Perception of surface contours and surface shape: from computation to psychophysics

David C. Knill

Department of Psychology, University of Minnesota, Minneapolis, Minnesota 55455

Received August 13, 1991; revised manuscript received February 19, 1992; accepted February 21, 1992

Contours projected from surface markings provide information for the perception of surface shape. The nature of this information depends on how the shapes of surface marking are constrained relative to the shapes of the surfaces upon which they lie. A natural constraint is that of figural regularity relative to the shape of an underlying surface. Such a constraint would be expressed in terms of the geodesic curvature of a marking, with markings having zero geodesic curvature (geodesics of a surface) being the prototypic regular figures. I propose a number of forms for a geodesic constraint and present psychophysical evidence from a contour-labeling experiment that the human visual system implicitly incorporates a geodesic constraint in the processing of reflectance contours.

INTRODUCTION

Image contours are particularly effective sources of information for the perceptual interpretation of surface shape and scene structure. One need simply look at the proliferation of line drawings used to represent three-dimensional (3-D) scenes to be convinced of this. This observation has resulted in the devotion of an entire body of work in computer vision to analyzing the information provided by different types of contour. Research has focused primarily on self-occlusion contours,[1-5] edges of polyhedral or piecewise smooth objects,[6-10] and shadow contours.[6,11] Contours that have received relatively little attention are those formed by the projections of extended markings on a surface, what Stevens[12] referred to as surface contours. Such contours are the subject of this investigation.

Surface contours are contours that project from intrinsic markings on surfaces, where intrinsic markings are taken to arise from physical processes acting on a surface independent of the viewing and the lighting geometry. Examples of intrinsic surface markings are sharp changes in surface reflectance and surface cracks.[13] Examples of nonsurface markings are self-occlusion boundaries (dependent on the viewing direction) and shadow boundaries (dependent on the lighting geometry). For the purposes of this paper, a further distinction is made between surface markings and texture markings upon a surface. Surface markings are extended markings upon a surface, while texture markings are characterized by having small spatial extents relative to the scale of smooth undulations of a surface. Much research has been done on the visual interpretation of shape from texture,[14-17] and this problem is not considered here.

Any reasonably strong inference of surface shape from surface contours must derive from prior assumptions about how the shapes of surface markings are constrained relative to the shapes of the surfaces on which they lie. Following the lead of other computational studies in vision, one could analyze the physics of how markings are formed on surfaces to derive what these constraints are. As Stevens[12] has pointed out, however, the wide range of different physical processes that underlie the formation of surface markings effectively precludes such a study. On the other hand, we can use the insights gained from phenomenal observations of the human perceptual interpretation of surface contours along with some geometric intuition to generate hypotheses about what assumptions the human visual system incorporates in its processing of surface contours (which may not, of course, correspond to the constraints that exist in the real world). Having generated such hypotheses, one can perform psychophysical experiments designed to test their validity as models of human perceptual processing. This is the strategy followed in this paper.

Section 1 presents the motivation for and a theoretical analysis of a constraint that I propose is used by the visual system in the interpretation of surface contours. A particular goal in this section is to formulate a constraint that unifies two different approaches that have been taken to inferring shape from surface contours: one focusing on the interpretation of surface shape from contours projected from smoothly curved surfaces[12,18,19] and the other focusing on the interpretation of surface orientation from contours projected from planar surfaces.[20-22] In Section 2 are presented the results of psychophysical experiments designed to test directly the validity of the proposed constraint.

1. PART I: COMPUTATIONAL ANALYSIS

The primary body of phenomenal evidence relating to the visual interpretation of surface contours comes from the appearance of various types of line drawing, as in Fig. 1. Stevens[12] has suggested that the visual system interprets the contours in such drawings as surface markings whose curvature is entirely, or mostly, attributable to the curvature of the underlying surfaces; that is, when possible, the contours are interpreted as geodesics of a surface. In the drawings shown here, with the exception of Fig. 1(f), this appears to be at least qualitatively accurate.

0740-3232/92/091449-16$05.00

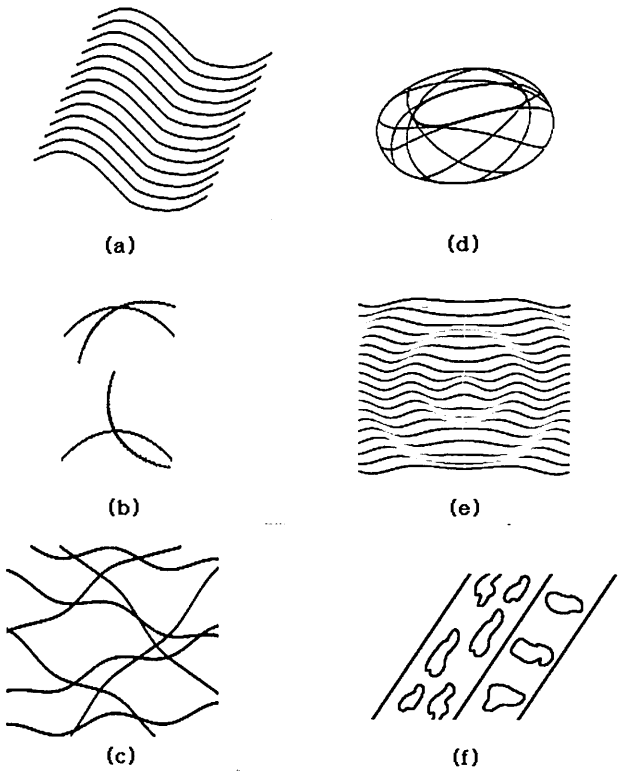

Fig. 1. Examples of line drawings that suggest 3-D surface shapes. (a) Sets of parallel contours are interpreted as lines of curvature upon a cylindrical surface. (b) Pairs of contours are often used to depict a doubly curved surface. The top pair suggests an elliptic surface, while the bottom pair suggests a hyperbolic surface patch. The validity of such interpretations depends on an assumption that the contours project from curves that are both geodesic and lines of curvature on a surface. (c), (d) Nonparallel, geodesic contours also strongly suggest the curvature of surfaces. (e) Contours need not be geodesic to suggest surface shape. Contours formed by planar sections of a surface are also a common device used to depict surface shape. Note that the shading of the contours in the image has been adjusted to remove any shading cue caused by the relative compression of contours. (f) The skew of planar figures can induce a percept of oriented surfaces in three dimensions.

In his research on analyzing the geometric constraints necessary to infer surface-shape properties from surface contours, Stevens drew particular motivation from figures such as Figs. 1(a) and 1(b), in which we seem to interpret the contours not only as geodesic but also as lines of curvature on their underlying surfaces. Using this constraint, he developed a detailed model for the accurate estimation of surface shape from sets of parallel contours [Fig. 1(a)]. The parallelness of a set of contours implies, under a general position argument, that the underlying surface is cylindrical. The assumption that the contours project from lines of maximal curvature on the surface then affords a solution for the surface shape up to one degree of freedom (corresponding to the global slant of the surface away from the line of sight). A line-of-curvature assumption also seems necessary for us accurately to infer, as we seem to, the local Gaussian curvature of a surface underlying intersecting pairs of contours, as in Fig. 1(b).

In general, the two constraints proposed by Stevens, that contours are interpreted as geodesics and that they are interpreted as lines of curvature, are not consistent, since lines of curvature are not generally geodesic. Where the two coincide is along lines of local surface symmetry (e.g., lines of curvature on cylindrical or spherical surfaces); thus our perceptual interpretation of figures such as Figs. 1(a) and 1(b) appears to be consistent with both constraints. On the other hand, Figs. 1(c) and 1(d), which were generated by drawing geodesics that are not lines of curvature of the underlying surfaces, also evoke strong percepts of surface shape. This suggests that perhaps a geodesic constraint is the more reasonable of the two possibilities as a general model of which constraints are incorporated into the human perceptual processing of surface contours. That the percept of surface shape evoked by the contours in Figs. 1(a) and 1(b) also matches a line-of-curvature constraint suggests that, when possible, the visual system imposes an additional constraint on top of a geodesic constraint. The additional constraint could be said to be either an assumption that surface markings are lines of curvature or an assumption that they are planar, since planar geodesics are, in fact, lines of curvature.[23] I prefer the latter description, because it provides for more economy of description within a general theory of surface-contour interpretation. The remainder of this section is devoted to the analysis and the development of a generalized geodesic constraint.

First I analyze how a strict assumption of geodesicity constrains the interpretation of surface shape from contours. I then use Figs. 1(f) and 1(g) to motivate a softer, generalized formulation of the constraint.

A. Strong Geodesic Constraint

One way to characterize the strategy used to interpret line drawings such as those in Figs. 1(a)–1(d) is to note that the visual system attributes all the curvature of the contours to the curvature of the underlying surface. Geodesics may be defined as those curves upon a surface whose curvature is entirely attributable to the curvature of the underlying surface; that is, those curves that have no curvature added to that imposed by the surface upon which they lie. Examples are straight lines, which are geodesics of planar surfaces, and the great circles of a

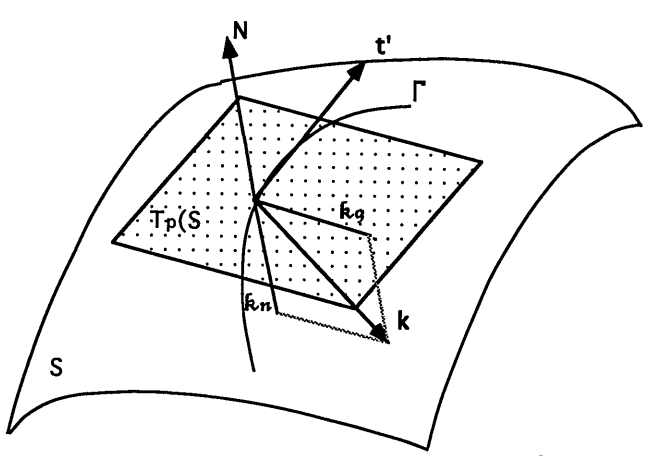

Fig. 2. Schematic of relationships between κ, κ_g, and κ_n. κ is the curvature of a curve Γ at a point p upon the surface S. The curvature vector $\mathbf{K}$ of the curve may be decomposed into a component in the tangent plane of the surface whose length is given by the geodesic curvature κ_g and a component parallel to the surface normal $\mathbf{N}$ whose length is the normal curvature κ_n of the surface at the point p traversed in the tangent direction t of the curve.

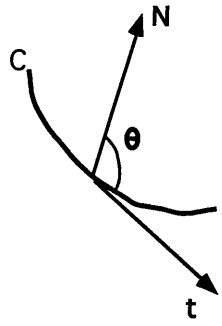

Fig. 3. Geometry for the derivation of the curvature-constraint equation. C is a contour projected from a curve upon a surface. $\mathbf{N}$ is the projection of the surface normal onto the image plane, and t is the tangent of the C. The angle between $\mathbf{N}$ and t is θ.

sphere (circles centered on the center point of the sphere), which are geodesics of the sphere.

