

Perceptual representation of visible surfaces

FLIP PHILLIPS

Skidmore College, Saratoga Springs, New York

JAMES T. TODD

Ohio State University, Columbus, Ohio

and

JAN J. KOENDERINK and ASTRID M. L. KAPPERS

Utrecht University, Utrecht, The Netherlands

Two experiments are reported in which we examined the ability of observers to identify landmarks on surfaces from different vantage points. In Experiment 1, observers were asked to mark the local maxima and minima of surface depth, whereas in Experiment 2, they were asked to mark the ridges and valleys on a surface. In both experiments, the marked locations were consistent across different observers and remained reliably stable over different viewing directions. These findings indicate that randomly generated smooth surface patches contain perceptually salient landmarks that have a high degree of viewpoint invariance. Implications of these findings are considered for the recognition of smooth surface patches and for the depiction of such surfaces in line drawings.

It has long been recognized that a convincing pictorial representation of an object can sometimes be achieved by drawing just a few critical lines. Consider, for example, the photograph of a Henry Moore sculpture that is presented in the left panel of Figure 1 and the line drawing of this object that is presented in the right panel. Although the line drawing leaves out many salient aspects of the photographic image, such as surface color and texture, it effectively conveys the overall three-dimensional (3-D) structure of the depicted scene. Indeed, research has shown that reaction times for recognizing objects from line drawings are no different from the response times obtained for natural photographs (Biederman & Ju, 1988).

When an illustrator creates a line drawing of a 3-D scene, the lines are carefully positioned to denote specific landmarks on a surface. One type of landmark to consider in this regard includes the edges that connect planar faces of polyhedral objects. Several such edges from the base of the Moore sculpture are highlighted in Figure 2. Edges are one-dimensional structures and are, therefore, denoted by lines, but there are also important zero-dimensional landmarks (i.e., points) that are often essential for the perceptual interpretation of drawings. Consider, for example, the point labeled *a* in Figure 2 that marks a vertex where three

edges coterminate. There have been numerous mathematical analyses to show how the pattern of vertexes in an image provides potential information about 3-D shape (Clowes, 1971; Draper, 1981; Huffman, 1977; Mackworth, 1973; Malik, 1987; Waltz, 1975a, 1975b), and there is also compelling empirical evidence that these landmarks are a critical source of information for object recognition by human observers (Biederman, 1987).

Another type of landmark that is frequently denoted in line drawings includes occlusion contours that define the boundary between visible and occluded regions of a surface. Theoretical analyses have shown that occlusion contours provide potential information about the sign of surface curvature in their immediate local neighborhoods (Koenderink, 1984; Koenderink & van Doorn, 1982b). There is also psychophysical evidence that human observers can exploit that information for the perception of 3-D shape (Koenderink, van Doorn, Christou, & Lappin, 1996) and that it can also be used in the perceptual interpretation of other sources of information, such as shading or texture (Reichel & Todd, 1990; Todd & Reichel, 1989). Occlusion contours are one-dimensional structures that are denoted by lines, but they also contain a variety of salient point structures that could be used as landmarks for human perception. To qualify as a landmark, a contour point must have some salient attribute that makes it stand out from its neighbors. One such attribute is the abrupt termination of a contour at a cusp or T-junction (Koenderink & van Doorn, 1982a; Malik, 1987). Others could be defined by extrema or zero crossings along various dimensions of differential structure. For example, suppose that one was asked to identify the point on a contour that is farthest to the right, such as the one labeled *b* in Figure 2. Although observers have no difficulty

This collaboration was supported in part by NATO Scientific Exchange Grant CRG 92065. In addition, J.T.T. and F.P. were supported by grants from NIH (R01-EY12432-01) and NSF (SBR-0079277). F.P. was supported by a grant from the Keck Foundation. The authors thank Victor J. Perotti of the Rochester Institute of Technology for his advice, consultation, and insight on many issues. Correspondence should be sent to F. Phillips, Department of Psychology, Skidmore College, Saratoga Springs, NY 12866-1632 (e-mail: flip@skidmore.edu).




Figure 1. Henry Moore's *Three Piece Reclining Figure: Draped*, 1975.

in making such judgments, the location of the resulting landmark would be highly unstable, because it would vary with object orientation. An alternative approach that overcomes this difficulty is to define landmarks by extrema of curvature, such as those labeled *c* and *d* in Figure 2 (Richards, Koenderink, & Hoffman, 1987).

There have been several experiments reported in the literature in which the perception of surface landmarks on object silhouettes has been examined (Hoffman & Singh, 1997; Norman, Phillips, & Ross, 2001; Siddiqi, Tresness, & Kimia, 1996; Singh, Seyranian, & Hoffman, 1999). This research has been focused primarily on the ability of observers to segment objects into perceptually distinct parts. In a typical experiment, observers are presented with a smoothly curved silhouette, and they are asked to identify the perceived part boundaries by marking points along its contour. For most of the objects that have been studied thus far, observers are able to perform this task with a high degree of both intrasubject and intersubject reliability. Moreover, as was originally hypothesized by Hoffman and Richards (1984), the perceived part boundaries tend to be located at extrema of negative curvature.

Unfortunately, there has been almost no research on the perception of landmarks in interior regions of smoothly curved surfaces. Although Hoffman and his colleagues have argued that negative curvature extrema should define part boundaries on surfaces as well as on silhouettes, the evidence to support this claim has been primarily anecdotal (see, e.g., Hoffman & Richards, 1984; Hoffman & Singh, 1997). Some indirect evidence for the existence of landmarks on smooth surfaces comes from recent experiments on the perceived correspondence relations between different views of the same object (Koenderink, Kappers, Pollick, & Kawato, 1997; Koenderink, van Doorn, Kappers, & Todd, 1997; Phillips, Todd, Koenderink, & Kappers, 1997). Suppose, for example, that an observer is presented with pairs of objects that are structurally identical, except that they have different random textures and are positioned at different orientations in depth. A single point on one of the objects is highlighted with a small colored dot, and the observer is required to identify the corresponding point at a different orientation on the second object. The ability of observers to identify the point-to-point correspondences over different orientations is surprisingly ac-

curate. For orientation differences up to 30°, the average errors within the object's projected image can be as small as just a few minutes of arc (Phillips et al., 1997). When asked to reveal their subjective impressions of this task, almost all observers describe a similar strategy. The target points are localized by identifying their positions relative to other salient landmarks, such as the "top of a bump" or the "edge of a cliff." These are then used to triangulate the same positions when viewed from a different orientation.

Given the consistency of observers' subjective impressions, we were curious to discover the specific attributes of surface structure by which these landmarks are perceptually defined, and the present series of experiments were designed to address this issue. Our working hypothesis as we began these experiments was that the perceived landmarks on a surface would be located at local extrema (i.e., maxima or minima) of some underlying geometric property of the surface structure (e.g., depth, slant, curvedness, etc.). These could include point singularities, such as the peaks of mountains, where the underlying dimension is at a maximum (or minimum) in all directions, or line singularities, such as ridges or valleys, where the dimension is at a maximum (or minimum) in all directions except one. It is important to recognize that local extrema can have varying degrees of stability across different levels of dif-




Figure 2. Some potential *landmark* points on Moore's sculpture: *a* is a coterminal of three edges; *b* is the point farthest to the right on the occlusion boundary; *c* and *d* represent curvature extrema.

ferential structure. Unlike extrema of depth or slant, the locations of curvature singularities do not change as a function of surface orientation. Thus, they could be especially useful for defining landmarks that are viewpoint invariant (see also Hoffman & Richards, 1984; Richards et al., 1987; Siddiqi et al., 1996).

