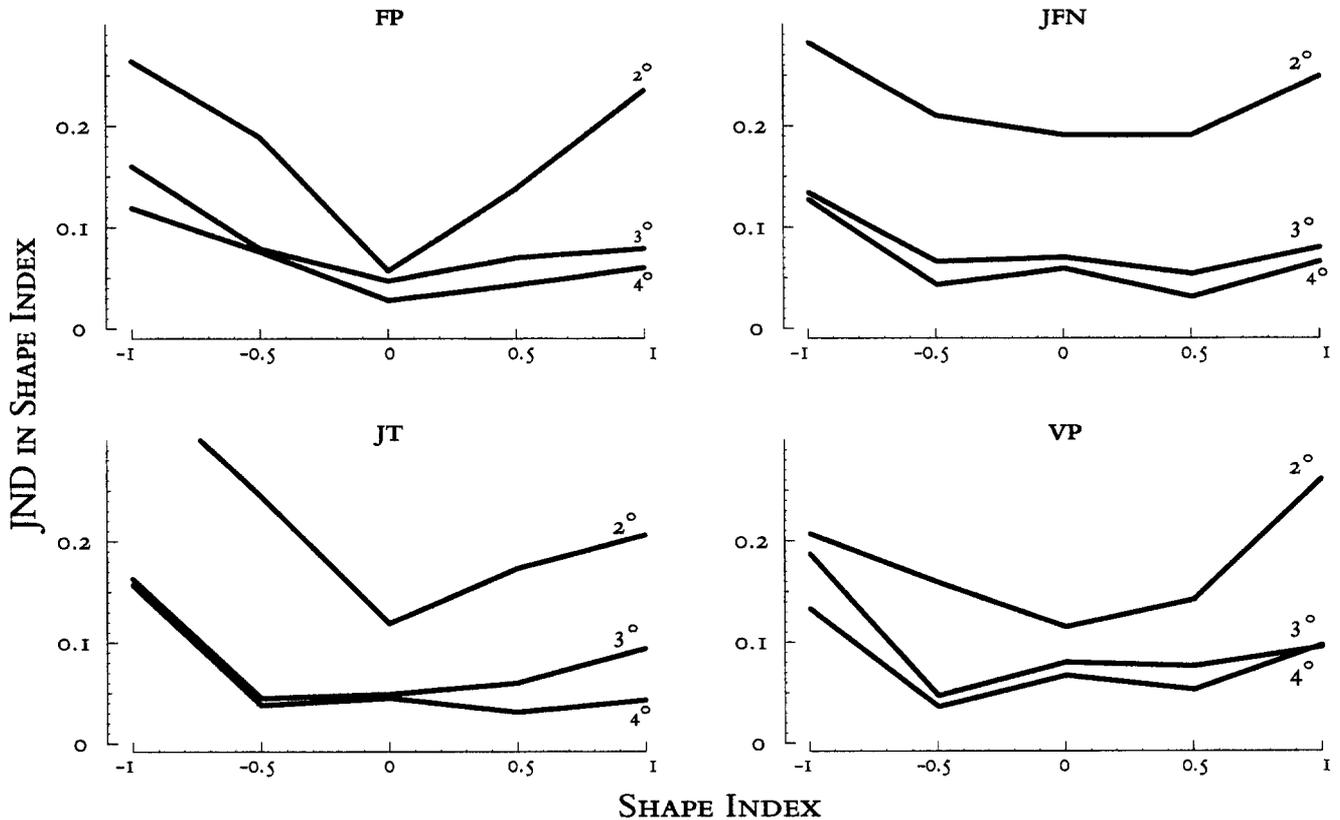


Figure 9. Individual results for Experiment 1 for observers F.P., J.F.N., J.T., and V.P. JND = just noticeable difference.

(1993). van Damme's results showed a strong W shape, with poor levels of performance at shape indices of ± 1.0 and 0.0 and better performance at ± 0.5 . Our results did not reflect as strong an effect, however, their stimuli were larger (6.7° and 6.5°) and, therefore, had more contextual information.

A combined average summary for all four observers is

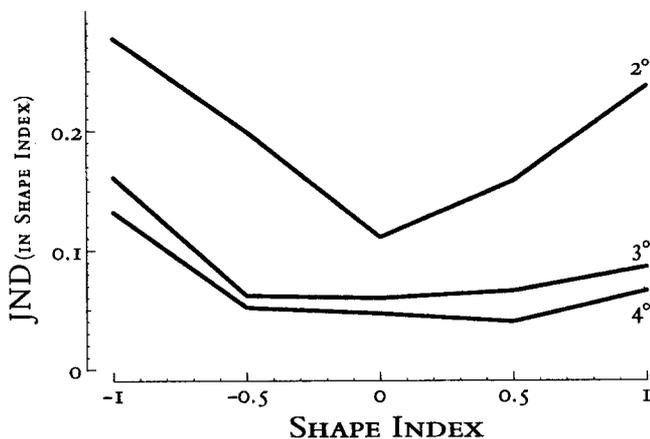


Figure 10. Summary results for Experiment 1. JND = just noticeable differences.

presented in Figure 10. These results show an interaction between the visual angle and shape index with the 2° case differing from the 3° and 4° cases. In the 3° and 4° cases, the differences between the various shape standards was significant, $F(4, 12) = 21.07, p < .001$, as was the effect of patch size, $F(2, 6) = 59.27, p < .005$.

These results suggest that there may be some optimal scale at which we are able to distinguish the features necessary to judge curvature and hence shape index. Just noticeable differences (JNDs) in van Damme's (1994) work with 6.5° patches were even lower than the JNDs we found for our largest case of 4° . From this collection of results, we can conclude that there is a certain amount of context that is necessary for us to make these judgments.

Furthermore, it is interesting to note that the JNDs for the better-performing cases of 3° and 4° are still somewhat high, indicating a rather poor performance in general on this task. A shape index difference of 0.1 (the JND for the bump case) is quite significant when we realize that this difference is $\frac{1}{20}$ th of the entire shape index scale. The lower JNDs around the monoclastic cases (-0.5 and 0.5) show that we are more sensitive when one of the curvatures is zero or flat. We seem to be more sensitive to the presence of the curvature than its magnitude, and we would appear to be slightly less sensitive when the curvatures are in the same direction.

Experiment 2

In Experiment 1, observers reported the subjective impression that they were not attempting to judge the local curvatures to determine shape, but were instead doing some variety of object recognition to find the patch that matched the standard. Note that we utilize the term *object recognition* as referring to the identification of some continuous, globally coherent surface patch. This is not the same as the more global definition of identifying, for example, a telephone. The studies of van Damme et al. (1994) showed that we could, in fact, reliably discriminate shape index of these smoothly varying patches (with 6.5° of context).