These ideas can be formalized as follows[24]: Let $\Gamma : \mathbf{x}(s) = [x(s), y(s), z(s)]^T$ be an arbitrary curve upon a surface S. The curvature vector of Γ at a point p may be decomposed into two orthogonal components, one in the direction of the surface normal $\mathbf{N}$ and one in the tangent plane of the surface $T_p(S)$ (Fig. 2). The magnitude of the component in the direction of $\mathbf{N}$ is the magnitude of the normal curvature, $|\kappa_n|$, of the surface in the tangent direction of Γ. Note that $|\kappa_n|$ is an intrinsic property of the surface.[25] The magnitude of the component in $T_p(S)$ is denoted by $|\kappa_g|$ and is referred to as the geodesic curvature of the curve. The total curvature of Γ at p is related to κ_n and κ_g by the relation

$$\kappa^2 = \kappa_n^2 + \kappa_g^2. \tag{1}$$

Since the magnitude of curvature of any curve passing through a point in a given direction is constrained to be no less than $|\kappa_n|$, we say that κ_n reflects the curvature imposed on a curve by the surface. The excess curvature is the geodesic curvature, given by κ_g. A geodesic is defined by the condition that $\kappa_g = 0$ everywhere along its length; thus, at each point along a geodesic, it is the straightest possible curve passing through the point in its tangent direction. It also follows that the normal to the curve at a point is parallel to the surface normal at the point. This characteristic of geodesics, along with the fact that a geodesic's curvature is equivalent to the normal curvature of an underlying surface, leads us to characterize geodesics as having locally the same shapes as their underlying surfaces.

The formulation of the constraint imposed on surface-shape interpretation by contours projected from geodesics follows immediately from the equivalence between the local shape of a geodesic curve and the local shape of the surface upon which it lies. It is the same as the constraint imposed by a contour on the 3-D shape of the curve from which it projects. Assuming orthographic projection, the constraint can be expressed as an equation relating the contour curvature κ to surface orientation and normal curvature:

$$\kappa = -\kappa_n \frac{\sin\theta \sin\sigma(\cos^2\sigma + \sin^2\sigma\cos^2\theta)}{\cos^2\sigma}, \tag{2}$$

where the normal curvature κ_n is computed in the tangent direction of the curve. θ is the angle between the projection of the surface normal in the image plane and the tangent direction of the contour (the tilt of the surface relative to the tangent direction of the contour), and σ is the slant of the surface out of the image plane. The geometry is summarized in Fig. 3.

Although Eq. (2) reduces by only one the number of degrees of freedom needed to describe the local shape of a surface at a point,[26] it leads to a number of qualitative constraints on the relationship between the contour shape and the surface shape. The first of these links the sign of curvature of a contour to the sign of curvature of the underlying surface:

Constraint 1
The surface underlying a contour projected from a geodesic curves in the same direction as the contour.

Figure 4 illustrates the constraint. If the projected surface normal points inward toward the curve, then the surface's normal curvature in the tangent direction of the contour is negative. If it points outward from the curve, the surface's normal curvature in the tangent direction of the contour is positive.

A second constraint applies to the position of inflection points in a contour and can be derived by setting Eq. (2) equal to zero:

Constraint 2
If p_c is a point on a contour C at which $\kappa = 0$ (generically, an inflection point of C) and p is the corresponding point upon a surface from which p projects, then, under the assumption of a general viewpoint, the surface shape at p matches one of the following two criteria:

1. The tangent direction of C, (x', y'), is an asymptotic direction of the surface ($\kappa_n = 0$).
2. The tangent of C is parallel to the direction of surface tilt; that is, the projection of the surface normal onto the image plane ($\sin\theta = 0$).

Local contour information in a static image does not uniquely determine which of the two criteria given above leads to the occurrence of an inflection point. However, the first condition is distinguished from the second by the fact that only at points matching that condition is the correspondence between p and p_c stable to small perturbations of the projective mapping, such as would be caused by rotational motion of the observer or of the surface.

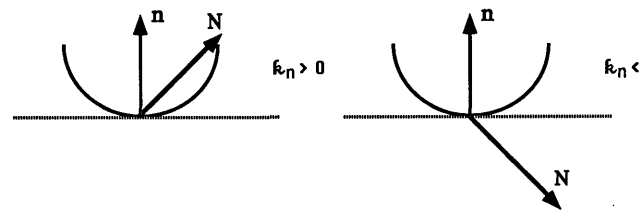

Fig. 4. Qualitative constraint on surface curvature and orientation imposed by a contour known to project from a geodesic. $\mathbf{n}$ is the contour's normal vector, $\mathbf{N}$ is the projection of the surface normal into the image, and κ_n is the surface's normal curvature computed in the tangent direction of the contour.

Rotations would cause inflection points in a contour resulting from the second condition to slide along the surface, whereas inflection points resulting from the first condition would move rigidly with the object. Inflection points in contours projected from geodesics of developable surfaces are particularly informative about surface shape, since inflection points that result from the first condition correspond to planar points of such surfaces (except in the degenerate case when the contour projects from a surface ruling and is therefore necessarily straight).

B. Softening the Geodesic Constraint

Figures 1(e) and 1(f) show two examples in which surface contours are not interpreted as geodesic but do provide information used by the visual system in the interpretation of surface shape. Figure 1(e) is an example of contours created as multiple planar sections of a circularly symmetric sine function. Figure 1(f) shows an example in which the skew of a curved (nongeodesic) figure upon a planar surface is used as information for the interpretation of surface orientation. In this subsection I discuss how the geodesic constraint could be generalized to be consistent with such phenomena.

Todd and Reichel[27] studied line drawings such as the one shown in Fig. 1(e) and showed that contours formed by such planar sections of surfaces can provide strong cues for the interpretation of surface shape. How can we reconcile this with the hypothesis of a geodesic constraint? One way to do so would be to note that the qualitative application of both constraints would lead to the same predictions of perceived surface shape. In particular, constraint 1, derived above for geodesics (that the relative sign of contour curvature and surface tilt determines the sign of surface curvature), holds for planar sections of surfaces as well, with the following caveat: For planar sections, the sign of the relationship between contour curvature and surface curvature may be reversed, depending on the orientation of the plane used to slice the surface. Along each contour, however, the sign remains fixed; thus inflection points in the contours are constrained to occur when a contour runs in an asymptotic direction of a surface. The qualitative shape predictions of the two constraints are therefore effectively equivalent.

While a qualitative application of a geodesic constraint generalizes well for handling contours projected from planar sections of surfaces, it could not explain such phenomena as the type of planar orientation from contour illustrated in Fig. 1(f). Computational accounts[20–22] of the phenomenon amount to the application of some form of figural regularity constraint in the interpretation of surface orientation. These models estimate surface orientation so as to extremize some measure of regularity (e.g., symmetry or curvature variation) applied to the deprojection of a contour onto a surface. They would, for example, interpret ellipses as circles on an oriented plane. Psychophysical studies have confirmed the psychological relevance of such ideas for the perception of surface orientation from texture, showing, as they have, that the local shapes of texture contours can be stronger cues than texture density and size gradients in determining perceived surface orientation.[28–31]

The combination of insights gained from the useful qualities of a geodesic constraint and a figural regularity constraint in the interpretation of surface contours leads to the formulation of a generalized version of both constraints in which figural regularity is measured in terms of the geodesic curvature of a surface. Examples that form the basis of our psychophysical investigation are two measures of the variation in a figure's geodesic curvature: the variance of geodesic curvature, given by

$$\mathrm{Var}(\kappa_g) = \frac{\int_C \kappa_g(s)^2 ds}{L} - \left[\frac{\int_C \kappa_g(s) ds}{L}\right]^2, \tag{3}$$

which is related to a measure of planar figural regularity proposed by Weiss,[22] and the integral of the squared derivative of geodesic curvature, given by

$$\mathrm{Int}(\kappa_g) = \int_C \left[\frac{\partial \kappa_g(s)}{\partial s}\right]^2 ds, \tag{4}$$

the latter of which is a generalized version of a figural regularity measure proposed by Barrow and Tenenbaum.[20] I refer to the two measures as the variance and the integral measures of geodesic regularity. Weiss provides a good discussion of the relative computational merits of the two types of measure in the context of planar figures, and the interested reader is urged to read his account.

Consistent with the observations that contours projected from geodesics of a surface are particularly effective in evoking shape percepts is the fact that both regularity measures are minimized for geodesics (as well as curves of any constant geodesic curvature). Furthermore, when a surface is known to be planar, the measures reduce to those proposed for planar figural regularity. The use of a geodesic regularity measure has the further advantage over a strict application of a geodesic constraint that it may be directly implemented in a working system that must integrate information from many different cues to surface shape.[31,32]

2. PART II: PSYCHOPHYSICS

Does the human visual system incorporate a geodesic constraint in the perceptual processing of surface contours? As we have seen, contours projected from geodesics are effective in eliciting accurate percepts of surface shape; however, this is not conclusive evidence for the hypothesis. It seems clear that the visual system does not assume the strong form of geodesic constraint: that all surface markings follow geodesic paths. Does it, however, incorporate a soft geodesic constraint similar to those incorporated in the geodesic regularity measures I have proposed? The phenomenal observations just alluded to do not resolve this question. In this part of the paper I present psychophysical experiments designed to test the hypothesis that the visual system assumes some form of geodesic regularity constraint on surface contours. In particular, we consider the case of reflectance contours formed by the projection of discontinuous changes in the reflectance of a surface.