EXPERIMENT 1

Our initial experiment was designed to provide an idea of what class of surface measures might be used to represent phenomenal feature locations on a 3-D surface.

At the grossest level, these geometric properties can be divided into two basic groups: those that depend on the position of the viewer relative to the object and those that are viewpoint invariant. Measures of the first type will vary as a result of the projection $\mathbb{R}^3 \rightarrow \mathbb{R}^2$ that occurs when 3-D objects in the world are projected onto the retina. That is, they will vary with the direction from which the object is viewed or the orientation of the object relative to the observer, whereas the descriptions of the second type remain constant over transformations of the viewer or the object.

At a gross level of analysis, there are characteristics that mostly remain constant over viewpoint and object transformations, such as the identity of the object, its color or texture, the relative location of “parts” on rigid objects, and so on. Similarly, there are characteristics that *will* vary with viewpoint, such as absolute and relative depth from the observer, ordinal and metric position of individual locations on the object, and so on. Any useful perceptual representation would most likely have to take both sorts into consideration, due to ambiguities in the sort of information provided. For example, in a representation of the Moore sculpture shown in Figure 2, we could use a viewpoint-dependent feature, such as *b*, or viewpoint-independent features, such as *a*, *c*, and *d*.

In this experiment, the nature of these landmark locations was examined: When we identify locations on objects, do we tend to do this in a viewer-centered or a viewpoint-independent way? More specifically, do viewer-centered

or viewpoint-independent measures best describe our perceptual representations of these locations?

Method

In this experiment, observers used a simple marking paradigm to indicate the local surface depth maxima and minima regions for a nontrivial 3-D object from a viewer-centered frame of reference. By presenting the object in various orientations relative to the observer, we can then compare the marked extrema with the actual minima and maxima in the various orientations as a measure of their accuracy.

Observers. The observers consisted of 4 adults, the authors and 1 additional laboratory member. All were aware of the purpose of the experiment and had normal or corrected-to-normal vision. All the observers had extensive experience with the concepts under investigation (viewer-centered depth and curvature extrema) and had served as observers in numerous other experiments dealing with these same concepts.

Stimuli. The experimental *probe surface* stimuli were similar to the turbulent surface patches used in our past research (Phillips & Todd, 1996; Phillips et al., 1997). Simply put, these surfaces are smoothly varying, self-similar, “bumpy” surfaces. A grid of bumps are created at some given spatial frequency, and additional grids are superimposed at various scales to yield the self-similarity. The nominal height of a given bump is determined by a random function, and each bump is smoothly interpolated with its neighboring bumps. The superposition of bumps of various (typically harmonic) scales yield our final stimuli.

More specifically, these surfaces are of the variety $z = f(x, y)$, a height function over the (x, y) plane—so called *Monge surfaces*. For these experiments, $f(x, y)$ is a two-dimensional, smoothly varying random wave function. By summing n octaves of this function, we end up with a class of self-similar surfaces often used to simulate natural phenomena, such as mountains, marble, fire, and clouds, in photo-realistic computer graphics images (see Peachey, 1985, and Perlin, 1985, for further details). Indeed, there is significant evidence that many natural objects possess this type of self-similarity (Mandelbrot, 1983; Thompson, 1992).

Equation 1 shows the specific function used in these experiments, where n is the number of octaves and f the frequency of the random wave function:

$$f(x, y) = \sum^n \frac{\text{noise}(x, y)}{f^n}. \quad (1)$$

The *noise* function is defined as follows. First, a two-dimensional discrete lattice of uniformly distributed random numbers is defined.




Figure 3. A depth map and surface plot of the first of two turbulent probe surfaces used in these experiments.




Figure 4. A depth map and surface plot for the second probe surface stimulus used in these experiments.

The distance from point to point in this lattice forms the underlying size of the lowest frequency component of the resulting surface (i.e., the carrier frequency for the rest of the noise). In order to obtain a value at a given real-valued (x,y) location, a bicubic interpolation is performed between the neighboring lattice points. This yields a smooth, differentiable function that is then summed over n octaves, according to Equation 1, to create the probe surfaces for the stimuli.

The base wavelength of the noise function was chosen so that the resulting probe surface had features (bumps and dimples) whose size subtended roughly 2° of visual angle when viewed frontoparallel. Two octaves were added, contributing 1° and 0.5° features, respectively. The resulting summation yielded a bumpy surface somewhat reminiscent of a mountainous landscape. An example of the stimuli used can be seen in Figures 3 and 4. This surface was mathematically infinite in extent in both x and y , and an arbitrary (x,y) position on the surface was chosen as the center of the test region used in this and subsequent experiments. For those interested in further details, a rigorous explanation of the generation and analysis of these stimuli is available in Phillips (in press).

Finally, the probe surface was smooth shaded and textured, using a uniform random checkerboard pattern of gray levels. The texture was applied so that equal areas on the surface had approximately equal areas of texture map associated with them. On each presentation, the texture was randomly displaced in (x,y) so that the texture as a whole could be used as a cue but no specific location on the texture could be used as a visual anchor point.

Procedure. Before the start of the experiment, all the observers discussed and agreed on the phenomenal geometric definition of viewer-centered depth, minima and maxima, as well as confirming their presence in a set of test images. All agreed that the extrema were point singularities that existed relative to the observers' line of sight.

Each observer was presented with one of the two probe surface stimuli in one of three possible orientations: frontoparallel (the 0 condition), slanted backward, away from the observer, 20° (the $-$ condition), or slanted forward, toward the observer, 20° (the $+$ condition). The observer then used the mouse to manipulate a monocularly presented cursor to locate all of the local depth minima and maxima present on the probe surface. The cursor was unconstrained and, therefore, free to move to any location on the probe surface. When the observer located a feature point at a desired location, a key on the keyboard was pressed to indicate the type of point located under the cursor. Each location selected by the observer was then marked, using a small dot, in order to eliminate the chance that a location would be classified more than once.

All the observers marked locations on both shapes at all three orientations twice per condition, resulting in a 2×3 design. Each session consisted of one of the six possible conditions in which the ob-

server marked as many or as few features as they interpreted as being present. No time limit was placed on the session, and no specific ordering of marking (maxima before minima, topmost before bottommost, etc.) was imposed. The probe surfaces and orientations were randomized across trials and across observers in order to reduce bias that might arise as a result of learning any one of the surfaces in a particular orientation.

Apparatus. All the stimuli utilized in this experiment were created and presented using QuickDraw 3D v 1.5.1 and Sprockets v 1.0 on an Apple Power Macintosh 9600/200 workstation, which was expanded to utilize a second monitor for stereo presentation. A stereoscope was constructed using off-the-shelf optical components. Four first-surface mirrors were mounted on individual adjustable stages to allow for the adjustment of each individual's interpupillary distance. The entire mirror assembly was mounted to an optical bench strip that was, in turn, mounted to the table holding the monitors. Each monitor was a 15-in. Sony Trinitron, color matched using a Light-source ColorTron 32-band spectrophotometer. All the images were displayed at $1,024 \times 768$ resolution at 75 Hz in 32-bit color depth.

A chinrest, mounted on a separate table to avoid and isolate unwanted vibration, was utilized to steady the observer's head and maintain a constant viewing distance. The stimuli were approximately 16.5 cm in size, viewed at an effective distance of 86 cm, which resulted in stimuli that subtended roughly 11.4° of visual angle ($1^\circ \approx 1.5$ cm). All the trials took place with dim room lighting.