Unfortunately, most of the world, in general, is not made up of these mathematically nice pieces of geometry. Even articles that we associate as being spherical are not usually entirely spherical, even at low levels of resolution. A baseball, for example, has stitches that create courses and runs along the roughly spherical leather surface. However, we still see the ball as a round object. Clearly, there is a scale-dependent process at work that allows us to make these judgments. Objects in nature often times have different appearances when viewed at different scales. A tree, for example, might appear just as a two-dimensional T-shaped cutout from sufficiently far away.

Fractal geometry has been of interest in the mathematics community for some time. For many years, the computer graphics community has used fractal techniques for the simulation of natural phenomena (Mandelbrot, 1983; Pentland, 1983). Perlin (1985) and Peachy (1985) simultaneously developed image synthesis techniques based on the fractal-like combination of random noise, opening up a huge class of natural phenomena such as fire, smoke, clouds, marbling, and wood grains to successful visual modeling. In addition, it has been shown by Gilden, Schmuckler, and Clayton (1993) and Knill et al. (1990) that the human visual system is capable of making reliable discriminations of fractal images.

In our second experiment, we chose not only to restrict the context of the target patch, but to also embed it in a context of fractallike turbulence that might provide distracting to the task of object recognition.

Method

Apparatus. The apparatus and setup were the same as in Experiment 1.

Stimuli. The central patch of the stimuli were the same smoothly varying surfaces as utilized in Experiment 1 (see Equation 4). The surrounds were no longer flat, but instead, were generated using a turbulent self-similarity procedure.

A turbulent surface was generated by combining multiple octaves of two-dimensional noise (as per Peachey, 1985, and Perlin, 1985) using a $1/f^\beta$ power spectrum. This type of fractal turbulence is often used to simulate natural phenomena such as mountains, marble, fire, and clouds in computer graphics displays (see Figure 11).

To generate these surfaces, a two-dimensional integer lattice of random noise is generated. We will refer to it as $\lambda(x_i, y_i)$, where we define x_i as the integer part of x and x_r as the fractional part of x (and similarly for y). The distance between the integer lattice points constitutes the base wavelength of the resulting surface, which we will define below. To obtain a value at a given real-valued location, a bicubic interpolation is performed, using the smooth Hermite cubic,

$$h(x) = 3x^2 - 2x^3, \quad (9)$$

and the interpolation function u ,

$$u(\alpha, \beta, \delta) = [h(\delta) \cdot \alpha] + [h(1.0 - \delta) \cdot \beta], \quad (10)$$

where α and β represent the two values to be interpolated and δ represents the percentage along that line. The resulting real-valued noise function is then defined as:

$$\text{noise}(x, y) = u\{\lambda(x_i, y_i), \lambda(x_i + 1, y_i), x_r\}, \\ u\{\lambda(x_i, y_i), \lambda(x_i, y_i + 1), x_r, y_r\}. \quad (11)$$

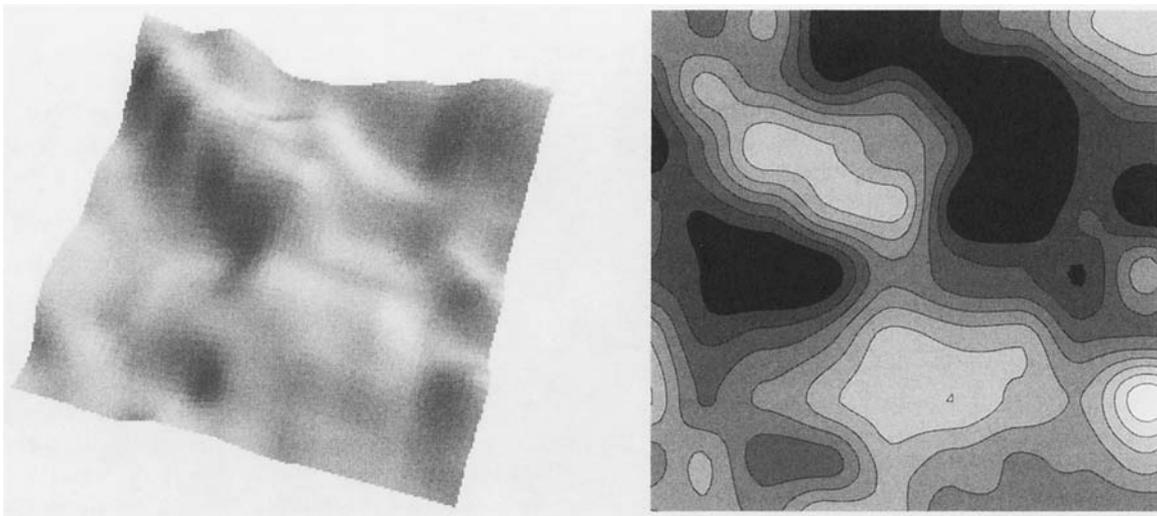


Figure 11. A turbulent surface and its contour plot.

This yields a smooth, differentiable function that can then be used to create the surfaces for the stimuli. To generate turbulent surfaces, n octaves of this noise are summed, yielding a fractal-like Monge form surface, or more formally

$$z_{turb}(x,y) = \sum_1^n \frac{noise(x,y)}{f^n}. \quad (12)$$

These surfaces differ from true fractals in that they are everywhere continuous, and thus, differentiable. Note that by increasing n and, therefore, adding more octaves of noise, smaller features result on this surface.

An example of a turbulent surface created with three octaves of noise is shown in Figure 11, along with its associated contour plot.

For this experiment, three octaves of noise were used, and the spatial frequencies of the resulting features generated by the turbulence were matched to the size of the test patch by defining the integer lattice λ appropriately.

Using a feature analysis algorithm, locations on this infinite surface were sampled and correlated with the current smooth test patch. This algorithm attempted to find regions on the turbulent shape that resembled the probe shape by using a template matching algorithm. Fits were calculated as height correlations between the desired smooth patch and the turbulent surface. When a location of acceptable fit was found, the smooth probe patch was embedded within the turbulent patch using the same blending operation as in Experiment 1. In this case, the resulting surface equation is more complicated because the smooth patch $z(r, \theta)$ is blended with the turbulent surface $z_{turb}(r, \theta)$ as is shown in Equation 13,

$$z'(r,\theta) = z(r,\theta) \cdot \{1 - h[t(r,2,3)]\} + z_{turb}(r,\theta) \cdot h[t(r,2,3)]. \quad (13)$$

As in Experiment 1, the stimuli were presented with motion, shading, texture, and viewed stereoscopically. A blue dot was placed at the probe point.