An important aspect of any quantitative study of surface-contour processing is that naturalistic images be used as stimuli, so that the contours on which subjects

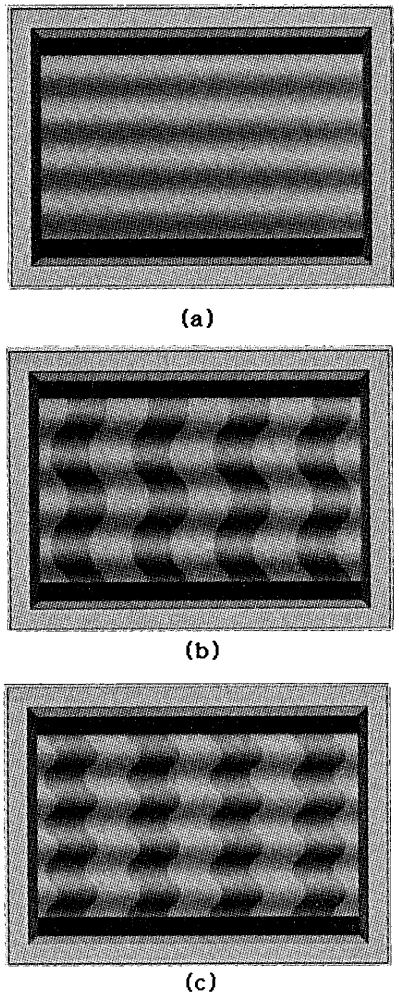

Fig. 5. Differently shaped stripes painted over the same shading pattern evoke different shape percepts. In this case shading pattern (a) only ambiguously determines the shape of the underlying surface. It could be two bumps illuminated from behind the camera, as in (b), or it could be four bumps illuminated from above, as in (c). The shapes of the reflectance contours in (b) and (c) determine the perceptual interpretation.

base their perceptual judgments are formed by the same types of luminance singularity as found in real images. What little research has been done on the problem of surface-contour interpretation has relied on the use of line drawings, as exemplified by the phenomenal demonstrations in Section 1. The extent to which such results generalize to contours in natural images is unclear. Figure 5 shows that some of the effects obtained for line drawings of surface contours do generalize to a more naturalistic image, in which the contours are formed by sharp changes in surface reflectance. We might expect, however, that studies that rely on line drawings would underestimate the role played by surface contours in shape perception, since any configuration of contours in a line drawing must be a strong enough cue for shape to jump start the shape-from-contour system in the first place.

The use of naturalistic images has its drawbacks. In particular, such images contain a multitude of cues for surface shape besides surface contours, so that any effective study of the effects of surface contours on shape perception would have to be done in the context of cooperative interactions between the contour information and other shape cues (e.g., shading). In the experiments presented here, I avoid this problem by using a contour-labeling task to probe the constraints imposed by the visual system in the interpretation of reflectance contours. In the task, subjects are asked to report the physical cause of a set of contours in an image. Since such a task involves a small number of discrete choices about the global labeling of a contour, we can easily isolate the effect of changes in the geodesic regularity of a contour on subjects' performance.

The specific task was one in which the subjects were asked to label closed contours in rendered images of a surface as either reflectance contours projected from the surface or the occluding contours of a transparency placed between the viewer and the surface. I refer to the first type of labeling as a reflectance contour labeling and to the second as a transparency contour labeling. Figure 6 shows examples of the type of stimulus used in the experiments. In both images the interpretation of the middle darkened patch is ambiguous: it could result from a change in reflectance on the vaselike surface or from a second transparent surface floating over the surface. The geodesic regularity hypothesis makes a clear prediction about how human beings will label such patches: all other things being equal, greater geodesic regularity of the bounding figure of a patch relative to the background surface will produce greater proportions of reflectance contour judgments.

Two experiments were done to investigate the effect of contour shape on the labeling of contours as reflectance contours or as transparency contours. The first experiment may be considered preliminary in the sense that it was designed simply to determine whether any such effect could be obtained as well as to illuminate the types of strategy used by subjects to perform the task. The second experiment was designed to isolate the effect of contours' geodesic regularity on labeling judgments. Since the method of stimulus generation and the experimental procedures were the same for both experiments, I describe them before going into the specifics of each experiment.

A. General Methods

1. *Stimuli*

Each stimulus image used in the experiment was created by a nonlinear combination of a pair of base images. The first image of a pair was a shaded image of a surface created by computer rendering of a 3-D object model. I

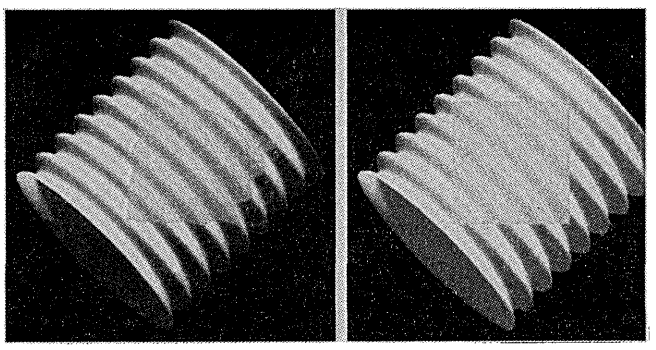

Fig. 6. Examples of the stimuli used in the experiments.

refer to this image as the surface image for a stimulus. The second image was a binary image representing the closed region in the stimulus image in which a patch was to be painted. This image took on a value of 1 inside the patch region and 0 outside it. This image is referred to as the patch image for a stimulus.

Stimuli were constructed from pairs of surface and patch images by using the following formula for pixel intensities:

$$I(x,y) = \tau_p I_p(x,y) I_s(x,y) + 255\rho I_p(x,y) + [1 - I_p(x,y)] I_s(x,y), \quad (5)$$

where $I_p(x,y)$ is the binary value of pixel (x,y) in the patch image and $I_s(x,y)$ is the intensity of pixel (x,y) in the surface image. The parameters τ and ρ may be interpreted, respectively, as the transmittance and the reflectance of a transparent filter placed over the surface[33] or as parameters specifying a change in the reflectance properties of the surface. To maintain consistency with their interpretation as parameters of a transparent surface, τ and ρ as well as their sum $\tau + \rho$ must be less than 1. The values used for these parameters in the experiments conformed to this constraint.

Surface images were generated by 3-D rendering of polygonized models of smooth surfaces. The surfaces were modeled as having matte reflectance and were illuminated by both a point source at infinity and an ambient light. The resulting images were downsampled, with averaging, by a factor of 3, to provide some antialiasing of the surface borders. The specific details of the surface and the lighting models used for each experiment are given in the individual methods subsections. The methods of constructing the patch images differed for the two experiments; therefore, they too are described in the individual methods subsections. The background luminance of the display was fixed at a midgray of 33 cd/m^2.

2. *Apparatus*

Surface and patch images were generated on a Stellar GS1000 computer. The surface images were rendered by using Stellar's Automatic Visualization System. An Apple Macintosh IIX workstation was used to run the experiments. Stimulus images were created by combining surface and patch images dynamically during an experimental run and were presented on the workstation's 8-bit color display monitor. The nonlinearity of the screen ($\gamma = 2.05$) was corrected through modification of the lookup table. The correction lowered the effective resolution of the screen to 7.5 bits. The spatial resolution of the screen for the viewing distance used in the experiments (0.5 m) was 0.0477°/pixel. The subjects viewed the screen monocularly through a reduction screen that hid the borders of the display and with their chins resting on a chin rest. They made responses with the workstation's mouse.

3. *Procedure*

Subjects were presented with a stimulus image on each trial and asked to judge whether the dark patch in the middle of the surface in the image was a reflectance change on the surface or a transparent surface floating in front of the surface. Since preliminary studies showed that subjects might see a patch as a shadow on the surface, subjects were asked to record such an interpretation as a transparency. Subjects were given an unlimited viewing time and were asked to make their judgments by selecting one of two icons below the stimulus image, using the computer's mouse. Figure 7 shows an example display. Subjects in each experiment were split into two equal-sized groups, one of which was presented with the icons in one order (left—transparency, right—reflectance) and the other of which was presented with the icons in the reverse order (left—reflectance, right—transparency).

Each experiment consisted of multiple presentations of a set of stimulus images. The trials were organized into blocks, each of which contained a randomly ordered presentation of all the images in the stimulus set. To control for the effects of stimulus presentation order, a different random number seed was used to control the order of presentation for each subject in an experiment. Before the main experiment began, subjects were given one block of practice trials to accustom them to the task. Subjects' scores consisted of the number of transparency and reflectance judgments made for each stimulus image. The trials were run in the dark. The time taken to explain the experimental task and to run through the practice trials was typically more than 5 min, providing ample time for the subjects to adapt to the background luminance of the display.

B. Experiment 1

The first experiment was designed to determine whether contour shape plays a role in subjects' labeling of contours as well as to gain insight into how subjects perform the task. The contour-labeling task could be subject to three

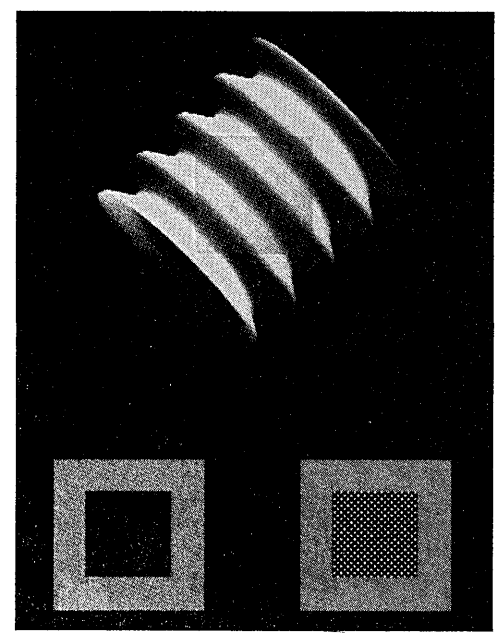

Fig. 7. Example of the display presented to the subjects in the experiment. The solid icon was used to indicate a reflectance judgment, and the icon filled with the cross-hatched pattern was used to indicate a transparency judgment. After recording a judgment on a trial, the subjects indicated that they were ready for the next trial by selecting, with the mouse, the black square below the two judgment icons.

types of prior perceptual constraint: one applied to reflectance edges, one applied to the occluding edges of thin-surfaces, and one that is simply the prior bias to see patches like those in the stimulus images as reflectance changes or transparent surfaces (presumably fixed for all the stimuli). Subject to the prior bias, the relative degrees to which a set of contours matches the constraints on reflectance edges and thin occluding edges determine the frequency of reflectance and transparency labelings.

The subjects were presented with two classes of stimulus designed to optimize the effect of contour shape on labeling. The first class contained patches whose bounding contours perfectly matched the proposed geodesic constraint; that is, which projected from geodesics of the surfaces in the stimulus images. The second class contained patches whose bounding contours were straight; that is, they were optimized for a putative figural regularity constraint on thin occluding edges. The patches were painted over surface images by using a multiplicative combination rule [$\rho = 0$ in Eq. (5)], so that the luminance relations across contours were consistent with interpretations of patches either as neutral-density filters floating over surfaces or as constant changes in surface reflectance. For these stimuli, the shapes of patch bounding contours provide the only information for determining the contour labeling.

1. Methods

a. Stimuli. Figure 8 shows the stimulus images used for the experiment. Two different surface models were used in the generation of the surface images: surfaces of revolution, whose generator curves were sinusoids of different frequencies, and Monge surfaces,[34] whose depth functions were the sums of sines of different frequencies. I refer to these as vase and crate surfaces, respectively, according to their qualitative appearance. The details of the surface and the shading models used to generate the surface images are given in Appendix A.

Four patch images, two of which had straight bounding contours and two of which had contours projected from geodesics of the surface, were generated for each surface image. One of the straight patches was a square oriented out of the image plane by the same angle as the surfaces, while the other was an irregular polygon. A matching geodesic patch was generated for each of the two straight patches by tracing geodesics on a surface between the points over which the corners of the straight patch appeared (see Appendix B for details of how the geodesics were computed). The geodesics were then projected into the image plane to create the bounding contours of the patch.