Prior to each session, a Nonius image was displayed with Vernier markings on the monitors for the two eyes. These marks were brought into alignment by the observer via adjustment of the monitors' geometry and the stereoscope's mirrors, resulting in an accurate, limited distortion stereographic image. Finally, the observers' response was obtained using a standard keyboard and mouse, located on a third table to avoid unwanted vibrations to the optical assembly or chinrest.

Results and Discussion

Because the observers were instructed to indicate all of the minima and maxima, using a viewer-centered frame of reference, we would expect there to be systematic variations in the observers' marking for each presentation condition, since the extrema changed relative to the viewer for each presentation condition. Figure 5 illustrates the viewer-centered depth maps for the frontoparallel and $+20^\circ$ slanted conditions.

We designed a local extrema finder, using Mathematica (Wolfram, 1991), that systematically located all of the local




Figure 5. Viewer-centered depth maps for Surface 1 (see Figure 3) from two viewing conditions: on the left, the frontoparallel viewing condition; on the right, the +20° slanted condition. White areas in the contour plot are closer to the observer. In the surface plot, the line of sight is along the z (vertical) axis. Note how many minima and maxima move or disappear entirely when the probe surface is tilted away from frontoparallel.

depth minima and maxima on the probe surface for each of the orientations relative to the viewer. The results of this search are shown in Figure 6.

The observers were free to mark any number of locations as extrema, and this presented a minor challenge in analysis. If there *were* direct one-to-one relationships between the marked and the actual extrema, it would be a simple matter to compute the euclidean distance offset as a measure of accuracy. To analyze performance in this experiment, a correlation technique was devised to measure a goodness of fit between the observed and the actual extrema. In our analysis, two surfaces were constructed, the first being the ground truth (actual extrema of the probe surface; hereafter, the *measure surface*) and the second de-

rived from the extrema indicated by the observers (hereafter, the *response surface*). The resulting surfaces were correlated, yielding a measure of goodness of fit between them. This solved the correspondence problem presented above.

The construction of the measure and response surfaces was done as follows. Initially, each surface was a plane the size of the probe area of the probe surface. On each surface, a unit-height Gaussian “bump” was added at the location of an actual or a judged extrema, respectively. The diameter of this bump was based on our previous research (Phillips & Todd, 1996), which investigated the perceptual salience of classes of geometric structures—most notably, bumps and dimples on which the local extrema sit in these experiments. In those experiments, we found that

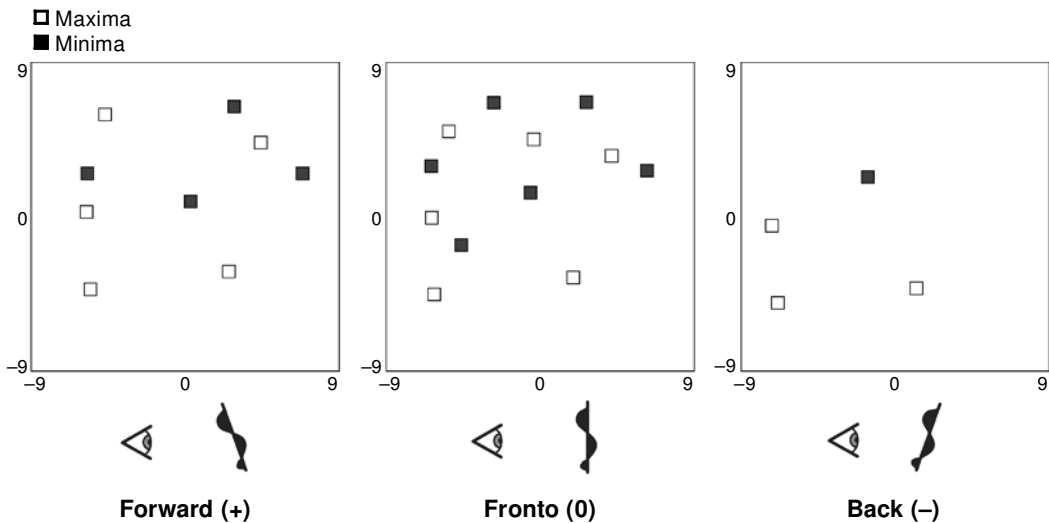


Figure 6. Locations of actual depth extrema on the probe surface for each viewing condition. Light squares represent maxima (peaks), whereas dark squares represent minima (dips) on the probe surface, relative to the observer. Note that several extrema move, appear, or disappear as the orientation of the probe surface changes. Units are in centimeters.




Figure 7. A measure surface generated from the actual probe surface extrema at each viewing condition. White areas represent peaks, whereas dark areas represent troughs. Areas of concentrated shading represent more steep peaks or troughs, whereas broader areas represent flatter regions. Units are in centimeters.

bumps and dimples that subtended approximately 2° of visual angle were optimally perceived when embedded in the same family of probe surfaces as those used here. To create the final maps, negative bumps (i.e., “dimples”) were added to the surface at the location of each minima, and positive bumps were added at the location of each maxima.

This process was executed separately for the actual and the observer-indicated extrema, resulting in two maps. One desirable side effect of this method of creating the maps was that the local height of the resulting map was scaled on the basis of the surrounding terrain. This resulted in a diffusion of the surface in areas that had small gradients of orientation (i.e., flat areas) and a sharpening of the surface in regions with high gradients (i.e., sharp peaks). The resulting measure surface construction for one of our stimuli is shown in Figure 7, and a construction of 1 observer’s response surfaces is shown in Figure 8. This technique for constructing implicit surfaces is similar in spirit to the techniques used by Blinn (1982) and expanded by others in the field of computer graphics.

In Figure 9 we show the observers’ responses at each orientation of the probe surface relative to the observer, along with the actual extrema. Critically, unlike the actual ex-

rema, the observers’ markings were very similar across orientations. Since we know that the actual extrema changed location or existence at each orientation, this suggests that the observers were not performing the task as specified. Thus, the observers’ markings seemed to be consistent with the extrema in the frontoparallel presentation condition, further suggesting a viewer-independent representation that utilized a fiducial frame of reference consistent with the probe surface’s height-field. This finding is consistent with our previous results (Phillips et al., 1997), wherein the observers appeared to be marking locations on a reference surface consistent with a particular privileged representation of the surface. In this experiment, as well as in our previous one, there was evidence that this frame of reference was the global ground plane of this terrain-like surface.

The R^2 between the response and the measure surfaces were as follows: forward slant (+), $R^2 = .46$; frontoparallel (0), $R^2 = .58$; backward slant (-), $R^2 = .35$. These results support the notion that the observers were not performing the task as requested—that is, they were apparently *not* using a viewer-centric frame of reference. The frontoparallel condition demonstrated acceptable performance, but the slanted presentations and adjustments were not as




Figure 8. Response surface generated from an observer’s markings for the location of local minima and maxima in each observation condition. Light areas represent marked maxima, whereas dark areas represent troughs. Areas of concentrated shading represent tighter clustering in responses, whereas more diffuse regions represent more variable responses. Clustered responses result in a steeper, narrower bump or dip, whereas diffuse responses result in a lower, more diffuse bump or dip. Units are in centimeters.

strong. A visual examination of the results show that the observers seemed to be marking the same locations regardless of the orientation of the probe surface. Table 1 shows that there was a high correlation between the judged extrema across the orientation conditions. The correlations between the *actual* extrema across the conditions were far weaker. For example, in Table 1, there is a very low association ($R^2 \approx .17$) between the actual extrema in the forward-slant (+) condition and the backward-slant (–) condition. However, there is a relatively good amount of agreement in the *judged* data, with an $R^2 \approx .59$. This relatively strong agreement carries across the orientation conditions consistently, whereas the actual agreement varies significantly. Since the observers appear to have been judging the same extrema regardless of orientation condition, this leads us to wonder *which* extrema or even what characteristic of the probe surface was actually being marked. Their judgments were far more consistent between conditions than would be suggested by the actual extrema. Whatever was being marked was, apparently, viewpoint independent for these surfaces.