An example stimulus is shown as a stereo pair (cross fuse on the left, divergent on the right) in Figure 12.

Observers. The displays were presented to the same four observers that served in Experiment 1.

Procedure. As with Experiment 1, observers were presented with two stimuli simultaneously and were asked to indicate the stimulus that contained the predefined standard shape patch by pressing a button on the mouse. Of the two stimuli displayed, one had the target patch embedded, and the other contained a test

patch. The observers were permitted to view the stimuli as long as necessary to make their decision. Feedback was provided in the form of a beep for correct responses.

The experiment was a factorial design of 5×2 , shape index by patch radius, and within subjects. The shape index was probed at the five critical locations: -1.0 , -0.5 , 0.0 , 0.5 , and 1.0 (see Figure 3), whereas curvedness was fixed at a constant 0.5 cm^{-1} . Only patch radii of 1.5 and 2 cm (3° and 4° , respectively) were used in this experiment, as the 2° case proved too difficult in the previous experiment.

As in Experiment 1, the two stimuli on each trial initially appeared in a frontoparallel orientation and were rotated back and forth about the vertical axis through an angular extent of $\pm 12^\circ$. The same adaptive staircase, as in Experiment 1, was run with 10 reversals to provide a 75% discrimination threshold. Four runs of 25 shapes at each level were performed by each observer.

As with Experiment 1, all observers were made aware of the patch radius and target shape index at the start of each staircase. The first display of each staircase consisted solely of the current target shape in both stimuli positions and at random orientations. Observers were allowed to view this display as long as they wished before they continued with the experiment.

Results and Discussion

The individual results for the four observers are shown in Figure 13. As with the previous experiment, the patches that subtended a larger visual angle elicited better performance.

However, the more complex surrounding region caused the judgments to worsen by about 50% across all factors, when compared with the results obtained in Experiment 1.

Similar to the results of Experiment 1, the patch-viewing angle's effect was significant, $F(1, 3) = 59.141$, $p < .005$, indicating a significant difference between the two viewing angles. Shape index was significant as well, $F(4, 12) = 27.55$, $p < .0003$.

The average results in Figure 14 show a stronger overall W shape than the prior experiment, suggesting that the ability to detect the presence or absence of curvature was relatively unaffected, but when judging curvatures that were synclastic or anticlastic, observers had more difficulty than in Experiment 1. Overall, the observers had more trouble



Figure 12. An embedded patch of shape index 0.0 and radius of 2.0.

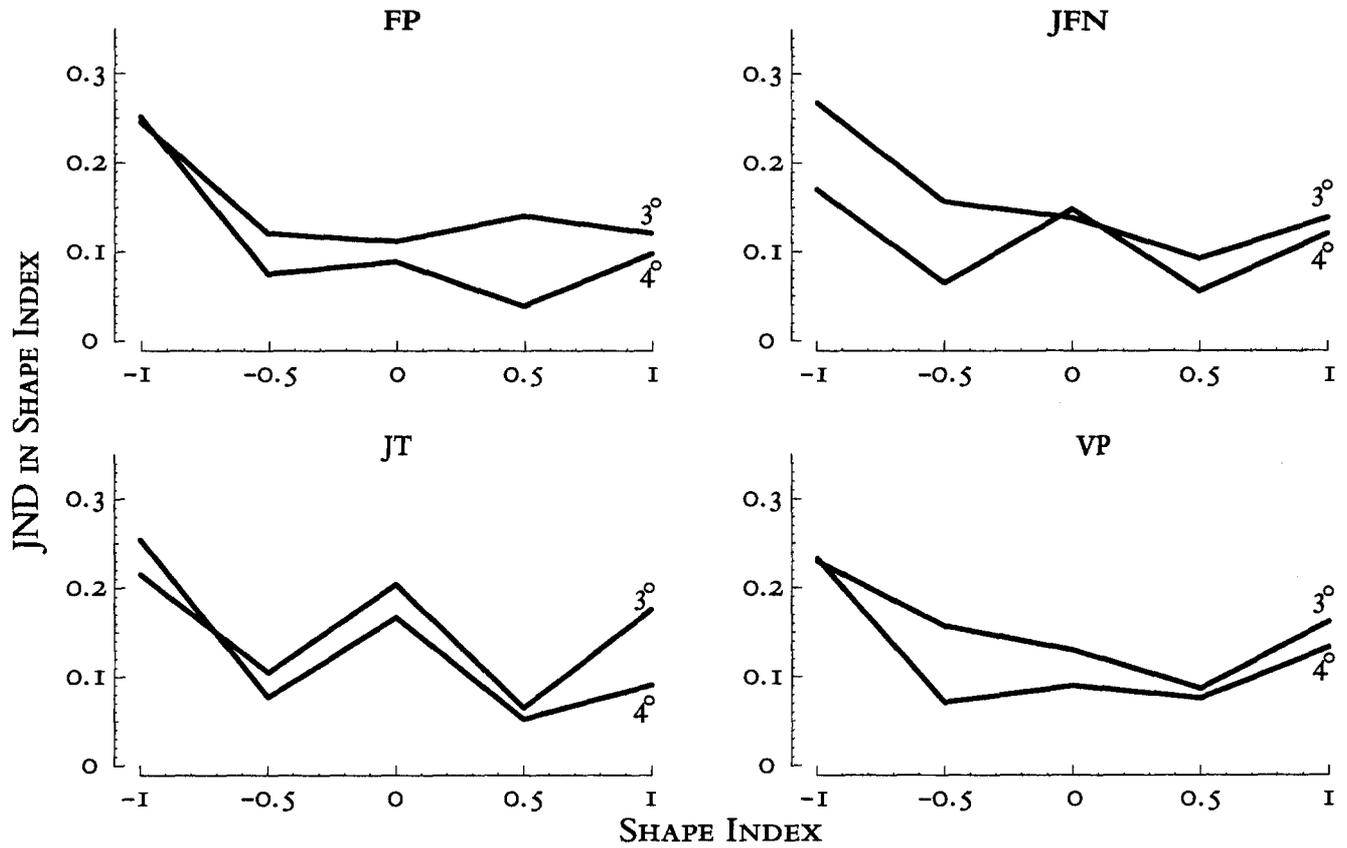


Figure 13. Individual results for Experiment 2 for observers F.P., J.F.N., J.T., and V.P. JND = just noticeable difference.

with this task, performing very poorly even with the large patch size case.