The combination rule given in Eq. (5), with $\tau = 0.3$ and $\rho = 0$, was used to combine the surface images with each of their corresponding patch images, resulting in 24 final stimulus images. The dimensions of the projected surfaces presented to the subjects were approximately 133 × 133 pixels along the sides, though this varied by small amounts between surfaces.

b. Procedure. The procedure given in Subsection 2.A was used to collect subjects' scores for frequency of transparency judgments for each stimulus image. Eight judgments were recorded for each of the twenty-four stimulus images.

c. Subjects. Twelve undergraduates from an introductory psychology class served as subjects for the experiment. Seven subjects were emmetropes, four were myopes with 20/20 corrected vision, and one was a hyperope with 20/20 corrected vision.

2. Results

Figure 9 shows the percentage of transparency judgments for each stimulus, averaged across all subjects. An effect of contour shape is clearly evident in the results for each of the six surfaces used, with straight patches consistently judged transparent more often than geodesic patches. Wilcoxon signed-ranks tests for each of the six surfaces showed the effect of straightness versus geodesicity to be highly significant for all the surfaces ($W = 0$, $p < 0.005$, for all surfaces). The difference in percentage of transparency judgments between vase (37.2%) and crate (55.5%) surfaces was also highly significant ($W = 3$, $p < 0.005$).

3. Discussion

The most notable feature of the results is that the subjects consistently judged straight patches to be transparent more often than geodesic patches. This suggests an implicit assumption of prior constraints on thin occluding edges and/or prior constraints on reflectance edges. Unfortunately, the results do not answer the and/or question, namely, which of the two types of edge does the visual system assume constraints on, or does it assume constraints on both? Either of the three possibilities could explain the data. For example, an assumption that straight occluding edges are more common than any given curved edge would, by itself, lead to an increase in transparency judgments for straight patches. An implementation of such an assumption is the commonly used straightness invariant: an interpretation rule stating that straight edges in an image project from straight edges in the scene.[35] By the same token, an assumption that geodesic reflectance edges are more common than any given nongeodesic curved edge would lead to an increase in reflectance judgments (or, equivalently, a decrease in transparency judgments) for the geodesic patches, since the straight patches have bounding contours that would not have projected from geodesics of the underlying surfaces.

To isolate and test the hypothesized geodesic regularity constraint on reflectance edges, we need to vary the geodesic regularity of patch bounding contours while holding the figural regularity in the image plane constant. This is done in experiment 2; however, before moving on to a description of that experiment, let us note a number of informal observations obtained through interviews with the subjects after the performance of the experiment. These will motivate a further change in the design of the experiment.

The most significant observation made by the subjects was that they had a strong bias to see all the patches as painted onto the underlying surfaces. Often the subjects reported that they initially based their judgments on how easily they could switch their perceptual interpretation of a patch to be transparent. They reported that, as the experiment wore on, these changes in percepts became more automatic; however, the initial observation raises the

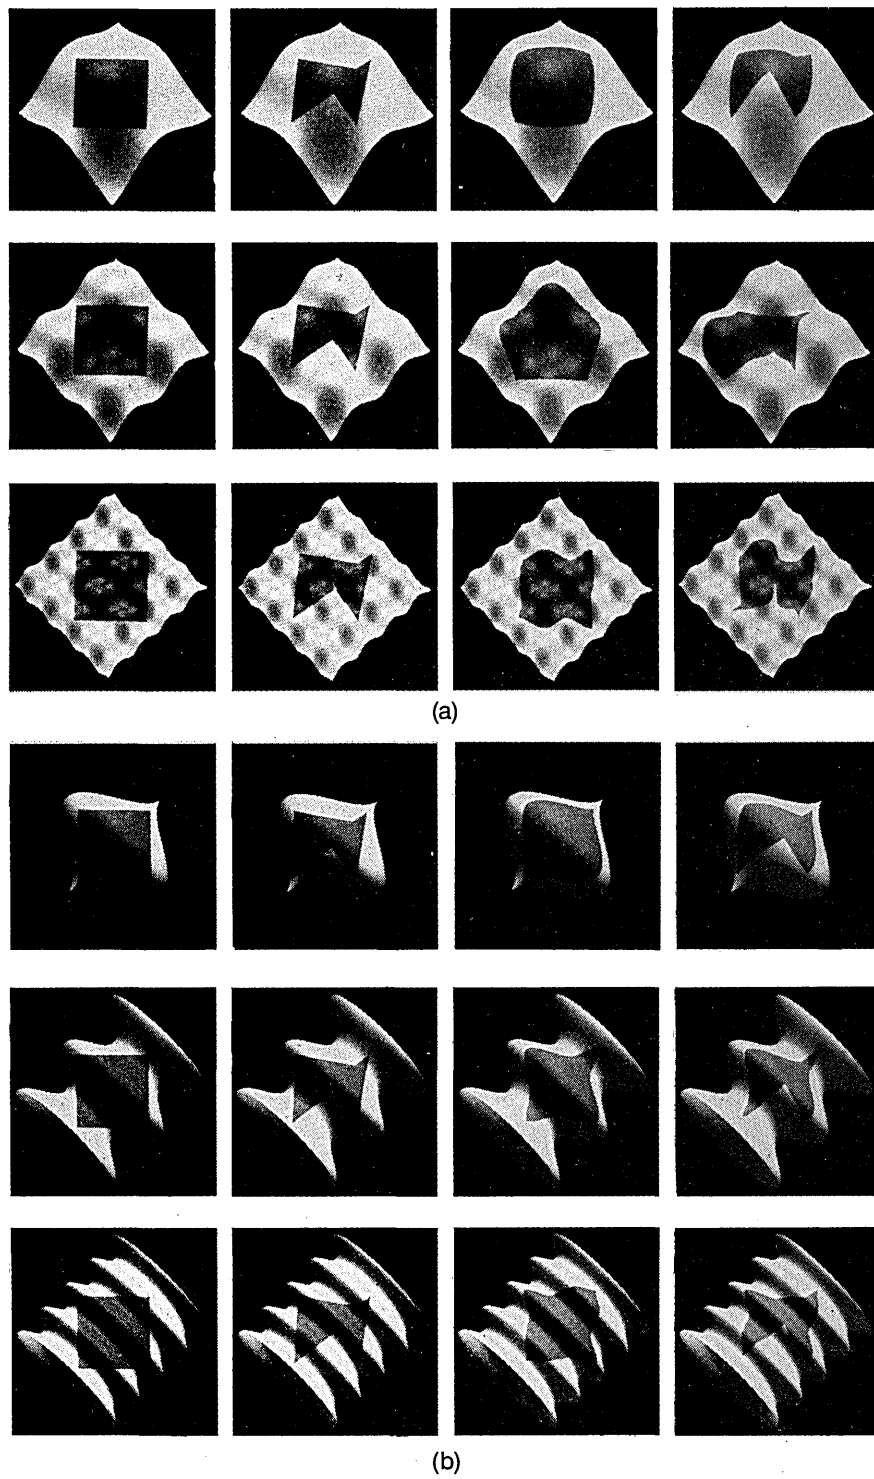

Fig. 8. Stimuli used in experiment 1. The first two columns in both (a) and (b) contain the straight patches, and the second two columns contain the geodesic patches.

problem that cognitive factors may have played a strong role in subjects' performance. This is further implicated by the fact that most of the subjects spontaneously reported their detection of the main experimental manipulation, expressing thoughts such as, "You were looking to see if I saw patches with straight boundaries as transparent more often than patches with curved boundaries, weren't you?" The design of the second experiment speaks to these issues.

C. Experiment 2

In experiment 2, I directly test the predictions made by the hypothesis that the visual system assumes a geodesic regularity constraint on reflectance edges. Within the labeling paradigm being used here, this requires varying the geodesic regularity of patch bounding contours while holding the shape of the contours in the image plane fixed. Two transformations that serve the desired function are rotations and translations of a patch in the image plane

relative to a fixed surface. While these transformations leave the shapes of patch bounding contours unchanged in the image, they lead to variations in contour shape relative to the underlying surface, since they change the positions of the contours relative to the surface. I chose to use rotations to vary patch orientation relative to a fixed surface. Measuring the geodesic regularity of the patch bounding contours in these stimuli as a function of patch orientation and relating this to subjects' transparency judgments provides a direct test of the prediction that decreases in geodesic regularity should lead to corresponding increases in transparency judgments.

A further change in the design was made to account for the problems of cognitive penetration raised in the discussion of the first experiment. In this experiment I varied a new factor in stimulus generation: the value of the additive constant ρ used in combining surface and patch images [see again Eq. (5)]. Remember that this was held constant at 0 in experiment 1. Varying the value of ρ serves two functions: First, the addition of a second factor serves to mask the main experimental manipulation (the different rotations of patches), reducing the possible influence of cognitive factors. The effectiveness of variations in ρ in this regard was nicely borne out in the informal reports given by the subjects after participating in the experiment (see the discussion of results). Second, increases in ρ shift subjects' labeling bias toward a transparency judgment, an effect confirmed by the results of a supplementary experiment described in Fig. 10. An intuitive explanation of the effect is that increases in ρ make the patch appear flatter and more opaque by reducing the luminance contrast within the patch. While the added constant luminance within a patch is consistent with an interpretation of the patch as a flat transparent surface with constant reflectance (when $\rho \gg \tau$, the patch appears opaque), it has no such clear interpretation as a change in surface reflectance properties.

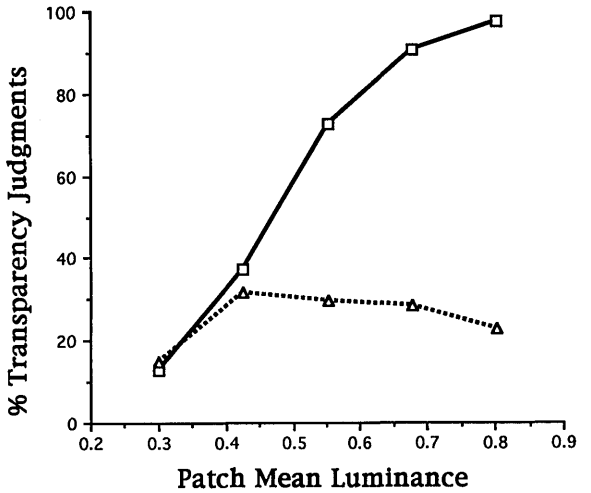

Fig. 10. Mean luminance of a patch manipulated either by varying the multiplicative constant τ (△) used in stimulus generation while keeping ρ constant at 0 or by varying ρ (□) while keeping τ constant at 0.3. The graph shows the percentage of transparency judgments as a function of patch mean luminance for the two conditions. When the mean luminance was increased by varying ρ, the percentage of transparency judgments increased dramatically [Friedman two-way analysis of variance: $\chi^2(4) = 37.92, p < 0.0001$]. The percentage of transparency judgments did not, however, increase significantly with increases in τ [Friedman $\chi^2(4) = 3.6, p > 0.2$].