These results suggest that, at least for the orientations utilized in this experiment, the observers appear to have been performing the identification task on the probe surface itself, rather than in the image—that is, using a viewer-independent rather than a viewer-centered frame

of reference. The observers were instructed to utilize a viewer-centered frame of reference to mark the extrema on our surfaces, and therefore, we would have expected that there would have been some systematic variation in the depth markings for the different orientation conditions; however, no such variation appeared. Similar results were obtained in a previous set of experiments (Phillips et al., 1997) that required identifying arbitrarily presented locations on a surface that was randomly reoriented relative to

Table 1
 R^2 Between All Pairs of Orientations for Actual and Judged Extrema

Slant Conditions	R^2	
	Judged Extrema	Actual Extrema
(+,0)	.6561	.5929
(+,-)	.5625	.1681
(0,-)	.7056	.2601

Note—The similar relationships between the observers’ settings and the large differences between the actual extrema belie the suggestion that the observers are indicating something other than the actual viewpoint-dependent extrema. Whatever they are marking is similar regardless of orientation of the probe surface relative to the observer. In the forward versus frontoparallel (+, 0) condition, there is a reasonably high R^2 between the locations of the extrema, but in the other conditions [(+, –) and (0, –)], this is not the case. The correspondence of the judged extrema is consistently high across all conditions.

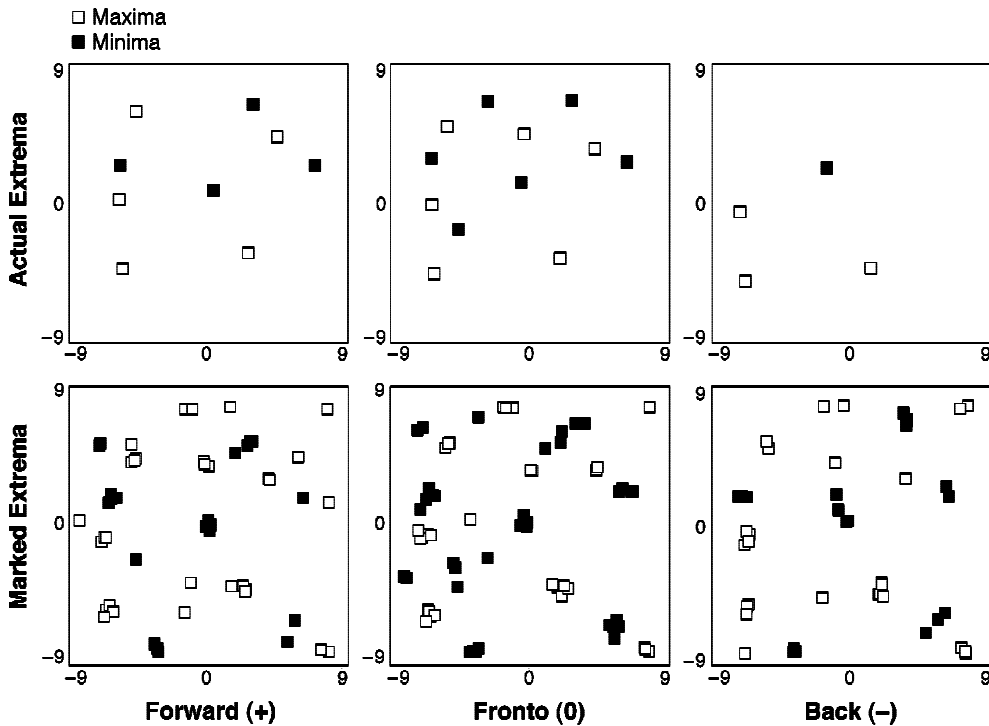


Figure 9. Actual extrema (top) and a composite of all the observers' markings (bottom) of the locations of local minima and maxima in each observation condition. Light squares represent marked maxima (peaks), whereas dark areas represent minima (pits). Note that, unlike the actual extrema, the observers' responses were very similar across viewing conditions; therefore, the observers were not performing the task as instructed. Thus, the observers' markings seemed to be consistent with the extrema in the frontoparallel presentation condition, further suggesting a privileged, viewer-independent representation of the surface.

the observer. The observers in that experiment did not exhibit any systematic bias as a function of the orientation of the surface, which agrees largely with the results of the present experiment.

Our main goal in this experiment was to determine whether observers would be able to make use of the structural information available to make extrema judgments from a specific frame of reference. If our representation of surface structure in this case were mainly viewer centered, we would expect variation in responses across the presented orientations. The results of the present experiment suggest a representation consistent with landmarks defined by viewer-independent features. Since these results suggest the possibility of a privileged frame of reference when locations on an object are identified, we should now consider what underlying geometric properties might give rise to the information used to frame this coordinate system.

EXPERIMENT 2

In Experiment 1, the observers were asked to identify local point singularities in surface structure, such as the peaks of mountains, where viewer-centered depth is at a maximum (or minimum) in all directions. Experiment 2, in contrast, was designed to investigate the perceptions of line singularities, such as ridges or valleys, where an under-

lying dimension of surface structure is at a maximum (or minimum) in all directions except one. One way of thinking about these two types of surface landmarks is that the point singularities are nodes on a graph and the line singularities are the edges that connect them. For example, in creating a city map, it is common to represent streets as lines (i.e., the edges of a graph) and individual buildings or intersections as points (i.e., nodes). A similar representation is also possible for smoothly curved surfaces, in which the peaks of mountains are represented as nodes and the ridge lines that connect them as edges. There is significant empirical evidence in the area of environmental psychology that humans represent their physical environment in a manner that is similar to a graph structure (cf. Lynch, 1960; Nasar, 1998).

Because the results of Experiment 1 had shown that observers could not successfully adopt a viewer-centered

Table 2
Results From the Constrained Judgment Task
to Test the Reliability of Observers' Markings

Observer	\bar{x}_σ	σ_σ
A.K.	0.190	0.100
J.K.	0.163	0.100
F.P.	0.215	0.142

Note—Units are in centimeters. All 3 observers' responses were tightly clustered within ≈ 0.2 cm (0.13° visual angle).




Figure 10. An example of the probe surface used in Experiment 2 and the marking of linear structures by Observer F.P. In this task, the observer was instructed to mark the phenomenal ridge lines and valleys of the surface, using as many locations as necessary to fully describe the structures, including their endpoints.

frame of reference in making judgments about local depth extrema, we chose, in the present study, to let them decide what constituted a ridge or a valley without providing explicit instructions about how those concepts should be perceptually defined. In an effort to determine whether observers could perform reliably with such vague instructions, a pilot experiment was conducted in which observers viewed a stereogram of a smoothly curved surface and were asked to draw the pattern of ridges and valleys onto a sheet of graph paper. A visual analysis of these drawings revealed a high level of agreement across different observers, thus suggesting that instructing them to mark ridges and valleys is a perceptually meaningful task.

Experiment 2 had three specific goals. First, we wanted to measure the precision with which observers can identify the locations of ridges and valleys on a smoothly curved surface at different orientations in depth. Second,

we wanted to assess the consistency of these judgments across different observers. Finally, we also wanted to determine the underlying dimensions of differential structure (e.g., depth, orientation, or curvature) by which these landmarks are perceptually defined.