In considering the W-shaped psychometric function obtained in this experiment and in the earlier studies of van Damme et al. (1994), it is interesting to note that the form of this function is quite similar to the overall distribution of shape indices on many natural objects. Figure 15 shows a histogram of the relative frequencies of different shape

indices for a sample of 100,000 points selected at random on the fractal noise surfaces used in this experiment. We have also performed similar analyses on several other types of objects, including some that are manmade, such as an automobile body, and others that occur naturally in the environment, such as a human skull. All of the histograms exhibited this same suspension-bridge shape. The fact that this same distribution occurs over a wide range of objects suggests that it may arise from the basic constraints of surface topology. In comparing this distribution with psychometric functions of human observers, it would appear that sensitivity to different shape indices is roughly proportional to their frequency of occurrence in the natural environment.

Although there were several redundant sources of information in these displays, including shading, texture, motion, and binocular disparity, from which the three-dimensional structure of the depicted surfaces could potentially have been determined, the ability of observers to discriminate local shape was nonetheless surprisingly poor. Note in Figure 14, for example, that the JND for negative synclastic points was 0.25, a quarter of the entire scale, even for the larger patch sizes of 4°. There is some reason to suspect, moreover, that these results may actually overestimate observers' performance.

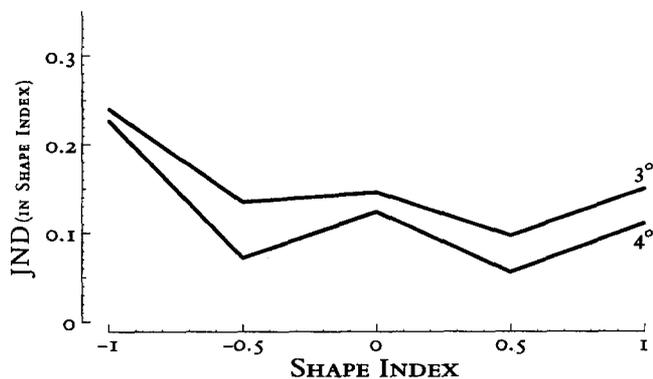


Figure 14. Summary results for Experiment 2. JND = just noticeable differences.

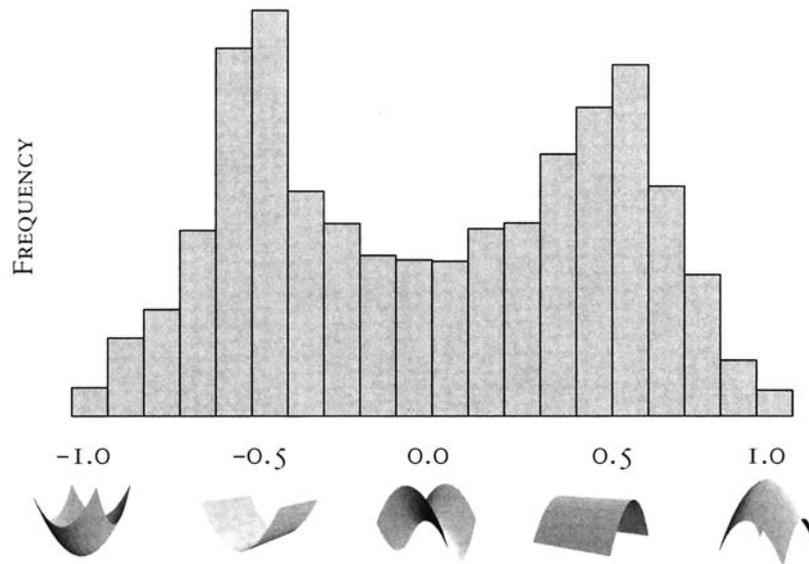


Figure 15. Histogram of shape index of turbulent surfaces.

An important characteristic of the quadric surface patches described in Equation 4 is that the tangent plane at the depicted probe point was always perpendicular to the observer's line of sight. There is a growing amount of evidence that observers' perceptions of three-dimensional structure in both real and pictorial displays can be systematically distorted such that intervals in depth may appear stretched or compressed relative to those in the frontoparallel plane (see Norman, Todd, & Phillips, in press, for a review). Let us suppose, for the sake of argument, that similar perceptual distortions occurred in this experiment. For most points on an object's surface an affine distortion of its perceived structure in depth would differentially affect κ_1 and κ_2 , and should, therefore, result in a systematic distortion in perceived local shape. However, for those points with a frontoparallel tangent plane, an affine distortion in depth would have identical effects on both κ_1 and κ_2 , so that perceived local shape would remain invariant. In other words, it is reasonable to suspect that the sample of probe patches used in this experiment may not have been representative of the vast majority of surface patches whose tangent planes were not perpendicular to the observer's line of sight.

Experiment 3

In an effort to avoid this problem in Experiment 3, we abandoned the use of imbedded quadric surface patches in favor of the more natural, irregular surfaces created by the turbulence function we used for the surrounding context in Experiment 2.

Because it appears that observers cannot reliably differentiate between different magnitudes of curvature (and thus shape index) as a patch's context becomes more restricted, we were curious if they could perform the conceptually simpler task of simply identifying the sign of the local

Gaussian curvature. This amounts to the modest task of estimating the signs of the two principal curvatures κ_1 and κ_2 . As shown in Equation 1, the magnitude of the Gaussian curvature will be positive if the κ_1 and κ_2 are in the same direction (synclastic), negative if they are in opposite directions (anticlastic), and zero if either of the curvatures is zero (monoclastic).

Whereas the shape index of a patch is only invariant under affine distortions in depth when its tangent plane is perpendicular to the line of sight, the sign of the Gaussian curvature is invariant under such distortions for all possible orientations of the tangent plane. Because observers were simply required to judge the sameness or differentness of two directions (not magnitudes) of curvature, we suspected that this task might be much easier than the discrimination of shape index employed in Experiments 1 and 2.

Method

Apparatus. The apparatus and setup were the same as in Experiments 1 and 2.

Stimuli. As noted above, we abandoned the globally smooth Monge form representation of the probe patches in favor of using the noise or turbulent surfaces used as surrounds in Experiment 2. For these surfaces there is no embedded probe patch, smooth or otherwise. Instead, we sample the surface in search of locations that meet certain criteria set forth in the Procedure section below.