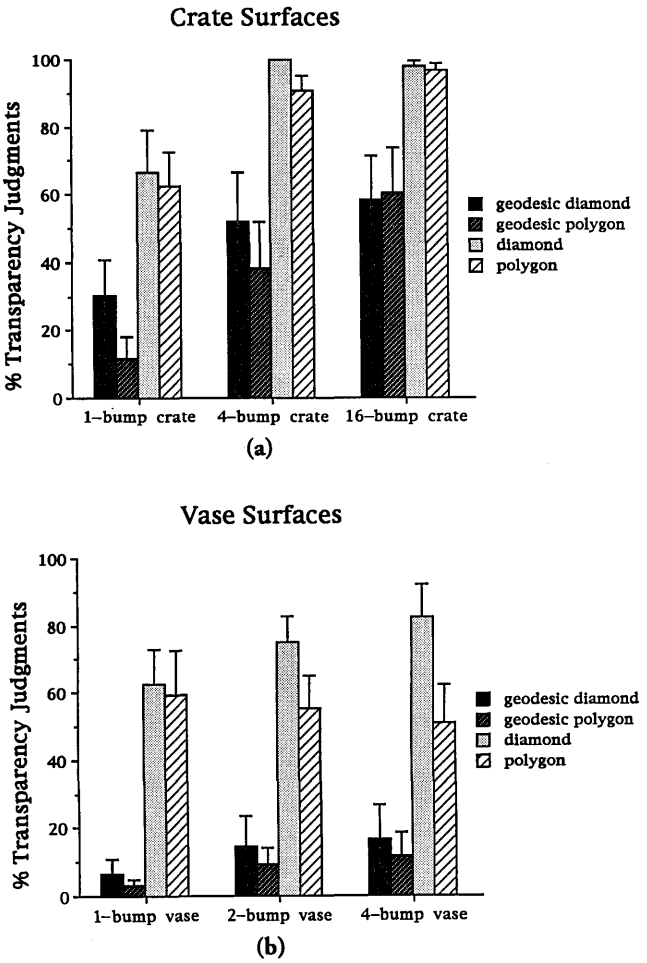

Fig. 9. Results of experiment 1 plotted as the percentage of transparency judgments for crate and vase surfaces.

1. *Methods*

a. Stimuli. The image of an eight-bump vase surface was used for all the stimuli in the experiment (see Appendix A for an exact specification of the surface model). Three different patch images were created from the vertices of a four-sided irregular polygon[36] in the following way: The vertices were deprojected onto three different surfaces and connected by tracing geodesics of the surfaces between the vertices. Each of these curves was then projected back into the image plane, and the resulting closed regions were filled in to create a patch image. One of the surfaces used for creating the patches was the eight-bump surface used in generating the stimuli, while the other two were four-bump and two-bump vase surfaces. Each of the three patches was rotated by 0°, 22.5°, and 45° in the image plane to generate three sets of stimuli whose patches all had equivalent bounding contour shape in the image plane. This gave a total of nine patches.

The nine patches were combined with the surface image by using the multiplicative–additive combination rule given in Eq. (5). Three different additive coefficients were used, resulting in a total of 27 different stimulus images. The multiplicative coefficient used was $\tau = 0.3$.

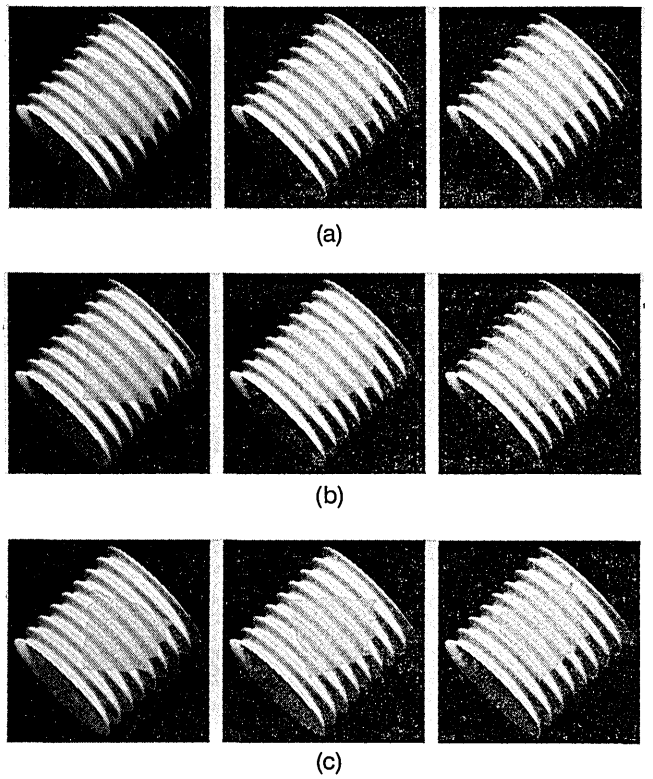

Fig. 11. Examples of stimulus images used in experiment 2 ($\rho = 0.0$): (a) images created with patch 1, (b) images created with patch 2, (c) images created with patch 3.

The additive coefficients used were 0.0, 0.15, and 0.3. Figure 11 shows the nine stimulus images generated with $\rho = 0$.

b. Procedure. The procedure given in Subsection 2.A was used to collect the subjects' scores for frequency of transparency judgments for each stimulus image. Eight judgments were recorded for each of the 24 stimulus images.

c. Subjects. Twelve undergraduates from an introductory psychology class served as the subjects for the experiment. Eight subjects were emmetropes, and four were myopes with 20/20 corrected vision.

2. Results

I computed both the integral [Eq. (4)] and the variance [Eq. (3)] measures of the geodesic regularity of patch bounding contours for comparison with subjects' labeling judgments. These are listed in Table 1 for each of the stimuli used in the experiment (Appendix C details the method of computing the two measures). As expected, the values of both measures applied to the patch boundaries of each patch varied as a function of the patch's rotation angle in the image. Note that only for patch 1 do the two measures make different predictions about labeling performance.

Figure 12 shows the percentage of transparency judgments for each patch plotted as functions of both the variance and the integral measures. The differences in percentage of transparency judgments, when averaged across values of ρ, were highly significant for all three patches [patch 1: Friedman two-way analysis of variance, $\chi^2(2) = 12.41, p < 0.002$; patch 2: $\chi^2(2) = 14.93$, $p < 0.001$; patch 3: $\chi^2(2) = 14.93, p < 0.001$]. For each of the three patches, the percentage of transparency judgments increased with increases in the variance of geodesic curvature. The same is true for patches 2 and 3 when the comparison is made with the integral of the derivative of geodesic curvature but does not hold for patch 1.

Informal reports of the subjects given after they took part in the experiment indicate that the addition of the factor ρ to the experimental design was successful in masking the main experimental manipulation of rotating the patches. When asked what hypothesis was being tested in the experiment, all the subjects guessed that I was looking for an effect of patch lightness on performance.

3. Discussion

The results support the hypothesis that the visual system assumes a geodesic regularity constraint on reflectance edge shape. The variance of geodesic curvature provided a better predictor of the ordinal ranking of the subjects' transparency labelings than did the integral measure; however, the quantity of data collected in the experiment is not large enough to permit us to test the relative performances of the two measures as models of the constraint incorporated in visual system processing.

A number of caveats should be made concerning the interpretation of the results. The first regards an implicit assumption made in the analysis, namely, that the shape of the surface in the stimulus images was perceived correctly and did not vary with variations in patch geometry. Presumably, any constraint on reflectance edge shape applied by the visual system would be applied relative to the perceived shape of the surface; thus exact predictions of performance would require an accurate estimate of the perceived surface shape and the geodesic regularity of patch bounding contours relative to that shape. The use of a highly symmetric and regular surface was an attempt to minimize any differences in actual and perceived shape, and I suspect that, at the least, the qualitative aspects of the shape (e.g., that it was a surface of revolution, where the bumps and valleys were, etc.) were accurately perceived.

A second caveat has to do with an aspect of contour shape that may have played a role in labeling but that was not considered here: the angles formed by the vertices in patch boundaries. A number of researchers have sug-

Table 1. Variance of Geodesic Curvature and Integral of Squared Derivative of Geodesic Curvature Measures Applied to Each of the Nine Patches Used in Experiment 2

Patch Number	Rotation (°)	Var(κ_g)	Int(κ_g)
1	0	11.42	1.34
	22.5	11.6	0.89
	45	13.18	1.91
2	0	12.72	1.15
	22.5	13.0	1.75
	45	14.63	2.2
3	0	0.02	0.08
	22.5	16.24	1.16
	45	21.5	2.38