Method

Procedure. The same surfaces and presentation conditions as those in Experiment 1 were used in this experiment—two different surfaces presented with their orientation frontoparallel and $\pm 20^\circ$ relative to the viewer. The observers were instructed to mark the phenomenal ridge lines and valleys of the surfaces. As with the first experiment, the stimuli were presented stereoscopically with texture and shading. A monocular cursor was manipulated with the mouse, which the observers used to mark as many points on the ridge and valley structures as they felt necessary to define the extent and shape of the structure. There was no constraint on the position of the cursor, and the ridges and valleys could be marked in any order. The stimuli, apparatus, and observers were the same as those in Experiment 1.

In order to obtain a more fine-grained measure of the precision of these judgments, 3 of the observers (A.K., J.K., and F.P.) also performed a more constrained version of the same task, in which the cursor could be moved only vertically along a single scan line. Ten scan lines were presented for each surface at each of the three possible orientations in depth, and each scan line was repeated on three separate trials. The observers were instructed to mark all the points at which the scan line crossed a ridge or a valley.

Results and Discussion

Constrained judgments. Let us first consider the results obtained on the constrained judgment task, in order to evaluate the overall reliability and precision of the perceived locations of ridges and valleys. It is important to keep in mind that, for a given surface, the locus of surface points along a vertical scan line was the same for all three possible orientations in depth. This made it possible to measure any subtle changes in the apparent locations of ridge lines that may have occurred as a function of changing viewing direction. As is consistent with the findings from Experiment 1, however, the observers' judgments re-




Figure 11. Construction of the response surface. Shown on the left is the composite marking data for all the observers. A Gaussian bump is moved across each separate, cubic-interpolated path, and these bumps are summed to create the surface seen on the right. Both the contour plot and the surface plot show broader and shallower areas, where there is less concentration in the observers' markings, and higher sharper areas, where there is a large amount of consistency and, thus, overlap.

Table 3
 R^2 Between Individual Observers' Response Surfaces

	Observer			
	A.K.	F.P.	J.K.	V.P.
F.P.	.664			
J.K.	.765	.706		
V.P.	.635	.781	.691	
All	.856	.894	.910	.870

Note—The last row contains correlations of each observer with the mean responses of all observers. There is reasonably strong agreement between observers, with variance likely being due to the constraints of interobserver reliability and differences in criterion for the endpoints of the linear structures.

mained remarkably stable over the different surface orientations, which spanned a 40° range.

In order to measure the spread of these judgments, a clustering analysis was performed to isolate individual ridge points and to measure the variance in the vertical positions of the observers' markings for each one. The results of this analysis are shown in Table 2, which shows the average variance for individual ridge points, collapsed over orientations, and the variations in variance among different ridge points. These results show that the observers were able to mark the positions of ridges and valleys within ≈ 0.2 cm (0.13° of visual angle). This is a remarkably high degree of reliability for localizing these structures, given that they were not explicitly defined by the instructions and that the surfaces were presented over a large range of possible orientations.

Unconstrained judgments. For the second part of Experiment 2, the observers' markings were not constrained. An example consisting of a surface and the marked structures of one session of Observer F.P. are shown in Figure 10. A similar analysis was used as in Experiment 1, based on the correlation between a reference surface and a judged surface.

The response surfaces were constructed as follows. For a given set of ridge or valley markings by a given observer, a smooth cubic interpolating spline (à la Bezier) was fit through all the marked points. A linear Gaussian bump or dip was then added to the surface along this fitted curve, depending on the type of structure marked, resulting in a 3-D linear structure on the surface. The bump's width and height were selected in the same way as in Experiment 1, a choice further reinforced by the results of the pilot experiment above. Multiple markings in the same area reinforced and, thus, heightened the bumps, whereas stray marks or marks made a few number of times resulted in lower structures. Similarly, the clustering of the marks controlled the width of the structure: Tightly clustered marks resulted in a narrow and taller bump or dip, whereas more dispersed markings, indicative of less agreement on the position of the feature in question, resulted in a lower and more dispersed bump or dip on the constructed surface. Repeating this for each observer's markings resulted in a response surface consisting of all of the locations marked as ridges or valleys. As in Experiment 1, this technique is

similar to that of Blinn (1982) and others for creating implicit surfaces.

Response surfaces were constructed for each observer in each of the presentation conditions (+, 0, or $-$, as in Experiment 1). Figure 11 illustrates the construction of a surface, using the responses for all the observers.

Interobserver consistency. In our pilot experiment outlined above, we determined the reliability within a given observer and found it to be high. We compared the results across observers as well, giving a measure of reliability between observers as well. To obtain this measurement, the response surfaces for each observer were correlated with those for all the other observers and the summary surface for all the observers. Table 3 shows the results of this analysis.

The correlations show a reasonably strong agreement between the observers' markings. The remaining variability is most likely due to the constraints of the interobserver reliability and the difference in criterion for the endpoints of the linear structures. For example, a piece of paper with a strong crease at one end that lessens over the length of the paper (see Figure 12) will create a gradually vanishing ridge structure whose presence or absence is subject to some threshold-like psychophysical function.

Underlying geometry. For this experiment, the task did not explicitly specify which geometric structures were to be marked. As a result, there is no "ground-truth" to compare the observations against. Instead, we need to compare them with various geometric surface properties in an attempt to discover which underlying structures might be providing the phenomenal information that leads to the markings. Responses based on viewer-centered features would, therefore, correlate with viewer-centered measures, such as depth. Viewer-independent responses would correspond to features that do not change with viewpoint, such as the gradient or curvatures of the surface.

At any location, P , on a surface, an infinite number of plane curves exist containing the normal vector, N_P there (so-called *normal curves*). Each of these curves can be assigned a measure of curvature, κ , at P . Two of these κ s are of primary interest, the minimum and maximum curvatures, commonly known as the *principal curvatures*. These curvatures, denoted κ_1 and κ_2 (sometimes κ_{\min} and κ_{\max}), have the additional characteristic that the curves that contain them are always orthogonal. The tangents of these curves, which define the tangent plane at P , along with




Figure 12. An example of the criterion problem in determining the beginning and ending of the linear structures. In this example, a surface strongly creased at one end and flat on the other creates a ridge that gradually fades away. Exactly where it disappears is variable on the basis of the observer's criterion.




Figure 13. Curvatures and several example reparameterizations. On the left are the two separate curvature values for Surface 1 used in our experiments. These curvatures, κ_1 and κ_2 , can be combined into other measures that reflect some curvature-based aspect of the surface geometry, as outlined above. For example, shape index and curvedness (S and C) represent scale-independent and scale-dependent aspects of the underlying shape, respectively. Mean and Gaussian curvature (H and K) show the average curvature and the product of the curvatures at a given point, the sign being a useful indicator of the class (but not the direction) of curvature in the latter case. Finally, the curvature contrast ($\Delta_{|\kappa|}$) shows a slight modification of mean and Gaussian curvature that captures the absolute local difference in curvature at a given surface location.

the normal, define a local orthogonal frame. For our purposes, it is important to note that these curvatures are intrinsic to the surface and are unchanged by viewpoint.

Traditionally, the two curvatures are combined into other measures that are useful for diagnosing the nature of the surface at a given location. Two of the most common are the *Gaussian curvature*,

$$K = \kappa_1 \kappa_2, \quad (2)$$

and the *mean curvature*,

$$H = \frac{\kappa_1 + \kappa_2}{2}. \quad (3)$$

A more formal description of curvature and the parameterizations of them can be found in the Appendix.