Each stimulus was a 12-cm square patch consisting of either two or three octaves of turbulent noise. The lowest spatial frequency of the turbulence was chosen to provide gross features of approximately 4 cm (and thus, 4° of visual angle) in size, as this was the optimal performance level in the previous experiments. The $1/f^\beta$ structure of the resulting harmonics created smaller features that ranged from approximately 1.0 cm to 4.0 cm in size and whose amplitude ranged from one quarter to one half of the amplitude of the gross features (see Figures 11 and 16).

As with the previous experiments, all stimuli were presented



Figure 16. Stereo pairs of example turbulent stimuli.

with motion, shading, texture, and were viewed stereoscopically. A blue dot was placed at the probe point at the center of the patch.

Observers. The same four observers that served in Experiments 1 and 2 took part in this experiment.

Procedure. For this task, observers were presented with a single patch with the probe point on the location under test. For all test points, the surface was translated such that the probe point was at the center of the screen and a constant distance (57.4 cm) away from the observer. The observers indicated if the surface at the probe location was of positive or negative Gaussian curvature by pressing one of two keys on the keyboard. As with the previous experiments, the observers were permitted to view the stimuli as long as necessary to make their decision. Feedback was provided in the form of a beep for correct responses.

This experiment was a 2×5 (turbulence octaves by absolute value of shape index), within-subjects factorial design. Two levels of turbulence, two and three octaves, were used. These two levels seem to create similar amounts of fractal information to that found in nature (Knill et al., 1990; Pentland, 1983). The three-octave case introduces small features, approximately 1 cm in radius, and thus, disrupts the regularity of the surface more than the two-octave case whose features are roughly 2 cm.

The noise function used for this set of experiments was infinite in extent. The turbulent patches were stochastically sampled across a $1,000 \text{ cm} \times 1,000 \text{ cm}$ region, and the local curvature information was calculated at each location. A set of points was collected that met the criteria of curvednesses of $0.5 \text{ cm}^{-1} \leq C \leq 1.0 \text{ cm}^{-1}$, and shape indices whose absolute value fell into one of six bins, 0.05 wide, centered at 0.143, 0.286, 0.429, 0.571, 0.714, and 0.857. Four blocks of 25 trials at each of the six levels of shape index and two levels of turbulence were run by each observer.

Notice in Figures 3 or 5 that, for shape indices on the interval $-0.5 < S < 0.5$, the two principal curvatures are in opposite directions, therefore, these shapes have a negative Gaussian curvature. Similarly, shape indices on the intervals $S < -0.5$ and $S > 0.5$ have principal curvatures in the same direction, thus, they possess positive Gaussian curvature. The six bins provided three levels of negative Gaussian curvature and three levels of positive Gaussian curvature.

Results and Discussion

The data were analyzed by using a probit-style procedure to obtain the best-fitting cumulative normal distribution for

each condition. The results for the three-harmonic case are shown in Figure 17. The data correlated to the cumulative normal function ($R^2 = 0.99, p < .001$).

The extreme shallowness of the psychometric function reflects the observers' poor performance for this task. There was also a leftward bias to the psychometric function, reflecting more trouble identifying the negative curvatures than the positive. The 25% and 75% thresholds were at shape indices of approximately 0.14 and 0.68, respectively.

The psychometric function for the two-harmonic case is shown in Figure 18. Phenomenological reports from the observers confirmed this case as being easier than the three-harmonic case. Indeed, performance improved somewhat, especially on the positive side of the scale. The 25% and 75% thresholds improved to 0.31 and 0.57, respectively.

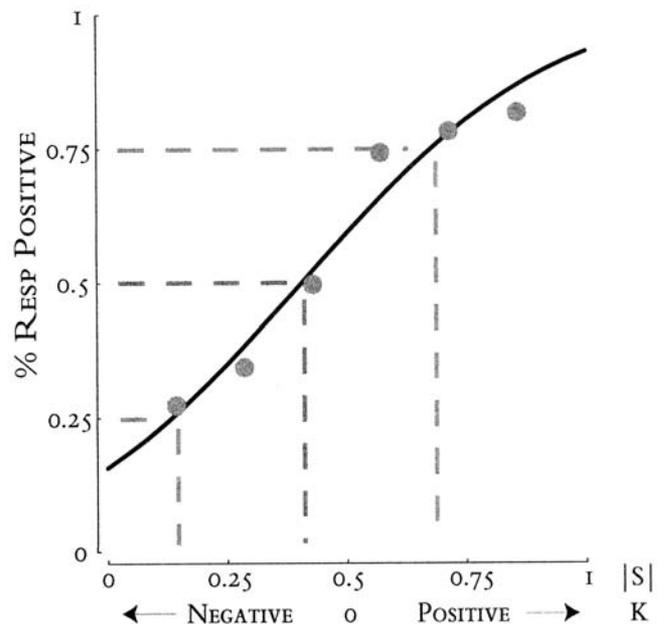


Figure 17. Summary results of three-harmonic case of Experiment 3. Resp = response.

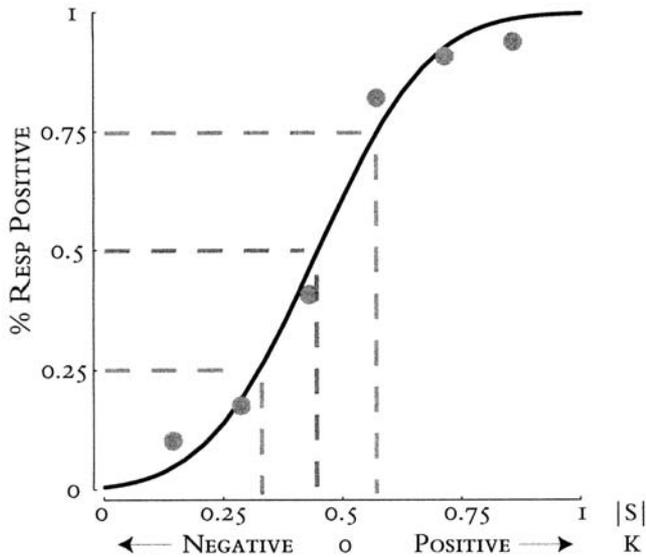


Figure 18. Results of the two-harmonic case of Experiment 3. Resp = response.

Finally, a Kolmogorov–Smirnov two-sample test was run to compare the two psychometric functions. This yielded $D(1200, 1200) = 0.22$, exceeding the critical value of 0.067, showing a significant difference of the two cumulative normal distributions with a $p < .01$.

As with the three-harmonic case, there was a slight bias toward the left, indicating that observers performed better at identifying the positive curvatures than the negative curvatures.

It would appear that in both cases the additional high-frequency features in the area surrounding the probe location interfered with the identification of the sign of the principal curvatures. In addition, because of the shift seen in the psychometric function, locations where the patches curved in opposite directions suffered more adversely.