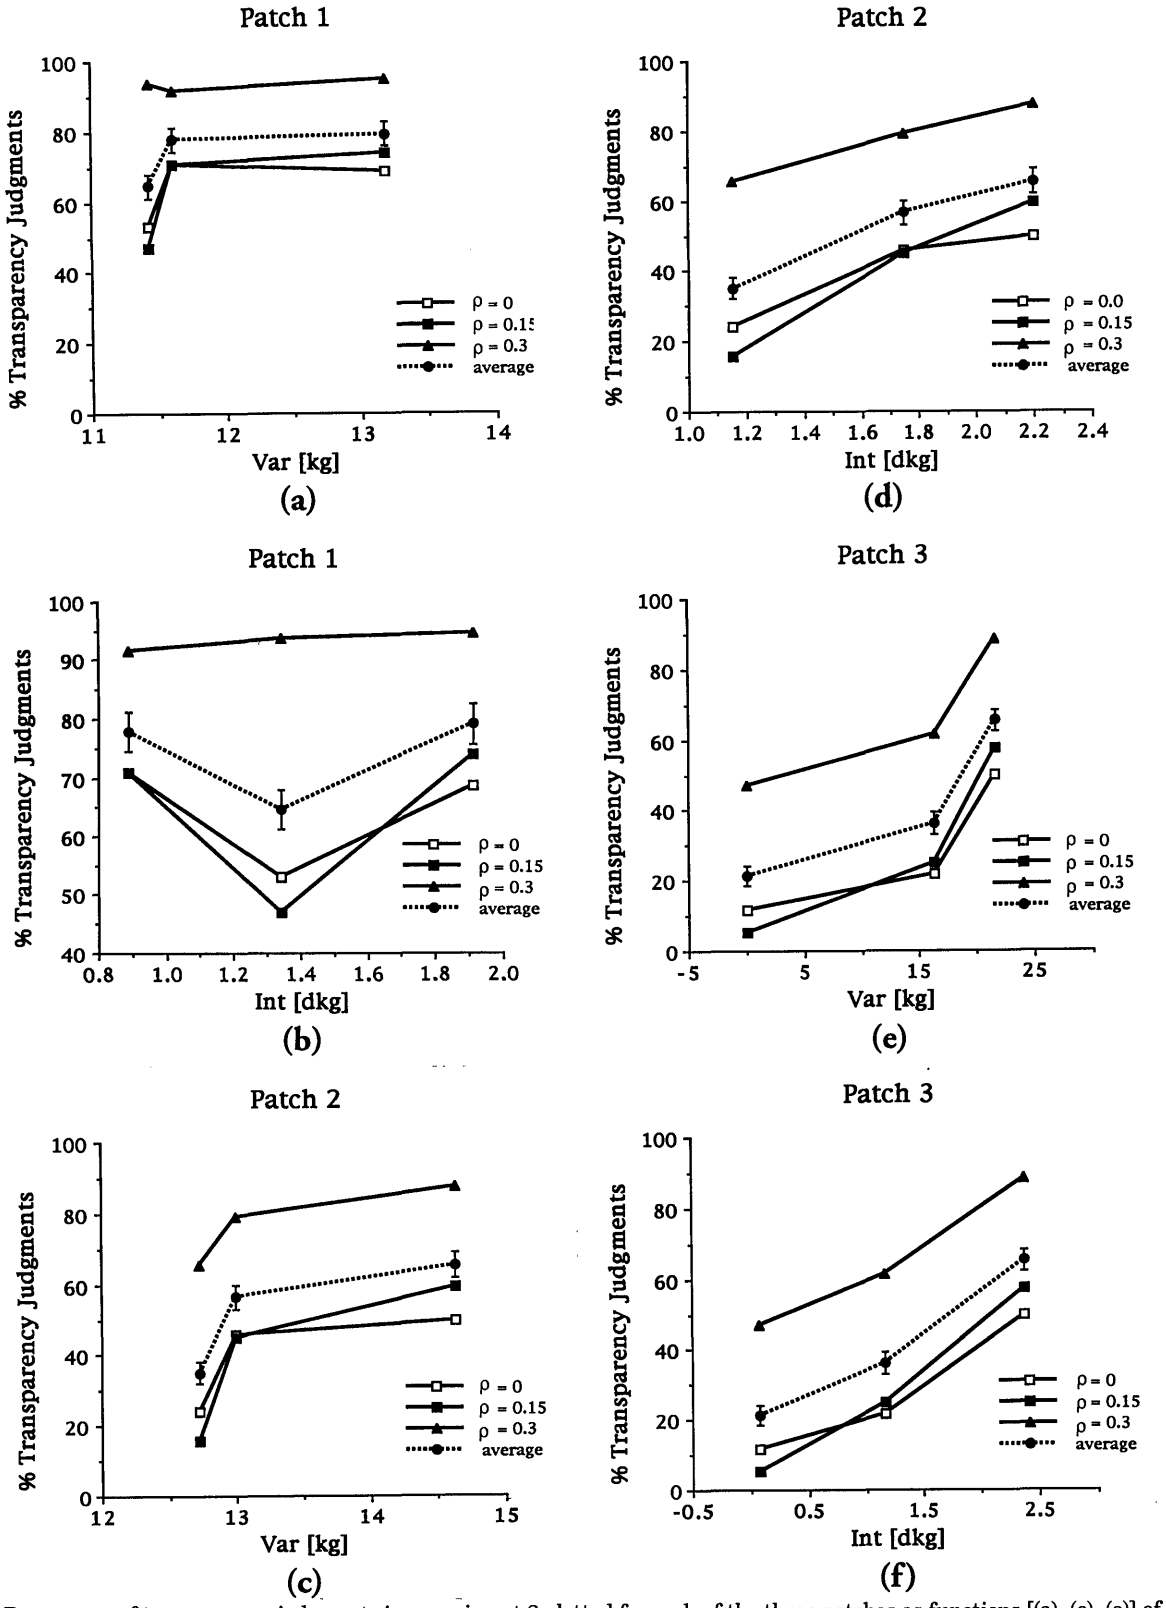

Fig. 12. Percentage of transparency judgments in experiment 2 plotted for each of the three patches as functions [(a), (c), (e)] of the variance of geodesic curvature of a patch's boundary and [(b), (d), (f)] of the integral of the squared derivative of geodesic curvature of a patch's boundary.

gested that the visual system assumes that corners of surface markings tend toward being perpendicular.[12] Such a constraint may have played some part in determining the subjects' labeling; however, I think that the contribution of the constraint in this experiment was small.

D. General Discussion

I have presented psychophysical evidence that the human visual system assumes some form of geodesic regularity constraint in the interpretation of reflectance contours. Previous studies relating to the problem, specifically those

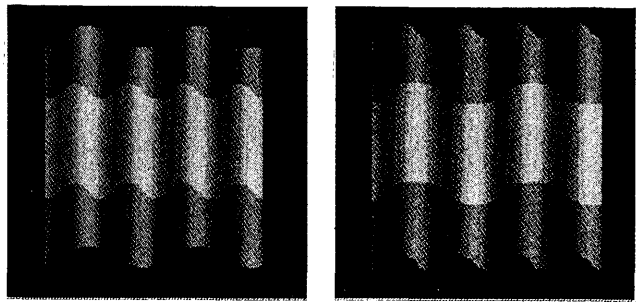

Fig. 13. Variants of the images shown in Fig. 5 in which both occluding contours and reflectance contours are present.

of Stevens[12] and Todd and Reichel,[27] relied on phenomenal demonstrations of perceived shape from line-drawing images in their development. In this, they differ from the more directed set of experiments using naturalistic images presented here. As pointed out in Part I of this paper, however, the assumption of a soft geodesic constraint is consistent with the type of phenomenon investigated by these researchers.

A number of considerations are important in the generalization of the present results. The first is that the stimuli used in the experiments were images of particularly simple types of surface, of the type that characterize our artifactual environment. This raises the possibility that the perceptual processes tapped into in the experiment were specialized for the interpretation of such surfaces. The possibility takes on more weight when one considers that the physical processes that underlie the creation of surface markings on artificial objects are ones that are likely to adhere to some sort of regularity constraint.[12] To test whether surface-contour constraints are general or more special purpose in the sense described here, further studies with irregular objects are required.

A second consideration is that some form of regularity constraint other than one based on the geodesic curvature of figures might predict the results. Though this could certainly be true, the geodesic constraint seems to be the most parsimonious and general predictor of a range of phenomena, including those cited in Part I of this paper. In this regard further tests of the geodesic constraint are needed. In particular, experiments need to be done to investigate the contribution of surface contours to shape perception per se. The contour-labeling paradigm used here does not provide a direct look at what would be the most significant functional role of the constraint: to support the interpretation of surface shape. Though line-drawing studies do focus on this aspect of surface-contour processing, I believe that more-naturalistic stimuli, as used here, should supplant the use of line drawings in quantitative studies of the perception of shape from contour.

1. *Interaction with Other Cues*

The weakness of the geodesic constraint as the sole determinant of surface shape suggests the importance of investigating the interactions between surface contours and other sources of information such as shading and texture in shape perception. An example of a seemingly comparable source of information is that provided by surface discontinuities, such as the occluding contours of a thin surface. Such contours are well-known strong cues for surface shape[20,37,38]; yet, just as with reflectance contours, strong assumptions about the shapes of the edges from which they project are needed for the inference of surface shape. Figure 13 shows images of two surfaces in which the shapes of the reflectance and the occluding contours provide conflicting evidence for surface shape. To this author, the shapes of the surfaces appear to follow the lowest frequency suggested by either of the contours, though they are both somewhat bistable. Clearly the interaction between the two types of contour is not a simple one and requires a more detailed investigation.

2. *Surface Contours and Surface Creases*

In this paper I have focused on surface contours that project from smoothly curved regions of a surface. Let us briefly discuss here the problem of how the visual system interprets surface contours at points where they intersect surface-crease contours (contours projected from discontinuities in surface orientation, such as the corners of a box). Figure 14 shows that the visual system does use the

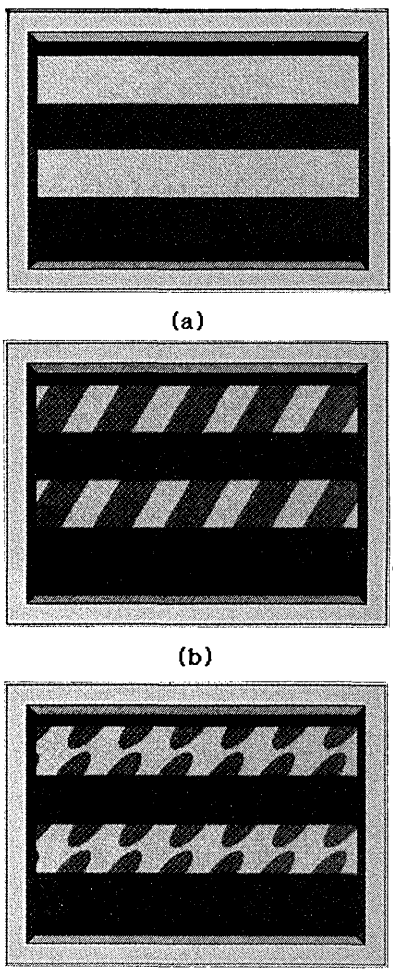

Fig. 14. Parallel stripes (a) take on the appearance of a folded surface when reflectance contours have tangent discontinuities at the luminance edges (b) and (c). The perceived degree of a surface's fold depends on the angle formed by the discontinuity in the reflectance contours. The effect does not depend on the contours being straight from the surface creases (c).

information provided by the behavior of surface contours at intersections with surface-crease contours in its interpretation of surface shape. Unlike for surface contours on smooth surface regions, however, an analysis of the behavior of surface markings at intersections with surface creases in the real world does lead to the formulation of an ecologically valid constraint that can be applied to contour interpretation.