For the analysis of the observers' response data, a set of measure surfaces were constructed from the depth, gradient, and curvature information of the probe surfaces. These surfaces were then compared with the observers' response surfaces to find a best-fitting match. For each measure, the reference surface was constructed where z represented the magnitude of the measure at the location (x, y) of the probe surface. For example, for Gaussian curvature, the resulting measure surface was of the form $z = K_{x,y}$, where K is measured at (x, y) on the probe surface. A depth-map measure surface relative to the viewer can be constructed as $z = Trans[f(x, y)]$ where $Trans$ represents the viewpoint transformation of the surface location, the surface defined by $f(x, y)$. Figure 13 shows maps for some of the curvature-based, viewer-independent measures.

Table 4
Results of Experiment 2 by Observer: R^2 of the Best-Performing Measures, Averaged Across All Presentation Orientations

Observer	Measure			
	H	S	$\Delta_{ \kappa }$	$\kappa_{ \max }$
A.K.	.405	.354	.397	.387
F.P.	.399	.347	.431	.412
J.K.	.472	.447	.479	.480
V.P.	.394	.337	.403	.390

Note—The measures show a strong consistency across observers. Highest R^2 s occur with mean curvature (H), shape index (S), curvature contrast ($\Delta_{|\kappa|}$), and the signed, maximum magnitude of the principal curvatures ($\kappa_{|\max|}$).

The results from Experiment 1 suggest that the point-features were identified in a viewpoint-independent manner; therefore, we chose to correlate the observers' markings with several of the curvature reparameterizations mentioned above, since these parameterizations are intrinsic to the surface—that is, they do not change with viewpoint. Specifically, we utilized the Gaussian and mean curvatures (K and H), shape index and curvedness (S and C), gradient magnitude ($|\nabla|$), total curvature magnitude ($\sum_{|\kappa|}$), curvature contrast ($\Delta_{|\kappa|}$), and signed curvature difference (Δ_{κ}). In addition, we looked at the signed maximum and the unsigned maximum of the two principal curvatures, $\kappa_{|\max|}$ and $|\kappa_{|\max|}$ (see the Appendix for details on these measures). For completeness, we also correlated the resulting markings with the viewpoint-dependent depth map.

In Table 4 and Table 5, the results are shown for each observer, partitioned into two sets by the strength of correlation. Within each measure, all the subjects were in general agreement, as were the measures between subjects. Across all observers and measures in the best-fitting cases, $R^2 \approx .4$, whereas with the poorly fitting measures, the R^2 remains close to 0 in all cases. The depth map shows little correlation as well. Even when assuming a frontoparallel depth map across all viewing conditions, these correlations are quite small with R^2 well below .05 in all cases (see the first column of Table 5 for details).

Since the observers' results were highly consistent, we also performed the correlations by orientation collapsed across observers. As would be expected, the results for the viewpoint-independent measures were consistent with the per-observer findings, with mean curvature (H), shape index

(S), curvature contrast ($\Delta_{|\kappa|}$), and signed maximum curvature magnitude ($\kappa_{|\max|}$) showing the most variance accounted for in Table 6. As with Experiment 1, there was little difference between orientations, although the backward-slanting condition showed less strength than the others. We suspect this may be explained by a slight criterion shift by all the observers due to the small change in the total amount of self-occlusions in this case (although we chose the range of orientations specifically to avoid wholesale self-occlusion of depth extrema).

The results from Experiment 2 generalize the findings of Experiment 1 from point structures to linear structures. The observers marked regions of the geometry where there were extremal values of intrinsic and, therefore, viewpoint-independent curvature and curvature-related measures. When representations that would support this sort of invariance are considered, the previously described graph structures can accommodate these findings nicely. It does, however, remain to be seen what the best constitution of this structure should be. It is reasonable to consider the “bump” and “dip” structures from Experiment 1 as nodes of such a structure, but the extended structures of Experiment 2 are more ambiguous. Does a ridge serve a *connecting* purpose, perhaps of two other geometric nodal locations, or does the ridge itself constitute a node?

GENERAL DISCUSSION

The research described in the present article was designed to investigate the perception of landmarks on smoothly curved surfaces and to determine the relative stability of these landmarks over different viewing directions. This research was motivated, in part, by an earlier experiment (Phillips et al., 1997), in which observers reported that surface landmarks, such as hills, valleys, and ridges, were used to determine the point-to-point correspondence relations across multiple vantage points. Our working hypothesis, as we began these experiments, was that the perceived landmarks on a surface would be located at local extrema (i.e., maxima or minima) of some underlying geometric property of the surface shape (e.g., depth, slant, curvedness, etc.). Our initial experimental strategy, therefore, was to measure the precision with which observers could identify local extrema at varying levels of differential structure.

Table 5
Results of Experiment 2 by Observer: R^2 of the Best-Performing Measures, Averaged Across All Presentation Orientations

Observer	Measure						
	Depth	K	C	$ \nabla $	$\sum_{ \kappa }$	Δ_{κ}	$\kappa_{ \max }$
A.K.	.048	.011	.000	.018	.000	.003	.000
F.P.	.025	.007	.000	.000	.000	.002	.000
J.K.	.0477	.009	.000	.007	.000	.004	.000
V.P.	.0552	.017	.002	.004	.004	.002	.001

Note—The measures show a strong consistency across observers. Absolute depth, Gaussian curvature (K), curvedness (C), gradient magnitude ($|\nabla|$), total curvature magnitude ($\sum_{|\kappa|}$), signed curvature difference (Δ_{κ}), and the unsigned principal curvature maximum ($\kappa_{|\max|}$) show little or no agreement with the observers' markings.

Table 6
 R^2 for Each of the Best-Performing Measures,
Collapsed Across Observers, for Each Orientation

Orientation	Measures			
	H	S	$\Delta_{ c }$	$\kappa_{[\max]}$
Frontoparallel (0)	.607	.570	.613	.603
Forward (+)	.612	.579	.603	.603
Backward (-)	.546	.514	.568	.558

Note—In all cases, R^2 is highest in the frontoparallel case, indicating that the observers were performing the task without regard for the global orientation of the surface.

In Experiment 1, the observers were instructed to identify the local depth maxima and minima on a stereoscopically presented surface in different orientations. Because depth extrema vary as a function of viewing direction, the accurate performance of this task would produce significant differences between the marked locations in the different orientation conditions. That is not what occurred, however. When we examined the distributions of observers' responses, the perceived surface extrema in the different orientations were all highly correlated with one another. That is to say, the selected landmarks were all viewpoint invariant. Apparently, the observers were unable to adopt a viewpoint-dependent frame of reference to identify local extrema, even though they were specifically instructed to do so. Whereas Experiment 1 was concerned with the zero-dimensional landmark points on visible surfaces, Experiment 2 was designed to investigate the higher order one-dimensional structures that are commonly referred to as ridges and valleys. Given the results from Experiment 1, we decided not to provide the observers with explicit instructions about how these structures are geometrically defined. All they were told was to mark the tops of the ridges and the bottoms of the valleys, and the precise definitions of these terms were left to their own intuitions.

Despite the vagueness of these instructions, the observers were able to perform these tasks with a high degree of confidence, and the overall pattern of their responses was remarkably consistent. As in Experiment 1, the landmarks they selected were minimally affected by changes in surface orientation, thus indicating that their judgments were based on some property of the surface that was viewpoint invariant. In an effort to reveal the specific attribute of surface structure, we correlated their pattern of responses with a wide variety of geometric measures at varying levels of differential structure.