These effects were greatly reduced in the two-harmonic case, which reduced the spatial frequency of these additional confounding features. This further reinforces the results of Experiments 1 and 2, showing that there is a large region of smooth context necessary for the identification of the local surface curvatures.

Experiment 4

Experiments 1–3 suggest that there is some contextual area over which these curvature judgments are made. Because the curvature metrics presented so far have all been local, they are only valid at the probe location, and therefore, may change rapidly from point to point on the surface. This reflects the nature of the turbulent surfaces presented in Experiments 2 and 3. The smooth probe patches utilized in Experiments 1 did not suffer from this rapid variance. In Experiment 2 we introduced a turbulent context and still used the smooth probe patch from Experiment 1. Experiment 3 introduced a turbulent probe location as well as the

surround, but it continued to use a purely local shape metric (sign of Gaussian curvature), which was defined only at the probe location.

In Experiment 4 we define a metric whose neighborhood is no longer local, and therefore, can be used to investigate the region over which this contextual information is necessary.

Method

Apparatus. The apparatus and setup were the same as in Experiments 1, 2, and 3.

Stimuli. Stimuli for Experiment 4 were the same turbulent surfaces used in Experiment 3. Because, in Experiment 3, performance was best in the two-harmonic case, it was used only in this experiment.

Observers. Three of the four observers that served in Experiments 1, 2, and 3 took part in this experiment.

Procedure. For this task, observers were presented with a single patch with the probe point on the location under test. For all test points, the surface was translated such that the probe point was at the center of the screen and a constant distance (57.4 cm) away from the observer. The observers indicated if the surface at the probe location was of positive or negative Gaussian curvature by pressing one of two keys on the keyboard. As with the previous experiments, the observers were permitted to view the stimuli as long as necessary to make their decision.

An extension of the shape index was created, yielding a metric that took into account a larger region than a strictly local neighborhood. A disk of some radius r , centered at the probe point, was randomly sampled with the resulting set of surface normals and frames used in the principal curvature calculations. In essence, this results in average principal curvatures for some region of radius r , centered at the probe point. This had the net effect of canceling out any smaller perturbations of the surface in favor of the more statistically dominant feature. For example, if a concave synclastic patch of some chosen radius was dotted with smaller, higher frequency bumps, the operator described above would report the lower frequency depression if the radius of the disk were roughly the size of the depressed region. If the radius of the disk were as small as the high-frequency bumps, it would instead reflect the shape of these features.

This experiment was a 3×5 , sample disk radii by absolute value of shape index, within-subjects factorial design. All displays had two octaves of turbulence with a base wavelength of 4° of visual angle, thus, the resulting harmonics created smaller features up to 2° in size.

Three sampling disk radii, 0.5 cm, 1.0 cm, and 2.0 cm, yielding visual angles of 1° , 2° , and 4° , respectively were used. As in Experiment 3, the turbulent patches were stochastically sampled and the local curvature information was calculated at each location using the area-based method described above. A set of points was collected that met the criteria of curvednesses of $0.5 \text{ cm}^{-1} \leq C \leq 1.0 \text{ cm}^{-1}$, and shape indices whose absolute value fell into one of six bins, 0.05 wide, centered at 0.143, 0.286, 0.429, 0.571, 0.714, and 0.857. Three blocks of 25 trials at each of the six levels of shape index and three disk radii were run by each observer.

On each trial, observers were presented with a single patch and responded as to the sign of the Gaussian curvature at the probe point. The keyboard was used for responses, and feedback was provided for correct responses. At the beginning of each block observers were informed of the sampling disk radius that would be used. They were also given feedback after each trial to help them

learn to integrate information over an appropriate sized neighborhood to be used in making their responses.

Results and Discussion

The psychometric functions from all three disk radii are displayed in Figure 19. All three functions fit the cumulative normal with R^2 's of less than 0.94, $p < .005$. The Kolmogorov-Smirnov test showed a significant difference ($p < .01$) between the 2° and 4° cases and the 1° and the 4° cases. There was not a significant difference between the 1° and the 2° cases.

Observers reported some frustration with this task, especially with respect to the feedback, which at times felt inconsistent. As the radius of the neighborhood was increased, the psychometric function became shallower. Only for the best case, seen on the left of Figure 20, does the psychometric function roughly equal the performance of the purely local shape function. A Kolmogorov-Smirnov test between the 1° case (which used two harmonics) from this experiment and the purely local, two-harmonic case of Experiment 3, showed no significant difference.

Thus, it would seem that observers are probably not averaging the local shape information over a fixed neighborhood as we may have expected from the results of Experiments 1-3. Because a fixed sampling radius was used in each block, this also may suggest that the observers were not using a fixed scale to determine the features of the point under investigation.

Indeed, much of the observer's frustration seemed to stem from the fact that it was difficult to know the size of the feature whose curvatures were being evaluated. When considering the results from the previous three experiments, it would appear that the area over which we make these curvature judgments depends on the size of the feature in question. Unfortunately, as we have seen in Experiments 3

and 4, we are now faced with the problem of deciding exactly what constitutes a feature.

General Discussion

Our informal description of shape is just that, a casual, nonspecific naming of the configuration of the constituent parts of the object. A more rigorous definition of shape would require some regional metric descriptions, which, in combination, map to our higher order nominal descriptions. Most models of solid shape perception proposed to date have relied on a representation similar to that of the $2\frac{1}{2} D$ sketch.

One problem with these models is that they are sensitive to variations in orientation and scale. Two similar but different sized objects will require two completely different representations because the depths of points at nominally equal locations will not be the same. Also, a simple rotation of the object destroys our ability to recover the configuration.

Pinker (1985) points out that this scheme is not well suited to a representational system because it would require an excessive number of object representations with which to match. Every possible set of orientations or depths would require a separate representation because they vary excessively whenever the object or viewer move relative to one another. As the result of several experiments performed on the observers' ability to judge slanted planes, Gibson (1979) later withdrew his support for such an idea.

Similar results have been obtained for local orientation judgments on more complex surfaces. Recent experiments by Koenderink, 1994 and Koenderink, van Doorn, and Kappers (1992; 1995), Koenderink and van Doorn (1995), Koenderink, Kappers, Todd, Norman, and Phillips (1996), Norman et al. (in press), and Todd, Koenderink, van Doorn, and Kappers (1996) show that, across observers and viewing conditions, there is an extraordinary amount of variance in the surface reconstructed from the perceived orientations.