Consider the case of creases formed by folding a developable surface (a surface that can be created by bending, without stretching or compressing, a flat surface). Any markings that intersect such a crease will necessarily have a discontinuity in their tangent directions. Moreover, the angle through which a marking's tangent is taken by this discontinuity is determined by the way in which the surface was folded at the crease. Let us specify the applicable constraint in the following way: Define a pseudosurface normal vector $\mathbf{N}_p$ at the crease as the bisector of the surface normals at either side of the crease:

$$\mathbf{N}_p = \frac{\mathbf{N}^+ + \mathbf{N}^-}{|\mathbf{N}^+ + \mathbf{N}^-|}, \tag{6}$$

where $\mathbf{N}^+$ and $\mathbf{N}^-$ are the surface normals immediately to either side of the crease at a point. Similarly, for a discontinuity in a curve's tangent, we can define the pseudonormal vector $\mathbf{n}_p$ of the curve as the bisector of the tangent vectors of the curve at either side of the discontinuity:

$$\mathbf{n}_p = \frac{\hat{t}^+ + \hat{t}^-}{|\hat{t}^+ + \hat{t}^-|}, \tag{7}$$

where t^+ and t^- are the unit tangent vectors of the curve as the discontinuity is approached from either side. Where a surface marking intersects a surface crease formed by folding a developable surface, the pseudonormal of the marking is constrained to be parallel to the pseudonormal of the surface crease; that is, $\mathbf{n}_p = \pm\mathbf{N}_p$ (note the similarity to a geodesic constraint for smooth surfaces). The constraint does not require that a surface marking be straight in the neighborhood of a crease, just that it be smooth.

The phenomenal appearance of the surfaces in Fig. 14 suggests that the visual system does assume something similar to the pseudonormal constraint in interpreting surface contours at intersections with surface-crease contours.

E. Summary

Surface contours provide information that the visual system uses in the interpretation of surface shape. I have analyzed the information provided by such contours under the assumption that the contours project from geodesics of surfaces and have proposed a generalized, soft version of a geodesic constraint that could apply to the interpretation of nongeodesic contours. Psychophysical results based on contour labeling that are consistent with the hypothesized constraint have been presented. Taken together with phenomenal observations on the effectiveness of geodesic contours as cues for surface shape, these results lead to the conclusion that the visual system applies some form of geodesic constraint in the interpretation of surface contours. I also considered the information provided by surface contours at intersections with surface-crease contours and showed that, for a special class of surface creases formed by folding developable surfaces, the behavior of surface markings was lawfully constrained in a manner similar to that of geodesics on smooth surface patches.

APPENDIX A: SURFACE MODELS

This appendix summarizes the surface and the shading models used in generating the surface images for experiments 1 and 2.

The crate surfaces used in experiment 1 were modeled as graphs of the function $A[\cos(2\pi f_u u + \phi_u)\cos(2\pi f_v + \phi_v)]$ parameterized as

$$\mathbf{X}(u,v) = \begin{bmatrix} u \\ v \\ A[\cos(2\pi f_u u + \phi_u)\cos(2\pi f_v v + \phi_v)] \end{bmatrix} \quad (0 \le u < x_{\max}, 0 \le v < y_{\max}), \tag{A1}$$

where $x_{\max}$ and $y_{\max}$ specify the width and the height of the surface and $2A$ is the maximum depth difference between peaks and troughs in the surface, the frequency parameters f_u and f_v determine the bumpiness of the surface, and the phase parameters ϕ_u and ϕ_v determine the positions of the bumps. Table 2 summarizes the parameters used to generate the three different crate surfaces used in experiment 1.

The vase surfaces used in both experiment 1 and experiment 2 were modeled as surfaces of revolution with a sinusoidal generating curve, parameterized as

$$\mathbf{X}(u,v) = \begin{bmatrix} x_0 - [R + A\cos(2\pi f v + \phi)\cos u] \\ v \\ R + A\cos(2\pi f v + \phi)\sin u \end{bmatrix} \quad (0 \le u < 2\pi, 0 \le v < y_{\max}), \tag{A2}$$

where x_0 is the point at which the axis of revolution intersects the x axis, the width of the surface at its widest point is given by $2(R + A)$, and the height is given by $y_{\max}$. The frequency parameter f determines the bumpiness of the surface, and the phase parameter ϕ determines the positions of the bumps. Table 3 summarizes the parameters used in generating the vase surfaces for experiments 1 and 2.

Polygonized models of the surfaces were rendered in the Automatic Visualization System on a Stellar GS1000. Each crate surface was rotated by $-10°$ about the x axis (the z axis was negative in the direction of the model camera), $-15°$ about the y axis, and $-45°$ about the z axis, in that order, and each vase surface was rotated by $-10°$ about the x axis and $-45°$ about the z axis. Surfaces were shaded by using a Phong shading model, the parameters of which are given in Table 4.

APPENDIX B: COMPUTING GEODESICS

The geodesics of a surface S may be defined in the following way: Let $S \subset \Re^3$ be represented by a parameterization, $\mathbf{X}: U \subset \Re^2 \rightarrow S$, where a point on S is given by $\mathbf{X}(u,v)$. A geodesic of S is given by a curve $\mathbf{X}[u(s),v(s)]$ that

Table 2. Surface Parameters for Crate Surfaces

x_{max}	y_{max}	A	f_u	f_v	ϕ_u	ϕ_v
400	400	100	0.0025 (1 cycle/surface)	0.0025 (1 cycle/surface)	π	π
400	400	500	0.005 (2 cycles/surface)	0.005 (2 cycles/surface)	π	π
400	400	25	0.01 (4 cycles/surface)	0.01 (4 cycles/surface)	0	π

satisfies the pair of second-order nonlinear differential equations

$$u'' + \Gamma_{11}^{1}(u')^2 + 2\Gamma_{12}^{1}u'v' + \Gamma_{22}^{1}(v')^2 = 0, \quad \text{(B1)}$$

$$u'' + \Gamma_{11}^{2}(u')^2 + 2\Gamma_{12}^{2}u'v' + \Gamma_{22}^{2}(v')^2 = 0, \quad \text{(B2)}$$

where the Γ_{ij}^{k} are the Christoffel symbols for the surface defined at each point (u,v) by the system of equations

$$\Gamma_{11}^{1}\langle \mathbf{X}_u, \mathbf{X}_u\rangle + \Gamma_{11}^{2}\langle \mathbf{X}_u, \mathbf{X}_v\rangle = \langle \mathbf{X}_{uu}, \mathbf{X}_u\rangle, \quad \text{(B3)}$$

$$\Gamma_{11}^{1}\langle \mathbf{X}_u, \mathbf{X}_v\rangle + \Gamma_{11}^{2}\langle \mathbf{X}_v, \mathbf{X}_v\rangle = \langle \mathbf{X}_{uu}, \mathbf{X}_v\rangle, \quad \text{(B4)}$$

$$\Gamma_{12}^{1}\langle \mathbf{X}_u, \mathbf{X}_u\rangle + \Gamma_{12}^{2}\langle \mathbf{X}_u, \mathbf{X}_v\rangle = \langle \mathbf{X}_{uv}, \mathbf{X}_u\rangle, \quad \text{(B5)}$$

$$\Gamma_{12}^{1}\langle \mathbf{X}_u, \mathbf{X}_v\rangle + \Gamma_{12}^{2}\langle \mathbf{X}_v, \mathbf{X}_v\rangle = \langle \mathbf{X}_{uv}, \mathbf{X}_v\rangle, \quad \text{(B6)}$$

$$\Gamma_{22}^{1}\langle \mathbf{X}_u, \mathbf{X}_u\rangle + \Gamma_{22}^{2}\langle \mathbf{X}_u, \mathbf{X}_v\rangle = \langle \mathbf{X}_{vv}, \mathbf{X}_u\rangle, \quad \text{(B7)}$$

$$\Gamma_{22}^{1}\langle \mathbf{X}_u, \mathbf{X}_v\rangle + \Gamma_{22}^{2}\langle \mathbf{X}_v, \mathbf{X}_v\rangle = \langle \mathbf{X}_{vv}, \mathbf{X}_v\rangle. \quad \text{(B8)}$$

For a given set of initial conditions (u_0,v_0) and (u_0',v_0'), the differential equations have a unique solution. The initial conditions define a point on the surface, $\mathbf{X}(u_0,v_0)$, and an initial direction in the tangent plane of the surface at the point, $\mathbf{X}' = u_0'\mathbf{X}_u(u_0,v_0) + v_0'\mathbf{X}_v(u_0,v_0)$.

For a given initial point and initial direction, the numerical integration of the differential equations specified in Eq. (B2) provides an estimate of the points along a geodesic on the surface. The computation of the geodesic connecting two points on a surface requires searching over the space of possible initial directions from one of the points for the direction whose geodesic will contain the other point.

The following algorithm was used to generate the geodesic patches used in the experiments:

Select a pair of connected vertices in the image, (x_0, y_0) and (x_f, y_f).

Deproject the vertices onto the surface, using orthographic projection, to obtain $\mathbf{X}_0$ and $\mathbf{X}_f$.

Determine $(\mathbf{u}_0,\mathbf{v}_0)$ by using the inverse map $\mathbf{X}^{-1} : S \subset \Re^3 \rightarrow U$.

Search through (u_0',v_0')

  Numerically integrate Eq. (B2), using (u_0,v_0) and (u_0',v_0') as the initial conditions, to estimate points $\mathbf{X}_i(u_i,v_i)$ along the geodesic.

Until $\min_i(|X_i - X_f|) < 0.5$ pixel unit.

A modified Euler method was used for the numerical integration of Eq. (B2), in which the discrete step size for the derivatives was dynamically varied to keep the separation between the points of the computed geodesic on the surface below a fixed resolution limit. For the geodesics used in the experiments, the resolution was very high, enforcing the limits $|\Delta\mathbf{X}_i| < 0.1$ pixel unit and $|\Delta\mathbf{X}_i'| < 0.1$ pixel unit.

APPENDIX C: COMPUTING MEASURES OF GEODESIC REGULARITY

For each patch image used in the experiments, a fine-resolution (0.1-pixel-unit) representation of the edges of the patches was stored for use in calculating the two measures of geodesic regularity proposed in the text: the variance of geodesic curvature of the edges [Eq. (4)] and the integral of the squared derivative of geodesic curvature [Eq. (3)]. The geodesic curvature of an edge at each point was estimated by using discrete-difference approximations of the tangent and the curvature vectors of the edge. The curvature vector was approximated as

$$\mathbf{k}_i \approx \frac{\Delta t_i}{\Delta s_i} \quad \text{(C1)}$$

$$\approx \frac{t_{i+1} - t_{i-1}}{|\mathbf{X}_{i+1} - \mathbf{X}_{i-1}|}, \quad \text{(C2)}$$

where t_i is the tangent vector at a point, approximated by

$$t_i \approx \frac{\Delta\mathbf{X}_i}{\Delta s_i} \quad \text{(C3)}$$

$$\approx \frac{\mathbf{X}_{i+1} - \mathbf{X}_{i-1}}{|\mathbf{X}_{i+1} - \mathbf{X}_{i-1}|}. \quad \text{(C4)}$$

$\mathbf{X}_i$ is the position of point i on the surface. The geodesic curvature at a point was calculated by using the expression

$$\kappa_{gi} = \langle t_i \Lambda \mathbf{N}_i, \mathbf{k}_i\rangle, \quad \text{(C5)}$$

where $\mathbf{N}_i$ is the unit surface normal at the point, computed directly from the parameterized surface model.