There have been a number of previous suggestions in the literature that perceptually salient landmarks on a surface are likely to be located at curvature extrema. This hypothesis was first proposed by Hoffman and Richards (1984) for the perception of surface part boundaries—based on a mathematical constraint that when two arbitrary shapes are made to interpenetrate one another, they must always meet at a discontinuity of negative curvature (see also Hoffman & Singh, 1997; Richards et al., 1987). A similar idea was later proposed by Phillips et al. (1997) in order to account for the ability of observers to make point-

to-point correspondence judgments over multiple vantage points. The results of the present experiments are, in many ways, supportive of these hypotheses. The judged landmarks showed a high degree of viewpoint invariance; they were poorly correlated with the patterns of extrema in depth or orientation, and they were highly correlated with the patterns of extrema for several measures of surface curvature.

There is, however, some remaining uncertainty in these data about how curvature is perceptually parameterized. Note in Tables 4 and 5 that some measures of curvature correlate quite highly with the observers' judgments, whereas others do not. Our initial intuition for this task was that perceived landmarks would most likely be located at extrema of curvedness, but the data did not support that prediction. Indeed, the correlations of curvedness with the observers' judgments were lower than those for any other measure of curvature we investigated. Because many of these measures are correlated with one another, it is difficult from the present data to draw any strong conclusions about what specific aspects of curvature are most salient for human perception. Of the measures we examined, mean curvature, curvature contrast, shape index, and signed maximum curvature magnitude were the best predictors of the observers' judgments, but additional research is obviously needed before any strong conclusions are drawn about the status of those measures as perceptual primitives.

The results of the present experiments indicate that randomly generated smoothly curved surfaces contain perceptually salient landmarks that have a high degree of viewpoint invariance (see also Phillips et al., 1997). These landmarks include point singularities, such as the apex of a mountain or the nadir of a crater, but they also include extended line structures, such as ridges and valleys. It is especially interesting that these landmarks exhibit a high degree of stability over changing viewing directions, whereas judgments of more generic local surface properties typically do not (e.g., Norman, Todd, Perotti, & Tittle,



Figure 14. An artist's depiction of one of the surfaces used in our experiments, using only line elements (i.e., no shading). The upper panels show the lines apparently used to denote occlusion and boundary contours, respectively. The lower panel is the original drawing. The individual contour drawings do not result in a satisfying three-dimensional percept; both sets of contours are required.

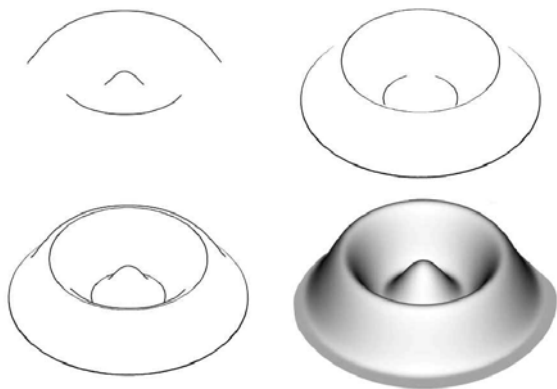


Figure 15. A formal example of the type of illustration shown in Figure 14, using a radial cosine surface. The top panels show the occlusion and boundary contours separately. The bottom left shows them integrated, and the bottom right shows the surface as represented using a contour texture. As with the artist-produced example, the collections of individual occlusion or boundary contours do not result in a compelling three-dimensional object, but their conjunction does.

1996). This suggests that the topological arrangement of surface landmarks could be a representational anchor that underlies our phenomenal experience of shape constancy. In a formal sense, these topological relations are most easily characterized as graph structures, which are closely related to those that have been studied in other perceptual contexts, such as the patterns of edges and vertexes on polyhedral objects (Clowes, 1971; Draper, 1981; Huffman, 1977; Mackworth, 1973; Malik, 1987; Waltz, 1975a, 1975b) or the patterns of occlusion contours on smooth surfaces (Hoffman & Richards, 1984; Richards et al., 1987).

An interesting issue for future research is to examine the importance of curvature ridge lines for the pictorial representation of smoothly curved surfaces. It is important to keep in mind that in mechanical drawings of polyhedral objects, the lines are typically drawn to denote the location's occlusion contours or the edges that connect planar facets—which are extreme examples of curvature ridges. Could a similar denotation scheme be effective for the pictorial depiction of more smoothly curved surfaces? In an effort to address this question, we asked an artist to draw a depiction of one of our experimental surfaces without using any shading. The results of her efforts are shown in the bottom panel of Figure 14. Note that it includes an outline of the perimeter of the surface patch and also some internal lines to depict the ridges and valleys. These two types of line structures are shown separately in the top panels to demonstrate how both are needed for a perceptually convincing depiction.

Another more formal example of this same phenomenon is shown in Figure 15. The lower right panel of this figure shows a radial cosine surface that is depicted with shading. The occlusion contours for this surface are presented in the upper left panel, and its curvature ridge lines are presented in the upper right panel. Note that neither of

these patterns presented in isolation provides a perceptually convincing pictorial representation of the 3-D surface structure. However, when both types of contours are combined in the lower left panel, they provide sufficient information to specify the surface shape. Although we have observed similar effects for other types of surfaces used in our investigations, we have not attempted to systematically compare the perceptual effectiveness of curvature ridge lines with that of other types of surface structures, such as depth or orientation ridges. This could perhaps be a useful strategy for identifying the specific attributes of local surface (or image) structure by which surface landmarks are perceptually defined.

REFERENCES

- BIEDERMAN, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, **94**, 115-117.
- BIEDERMAN, I., & JU, G. (1988). Surface versus edge-base determinants of visual recognition. *Cognitive Psychology*, **20**, 38-64.
- BLINN, J. (1982). A generalization of algebraic surface drawing. *ACM Transactions on Graphics*, **1**, 235-256.
- CLOWES, M. B. (1971). On seeing things. *Artificial Intelligence*, **2**, 79-116.
- DRAPER, S. W. (1981). The use of gradient and dual space in line drawing interpretation. *Artificial Intelligence*, **17**, 461-508.
- HOFFMAN, D. D., & RICHARDS, W. A. (1984). Parts of recognition. *Cognition*, **18**, 65-96.
- HOFFMAN, D. D., & SINGH, M. (1997). Saliency of visual parts. *Cognition*, **63**, 29-78.
- HUFFMAN, D. A. (1977). Realizable configurations of lines in pictures of polyhedra. *Machine Intelligence*, **8**, 493-509.
- KOENDERINK, J. J. (1984). What does the occluding contour tell us about solid shape? *Perception*, **13**, 321-330.
- KOENDERINK, J. J. (1990). *Solid shape*. Cambridge, MA: MIT Press.
- KOENDERINK, J. J., KAPPERS, A. M. L., POLLICK, F. E., & KAWATO, M. (1997). Correspondence in pictorial space. *Perception & Psychophysics*, **59**, 813-827.
- KOENDERINK, J. J., & VAN DOORN, A. J. (1982a). Perception of solid shape and spatial lay-out through photometric invariants. In R. Trappl (Ed.), *Cybernetics and systems research* (pp. 943-948). Amsterdam: North-Holland.
- KOENDERINK, J. J., & VAN DOORN, A. J. (1982b). The shape of smooth objects and the way contours end. *Perception*, **11**, 129-137.
- KOENDERINK, J. J., VAN DOORN, A. J., CHRISTOU, C., & LAPPIN, J. S. (1996). Shape constancy in pictorial relief. *Perception*, **25**, 155-164.
- KOENDERINK, J. J., VAN DOORN, A. J., KAPPERS, A. M. L., & TODD, J. T. (1997). The visual contour in depth. *Perception & Psychophysics*, **59**, 828-838.
- LYNCH, K. (1960). *The image of the city*. Cambridge, MA: MIT Press.
- MACKWORTH, A. K. (1973). Interpreting pictures of polyhedral scenes. *Artificial Intelligence*, **4**, 121-137.
- MALIK, J. (1987). Interpreting line drawings of curved objects. *International Journal of Computer Vision*, **1**, 73-103.
- MANDELBROT, B. B. (1983). *The fractal geometry of nature* (2nd ed.). New York: Freeman.
- NASAR, J. L. (1998). *The evaluative image of the city*. Newbury Park, CA: Sage.
- NORMAN, J. F., PHILLIPS, F., & ROSS, H. E. (2001). Information concentration along the boundary contours of naturally shaped solid objects. *Perception*, **30**, 1285-1294.
- NORMAN, J. F., TODD, J. T., PEROTTI, V. J., & TITTLE, J. M. (1996). Visual perception of three-dimensional length. *Journal of Experimental Psychology: Human Perception & Performance*, **22**, 173-186.
- PEACHEY, D. R. (1985). Solid texturing of complex surfaces. *Computer Graphics*, **19**, 279-286.
- PERLIN, K. (1985). An image synthesizer. *Computer Graphics*, **19**, 287-296.