One solution to this problem is to use high-order position-invariant measures, such as shape index and Gaussian curvature to describe the surfaces of an object. Some previous studies have used the shape index in several different contexts; shape-from-stereo (de Vries, 1993), shape-from-motion-parallax (van Damme et al., 1994; van Damme & van de Grind, 1993), and shape-from-shading (Erens, 1993a; 1993b).

In the four experiments described in this article we subjected observers to various complex stimuli and evaluated their ability to judge two higher order measures of surface structure at a depicted probe point; shape index and sign of Gaussian curvature, which evolve from measurements of local curvature. We saw in Experiment 1 that our perception of the local curvatures, and thus the shape index, depended on how much of a surface we were shown. The local property of shape index seemed to require a significant neighborhood, not just its local information. A minimum of 2° to 3° of contextual information was needed to perform this task reliably. In previous research, Todd and Reichel

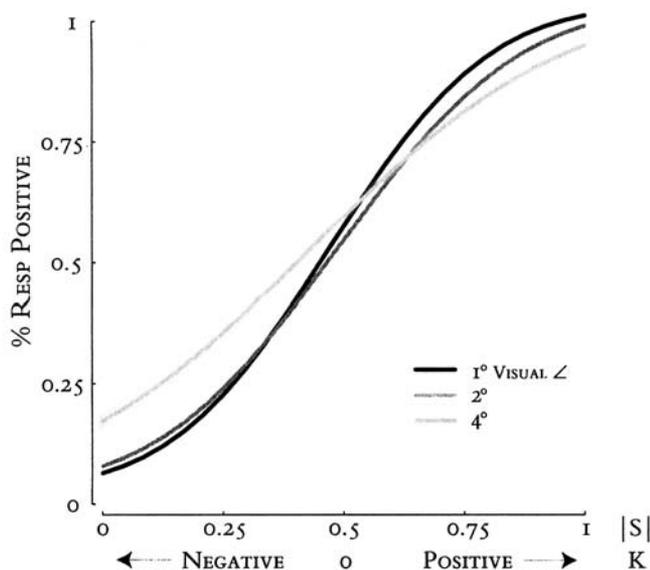


Figure 19. Summary of Experiment 4. Resp = response.

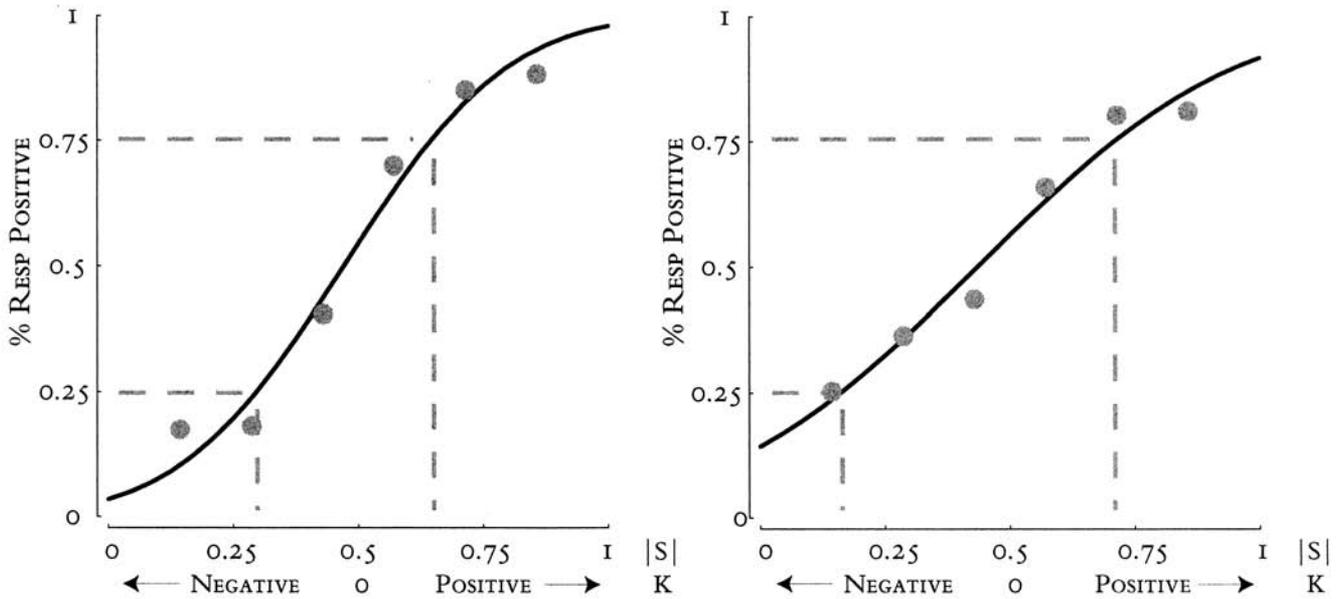


Figure 20. Results of the 1° (left panel) and 4° (right panel) cases. Resp = response.

(1989) have shown similar difficulty in deducing the ordinal structure of surfaces when context is restricted.

Experiment 2 helped to confirm the observer's reports that they were basing their responses on the global context rather than the local shape index. The contextual noise added by the turbulent patches reduced performance by approximately 50% relative to Experiment 1.

Because of the potentially nongeneric properties of the quadric patches used in Experiments 1 and 2, Experiment 3 used natural appearing, turbulent surfaces and probed points whose tangent surface was no longer confined to the frontoparallel plane. The task was simplified to judging the sameness or differentness of the directions of the principal curvatures. Here, we saw a significant decrease in the observer's ability to make this relatively simple judgment when the amount of turbulence, and thus, the number of higher frequency features surrounding the probe point was increased.

Finally, in Experiment 4 we drew on the results of Experiments 1 and 2 by attempting to take into consideration the quality of the surface around the probe points. A non-local, averaging curvature operator was developed in an attempt to describe the neighborhood-based curvatures in a given area of the patch. Using the turbulent stimuli from Experiment 3, observers still had difficulty in judging the sign of Gaussian curvature when a larger context was considered for the calculation of the curvatures. Because this experiment used fixed scales for the shape operator, it would suggest that observers were not using a fixed scale to make these judgments. As suggested by Koenderink (1984b), perhaps the observers were exploiting their ability to tune their outer and inner scales to match the features found in the stimuli.

Across these experiments, we have manipulated the neighborhood three different ways: by varying the amount

of smooth context surrounding the probed area, varying the amount of harmonic noise surrounding the patch, and finally by creating an operator that worked over varying neighborhood sizes. Despite all the cues given, including motion, stereo, texture, and shading, our results have shown that performance remains rather poor for these types of discriminations, and that it requires a significant amount of context from the surrounding neighborhood. Given this much contextual information, however, it would seem the task becomes that of recognizing a global surface patch, rather than its local curvature.