Both the variance of the geodesic curvature and the integral of the squared derivative of geodesic curvature of a

Table 3. Surface Parameters for Vase Surfaces Used in Experiments 1 and 2

Expt.	y_{max}	R	A	x_0	f	ϕ
1	400	100	100	200	0.0025 (1 cycle/surface)	π
	400	150	50	200	0.005 (2 cycles/surface)	0
	400	175	25	200	0.01 (4 cycles/surface)	0
2	400	175	25	200	0.02 (8 cycles/surface)	0

Table 4. Shading Parameters Used for Rendering the Two Types of Surface[a]

Surface Type	ρ	λ_p	λ_a	$\mathbf{L}$
Crate	1.0	204	51	$(0, 0.5735, -0.8192)^T$
Vase	1.0	204	51	$(0, 0.5, -0.866)^T$

[a] ρ is the surface reflectance, λ_p is the intensity of a point source at infinity given in gray-level units, λ_a is the intensity of the ambient light, and $\mathbf{L}$ is a unit vector in the direction of the light source. Images of the crate surfaces were generated by using a light source at a 35° slant away from the viewer. Images of the vase surfaces were generated by using a light source at a 30° slant away from the viewer.

patch edge were computed by using discrete sums to approximate the integrals involved. The estimates are given by

$$\mathrm{Var}(\kappa_g) = \frac{\int \kappa_g(s)^2 ds}{\int ds} - \left[\frac{\int \kappa_g(s) ds}{\int ds}\right]^2 \tag{C6}$$

$$\approx \frac{\sum_i \kappa_{gi}{}^2 \Delta s_i}{\sum_i \Delta s_i} - \left[\frac{\sum_i \kappa_{gi} \Delta s_i}{\sum_i \Delta s_i}\right]^2 \tag{C7}$$

$$\approx \frac{\sum_i \kappa_{gi}{}^2 |\mathbf{X}_{i+1} - \mathbf{X}_i|}{\sum_i |\mathbf{X}_{i+1} - \mathbf{X}_i|} - \left[\frac{\sum_i \kappa_{gi} |\mathbf{X}_{i+1} - \mathbf{X}_i|}{\sum_i |\mathbf{X}_{i+1} - \mathbf{X}_i|}\right]^2, \tag{C8}$$

$$\mathrm{Int}(\kappa_g) = \int \left[\frac{\partial \kappa_g(s)}{\partial s}\right]^2 ds \tag{C9}$$

$$\approx \sum_i \left(\frac{\Delta \kappa_{gi}}{\Delta s_i}\right)^2 \Delta s_i \tag{C10}$$

$$\approx \sum_i \frac{(\Delta \kappa_{gi})^2}{\Delta s_i} \tag{C11}$$

$$\approx \sum_i \frac{(\kappa_{gi+1} - \kappa_{gi})^2}{|\mathbf{X}_{i+1} - \mathbf{X}_i|}. \tag{C12}$$

The total measures for a patch listed in Table 1 were computed as the weighted sum of the measures for each side of the patch:

$$\mathrm{Var}_T(\kappa_g) = \sum_i^4 \lambda_i \, \mathrm{Var}_i(\kappa_g), \tag{C13}$$

$$\mathrm{Int}_T(\kappa_g) = \sum_i^4 \lambda_i \, \mathrm{Int}_i(\kappa_g), \tag{C14}$$

where λ_i is the length of side i of the patch and $\mathrm{Var}_i(\kappa_g)$ and $\mathrm{Int}_i(\kappa_g)$ are the calculated values of the two measures for side i of the patch.

Both measures were calculated by using a range of scales for the discrete-difference operations. The values varied only slightly when calculated over scales ranging from an approximately 0.1-pixel-unit separation between points to an approximately 2-pixel-unit separation between points.

ACKNOWLEDGMENTS

The author thanks Dan Kersten for many helpful discussions on the implications of a geodesic constraint on surface markings. Thanks go to Jim Anderson, Bill Warren, and Dan Kersten for their support of the dissertation from which this study grew. This research was supported by National Science Foundation grant BNS-8518675 and U.S. Air Force grant AFOSR 90-2074.

REFERENCES AND NOTES

1. J. J. Koenderink and A. J. van Doorn, "The shape of smooth objects and the way contours end," Perception **11,** 129–137 (1982).
2. J. J. Koenderink, "What does occluding contour tell us about solid shape," Perception **13,** 321–330 (1984).
3. J. J. Koenderink and A. J. van Doorn, "The singularities of the visual mapping," Biol. Cybern. **24,** 51–59 (1976).
4. V. S. Nalwa, "Line-drawing interpretation: a mathematical framework," Int. J. Comput. Vision **2,** 103–124 (1988).
5. J. M. H. Beusmans, D. D. Hoffman, and B. M. Bennett, "Description of solid shape and its inference from occluding contours," J. Opt. Soc. Am. A **4,** 1155–1167 (1987).
6. D. Waltz, "Understanding line drawings of scenes with shadows," in *The Psychology of Computer Vision,* P. H. Winston, ed. (McGraw-Hill, New York, 1975), pp. 19–91.
7. D. A. Huffman, "Impossible objects as nonsense sentences," Mach. Intell. **6,** 295–323 (1971).
8. D. A. Huffman, "Realizable configurations of lines in pictures of polyhedra," Mach. Intell. **8,** 493–509 (1971).
9. M. B. Clowes, "On seeing things," Artif. Intell. **2,** 79–116 (1971).
10. J. Malik, "Interpreting line drawings of curved objects," Int. J. Comput. Vision **1,** 73–103 (1987).
11. S. A. Shafer and T. Kanade, "Using shadows in finding surface orientations," Comput. Vision Graphics Image Process. **22,** 145–176 (1983).
12. K. A. Stevens, "The visual interpretation of surface contours," Artif. Intell. **17,** 47–73 (1981).
13. Surface cracks may be considered markings upon a smooth surface when their width is of a scale well below that of the smooth undulations in the surface.
14. J. J. Gibson, *Perception of the Visible World* (Houghton Mifflin, Boston, Mass., 1950).
15. K. A. Stevens, "The information content of texture gradients," Biol. Cybern. **42,** 95–105 (1981).
16. A. P. Witkin, "Recovering surface shape and orientation from texture," Artif. Intell. **17,** 17–47 (1981).
17. A. Blake and C. Marinos, "Shape from texture: estimation, isotropy and moments," Tech. Rep. OUEL 1774/89 (Oxford U. Press, Oxford, 1989).
18. K. A. Stevens, "Inferring shape from contours across surfaces," in *From Pixels to Predicates: Recent Advances in Computational Vision,* A. P. Pentland, ed. (Ablex, Norwood, N.J., 1986), pp. 93–110.
19. K. A. Stevens, "The line of curvature constraint and the interpretation of 3D shape from parallel surface contours," in *Natural Computation,* W. Richards, ed. (MIT Press, Cambridge, Mass., 1988), pp. 83–98.
20. H. G. Barrow and J. M. Tenenbaum, "Interpreting line drawings as three-dimensional surfaces," Artif. Intell. **17,** 75–116 (1981).
21. M. Brady and A. L. Yuille, "An extremum principle for shape from contour," IEEE Trans. Pattern Anal. Mach. Intell. **PAMI-6,** 288–301 (1984).
22. I. Weiss, "3D shape representation by contours," Comput. Vision Graphics Image Process. **41,** 80–100 (1988).
23. Stevens[12] suggests that certain configurations of contours in an image, such as parallel contours, trigger the line-of-curvature constraint. They might also be said to trigger a planarity constraint that is applied in addition to a geodesic constraint.
24. M. P. Do Carmo, *Differential Geometry of Curves and Surfaces* (Prentice-Hall, Englewood Cliffs, N.J., 1976).
25. Normal curvature is defined for each direction through a point by slicing the surface with a plane containing the normal vector of the surface and computing the curvature of the resulting curve of intersection at the point.
26. In general, one requires 5 degrees of freedom to specify fully the local shape of a surface: 2 for the orientation and 3 for the curvature. The 3 degrees of freedom for curvature correspond to the two principal curvatures and the direction along which one of them is computed (the direction for the second is fixed at 90° from the first). The normal curvature of the surface traversed in any direction can be computed from the principal curvatures by the relation $\kappa_n = \kappa_1 \cos\theta + \kappa_2 \sin\theta$, where κ_1 and κ_2 are the two principal curvatures and θ is the direction of traversal relative to the direction of the first principal curvature.
27. J. T. Todd and F. D. Reichel, "The visual perception of smoothly curved surfaces from double projected contour pat-

terns," J. Exp. Psychol. Hum. Percept. Perform. **16,** 665–674 (1990).
28. Recall that I have distinguished texture markings as occurring at a scale much finer than the curvature of an underlying surface, so that individual markings can be approximated as lying upon planar surface patches.
29. H. E. Gruber and W. C. Clark, "Perception of slanted surfaces," Percept. Mot. Skills **6,** 97–106 (1956).
30. M. L. Braunstein and J. W. Payne, "Perspective and form ratio as determinants of relative slant judgements," J. Exp. Psychol. **3,** 585–590 (1969).
31. J. J. Clark and A. L. Yuille, *Data Fusion for Sensory Information Processing Systems* (Kluwer, Boston, Mass., 1990).
32. D. Kersten, "Transparency and the cooperative computation of scene attributes," in *Computational Models of Visual Processing*, M. S. Landy and J. A. Movshon, eds. (MIT Press, Cambridge, Mass., 1991), pp. 209–228.
33. Equation (5) is actually a quadratic approximation of a filter model of transparency; J. Beck, K. Praxdny, and R. Ivry, "The perception of transparency with achromatic colors," Percept. Psychophys. **35,** 407–422 (1984).
34. A Monge surface is one that may be represented as the graph of a two-dimensional function, $f(x, y)$.
35. Contrary to popular belief, the straightness invariant does not result solely from a general viewpoint assumption but rather requires an assumption of some structure in the environment. The assumption is that straight edges form a distinguished set of configurations of edges in the environment. For a discussion of this point, see B. M. Bennett, D. D. Hoffman, and C. Prakash, *Observer Mechanics: A Formal Theory of Perception* (Academic, New York, 1989); D. C. Knill and D. Kersten, "Ideal perceptual observers for computation, psychophysics and neural networks," in *Pattern Recognition by Man and Machine,* R. Watt, ed., Vol. 14 of Vision and Visual Dysfunction (Macmillan, New York, 1991), Chap. 7, pp. 83–92.
36. An irregular polygon was used in generating the patches to avoid the orientation-specific effects that might arise from the use of symmetric patches.
37. E. Mach, *Contributions to the Analysis of the Sensations*, C. M. Williams, trans. (Open Court, Chicago, Ill., 1987; original work published in 1890).
38. V. S. Ramachandran, "Perceiving shape from shading," Sci. Am. **256,** 76–83 (1988).