PHILLIPS, F. (in press). Creating noisy objects. *Perception*.

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

PHILLIPS, F., TODD, J. T., KOENDERINK, J. J., & KAPPERS, A. M. L. (1997). Perceptual localization of surface position. *Journal of Experimental Psychology: Human Perception & Performance*, **23**, 1481-1492.

REICHEL, F. D., & TODD, J. T. (1990). Perceived depth inversion of smoothly curved surfaces due to image orientation. *Journal of Experimental Psychology: Human Perception & Performance*, **16**, 653-664.

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.

SIDDIQI, K., TRESNESS, K. J., & KIMIA, B. B. (1996). Parts of visual form: Psychological aspects. *Perception*, **25**, 399-424.

SINGH, M., SEYRANIAN, G. D., & HOFFMAN, D. D. (1999). Parsing silhouettes: The short-cut rule. *Perception & Psychophysics*, **61**, 636-660.

THOMPSON, D. W. (1992). *On growth and form: The complete revised edition*. New York: Dover.

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

WALTZ, D. (1975a). Generating semantic descriptions from drawings of scenes with shadows. In P. Winston (Ed.), *The psychology of computer vision* (pp. 19-91). New York: McGraw-Hill.

WALTZ, D. (1975b). Understanding line drawings of scenes with shadows. In P. H. Winston (Ed.), *The psychology of computer vision* (pp. 19-91). New York: McGraw-Hill.

WOLFRAM, S. (1991). *Mathematica: A system for doing mathematics by computer* (2nd ed.). Redwood City, CA: Addison-Wesley.

APPENDIX
Differential Geometry of Surfaces

Given a regular surface $M \in \mathbb{R}^3$, consider the set of all planes containing a given point, P , and its normal vector N_P . Each plane will intersect M in a curve whose local curvature, κ , at P is defined by

$$\kappa = \frac{1}{r}, \tag{4}$$

where r denotes the radius of the osculating circle at P . At the limits, a $\kappa = 0$ on a curve represents a “flat” region (its osculating circle has an infinite radius), and as κ approaches infinity, the curve develops a “kink,” or discontinuity (see Figure A1). The sign of the curvature defines its relative concavity or convexity.

On a given M , at all P there are a pair of orthogonal curves whose κ s are at a minimum and maximum, relative to all others. These are the *principal curvatures*, denoted κ_1 and κ_2 (sometimes denoted κ_{\min} and κ_{\max} , respectively). By examining the relationship between them, we can define several classes of special locations on a surface.

Intuitively, a location with no curvature in either direction,

$$\kappa_1, \kappa_2 = 0 \tag{5}$$

is defined as *planar point*. A location with curvature in only one direction,

$$\kappa_1 = 0 \text{ or } \kappa_2 = 0, \tag{6}$$

is a *parabolic point*. And a location with equal, nonzero curvatures,

$$\kappa_1 = \kappa_2 \neq 0, \tag{7}$$

is an *umbilic point*. The *mean curvature*, H , is defined as

$$H = \frac{\kappa_1 + \kappa_2}{2}. \tag{8}$$

Points where $H = 0$ are *minimal points*. Minimal surfaces are surfaces where $H = 0$ everywhere. The *Gaussian curvature*, K , is defined as

$$K = \kappa_1 \kappa_2. \tag{9}$$

Points where $K < 0$ are *hyperbolic points*, whereas points with $K > 0$ are *elliptic points*.

Together, H and K define the differential nature of a surface at a given location. Koenderink (1990) has suggested a reparameterization of the principal curvatures that offer the further benefits of scale-dependant and scale-independent measures.

The *shape index*, S is the scale-independent measure, defined as

$$S = -\frac{2}{\pi} \arctan \frac{\kappa_1 + \kappa_2}{\kappa_1 - \kappa_2}, \quad \kappa_1 > \kappa_2. \tag{10}$$

When $S = \pm 0.5$, the location is a cylindrical (or *monoclastic*) “trough” or “ridge.” At $S = \pm 1.0$, the location is an elliptical (or *synclastic*) “dimple” or “bump.” And when $S = 0$, the location is a hyperbolic (or *anticlastic*) “saddle.” The sign specifies the orientation of the curvatures relative to

APPENDIX (Continued)




Figure A1. The upper left figure shows two points, P_1 and P_2 , with curvature, $\kappa = 0$ at P_1 and $\kappa > 0$ at P_2 . The right-hand P figure illustrates the osculating circle at P , whose radius, r , determines the curvature P on curve C . The osculating circle is defined as the circle that is described by P_1, P, P_2 as P_1 and P_2 approach P on C , at the limit.

the normal defined at that point. The scale-dependent complement to shape index is *curvedness*, C , defined as

$$C = \sqrt{\frac{\kappa_1^2 + \kappa_2^2}{2}}. \quad (11)$$

Curvedness is a single-valued analogue of curvature; when $C = 0$, the location is a planar point, and as C approaches infinity, the location is a crease or point (depending on S). Unlike the raw curvature, it is always positive; the direction of curvature is encoded in the shape index.

For the present set of experiments, we defined additional reparameterizations of the κ s to better understand the nature of the observers' responses. *Curvature difference*, Δ_κ , is the signed difference of curvatures, defined as

$$\Delta_\kappa = \kappa_1 - \kappa_2. \quad (12)$$

Curvature contrast, $\Delta_{|\kappa|}$, is the difference of curvature magnitudes, defined as

$$\Delta_{|\kappa|} = |\kappa_2| - |\kappa_1|. \quad (13)$$

This measure is concerned only with the magnitude of the difference, and not with the sign (and will always be positive, since $\kappa_2 \geq \kappa_1$). *Total curvature magnitude*, $\Sigma_{|\kappa|}$, is the complement to curvature contrast, defined as

$$\Sigma_{|\kappa|} = |\kappa_1| + |\kappa_2|. \quad (14)$$

As with curvature contrast, this measure uses the magnitudes and is always positive. It represents the total amount of curvature at a given location, similar to C .^{A1}

Finally, we defined two measures in the spirit of the *minima rule* of Hoffman and Richards (1984). The first is the *signed, maximum magnitude*, $\kappa_{|\max|}$, defined as

$$\kappa_{|\max|} = \begin{cases} \kappa_1 & \text{if } |\kappa_1| > |\kappa_2| \\ \kappa_2 & \text{otherwise.} \end{cases} \quad (15)$$

The second is the *maximum magnitude*, $|\kappa|_{\max}$, defined by

$$|\kappa|_{\max} = \max(|\kappa_1|, |\kappa_2|). \quad (16)$$

NOTE

A1. The signed analogue to total curvature magnitude is equivalent to the mean curvature measurement and, therefore, is not included.