Edge-based approaches, such as those of Biederman (1987) and Richards et al. (1987), seem to have a good deal of explanatory power when it comes to the edges and occlusion contours of an object. However, any smoothly curved internal structure will lack adequate description using these models. Can we extend the idea of geons being parts of objects to what might constitute a part primitive within the internal structure of a smoothly curved object?

One such approach is that of Koenderink (1977) and Koenderink and van Doorn (1980), who demonstrated that the local structural properties of an object can be mathematically determined through analysis of isophotes (i.e., lines of constant luminance), yielding connected regions of positive and negative Gaussian curvature. This establishes two important concepts: First, that the local structure is determinable using features other than occlusion contours, and second, that these features can be used in a viewer and light source invariant way. It remains to be seen if this characteristic of solid shape can be used empirically.

Another interesting possibility is that we are utilizing some multiresolution strategy where a set of uniform primitives are used but at varying scales. Koenderink (1984b) has also investigated the use of scale space as a method to describe the scale related information available in an image.

Scale can be thought of as consisting of two components, the first being the outer scale, essentially the extent of the image, and the inner scale, which specifies the resolution of the elements that make up the image. This idea can easily be extended to the geometry and topological features of a surface as well. As we pointed out in our discussion of the fractal surrounds, the scale utilized to describe an object only makes sense in a given context. A golf ball, for example, with its hundreds of dimples, is a sphere at one scale and a concave pitted surface at another. In this case, as with the stimuli from our experiments, the features we wish to examine exist at some scales and not at others. Our research suggests that perceived shape is most salient over spatial neighborhoods of approximately 3° to 4°. Whether this holds for other classes of objects and viewing conditions remains to be experimentally determined.

References

- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115–147.
- Blum, H. (1973). Biological shape and visual science, part 1. *Journal of Theoretical Biology*, *38*, 205–287.
- de Vries, S. C. (1993). *Binocular perception of spatial configuration*. Utrecht, The Netherlands: Universiteit Utrecht.
- Erens, R. (1993a). *Visual perception of shape from shading*. Utrecht, The Netherlands: Universiteit Utrecht.
- Erens, R. G. F. (1993b). Perception of local shape from shading. *Perception & Psychophysics*, *54*, 145–157.
- 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*. Hillsdale, NJ: Erlbaum.
- Gilden, D. L., Schmuckler, M. A., & Clayton, K. (1993). The perception of natural contour. *Psychological Review*, *100*, 460–478.
- Gray, A. (1993). *Modern differential geometry of curves and surfaces*. Boca Raton, FL: CRC Press.
- Hilbert, D., & Cohn-Vossen, S. (1952). *Geometry and the imagination* (P. Nemenyi, Trans.). New York: Chelsea Publishing.
- Hoffman, D. D., & Richards, W. A. (1985). Parts of recognition. In S. Pinker (Ed.), *Visual cognition* (pp. 66–96). Amsterdam: Elsevier.
- Knill, D. C., Field, D., & Kersten, D. (1990). Human discrimination of fractal images. *Journal of the Optical Society of America A*, *7*, 1113–1123.
- Koenderink, J. J. (1977). How an ambulant observer can construct a model of the environment from the geometrical structure of the visual inflow. In G. Hauske & E. Butenandt (Eds.), *Kybernetiks*, (pp. 224–247). Munich, West Germany: Oldenberg.
- Koenderink, J. J. (1984a). What does the occluding contour tell us about solid shape. *Perception*, *13*, 321–330.
- Koenderink, J. J. (1984b). The structure of images. *Biological Cybernetics*, *50*, 363–370.
- Koenderink, J. J. (1989). *Solid shape*. Cambridge, MA: MIT Press.
- Koenderink, J. J. (1994). On so-called paradoxical monocular stereoscopy. *Perception*, *23*, 583–594.
- 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*, *22*, 869–878.
- Koenderink, J. J., & van Doorn, A. J. (1980). Photometric invariants related to solid shape. *Optica Acta*, *27*, 981–996.
- Koenderink, J. J., & van Doorn, A. J., (1995). *Relief: Pictorial and Otherwise Image and Vision Computing*, *2*, 321–334.
- Koenderink, J. J., van Doorn, A. J., Gaussian, & Kappers A. M. L. (1992). Surface perception in pictures. *Perception & Psychophysics*, *52*, 487–496.
- Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. L. (1995). Depth relief. *Perception* *24*, 115–126.
- Lipshutz, M. (1969). *Theory and problems of differential geometry*. New York: McGraw-Hill.
- Mandelbrot, B. B. (1983). *The fractal geometry of nature*. (2nd ed.). New York: Freeman.
- Marr, D., & Nishihara, H. K. (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London B*, 269–294.
- Norman, J. F., Todd, J. T., & Phillips, F. (in press). The perception of surface orientation from multiple sources of optical information. *Journal of Experimental Psychology: Human Perception and Performance*.
- Peachey, D. R. (1985). Solid texturing of complex surfaces. *Computer Graphics*, *19*, 279–286.
- Pentland, A. (1983). Fractal-based description of natural scenes. *Proceedings of IEEE CVPR-Computer Vision and Pattern Recognition*, 201–209.
- Perlin, K. (1985). An image synthesizer. *Computer graphics*, *19*, 287–296.
- Pinker, S. (Ed.). (1985). *Visual cognition*. Cambridge, MA: MIT Press.
- 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.
- Rogers, D. F., & Adams, J. A. (1990). *Mathematical elements for computer graphics*. (2nd ed.). New York: McGraw-Hill.
- Todd, J. T., Koenderink, J. J., van Doorn, A. J., & 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*, *22*, 695–706.
- Todd, J. T., & Reichel, F. D. (1989). Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psychological Review*, *96*, 643–657.
- van Damme, W. J. M., Oosterhoff, F. H., & van de Grind, W. A. (1994). Discrimination of 3-D shape and 3-D curvature from motion in active vision. *Perception & Psychophysics*, *55*, 340–349.
- van Damme, W. J. M., & van de Grind, W. A. (1993). Active vision and the identification of three-dimensional shape. *Vision Research*, *33*, 1581–1587.
- Wolfram, S. (1994). *Mathematica* (Version 2.2.2) [computer software]. Champaign, IL: Wolfram Research.

Received September 15, 1994

Revision received March 17, 1995

Accepted June 7, 1995 ■