

SRI International



SHADING INTO TEXTURE

Technical Note No. 398

October 2, 1986

By: Alex P. Pentland

Artificial Intelligence Center
Computer Science and Technology Division

Shading into Texture*

Alex P. Pentland

*Artificial Intelligence Center, SRI International,
333 Ravenswood Ave., Menlo Park, CA 94025, U.S.A.*

Recommended by Michael Brady

ABSTRACT

Current shape-from-shading and shape-from-texture methods are applicable only to smooth surfaces, while real surfaces are often rough and crumpled. To extend such methods to real surfaces we must have a model that also applies to rough surfaces. The fractal surface model [6] provides a formalism that is competent to describe such natural 3-D surfaces and, in addition, is able to predict human perceptual judgments of smoothness versus roughness. We have used this model of natural surface shapes to derive a technique for 3-D shape estimation that treats shaded and textured surfaces in a unified manner.

1. Introduction

The world that surrounds us, except for man-made environments, is typically formed of complex, rough, and jumbled surfaces. Current representational schemes, in contrast, employ smooth, analytical primitives—e.g., generalized cylinders or splines—to describe three-dimensional shapes. While such smooth-surfaced representations function well in man-made, carpentered environments, they break down when we attempt to describe the crenulated, crumpled surfaces typical of natural objects. This problem is most acute when we attempt to develop techniques for recovering 3-D shape, for how can we expect to extract 3-D information in a world populated by rough, crumpled surfaces when all of our models refer to smooth surfaces only? The lack of a 3-D model for such naturally occurring surfaces has generally restricted image-understanding efforts to a world populated exclusively by smooth objects, a sort of “Play-Doh” world [1] that is not much more general than the blocks world.

Standard shape-from-shading [2, 3] methods, for instance, all employ the heuristic of “smoothness” to relate neighboring points on a surface. Shape-from-

*The research reported herein was supported by National Science Foundation Grant No. DCR-83-12766 and the Defense Advanced Research Projects Agency under Contract No. MDA 903-83-C-0027 (monitored by the U.S. Army Engineer Topographic Laboratory).

texture [4, 5] methods make similar assumptions: their models are concerned with markings on a smooth surface. Before we can reliably employ such techniques in the *natural* world, we must be able to determine which surfaces are smooth and which are not—or else generalize our techniques to include the rough, crumpled surfaces typically found in nature.

To accomplish this, we must have recourse to a 3-D model competent to describe both crumpled surfaces and smooth ones. Ideally, we would like a model that captures the intuition that smooth surfaces are the limiting case of rough, textured ones, for such a model might allow us to formulate a unified framework for obtaining shape from both shading (smooth surfaces) and texture (rough surfaces, markings on smooth surfaces). We require a 3-D model competent to describe a continuum of surfaces ranging from the smooth to the rough as is illustrated in Fig. 1. The fractal model of surface shape [6, 7] appears to possess the required properties.




FIG. 1. Spherical shapes with surface crenulations ranging from smooth (fractal dimension = topological dimension) to rough (fractal dimension \gg topological dimension).

Evidence for the adequacy of this model comes from several sources. Recently conducted surveys of natural imagery [7, 8], for instance, found that the fractal model of imaged 3-D surfaces furnishes an accurate description of how most homogeneous textured or shaded image regions change over scale (change in resolution). Perhaps most convincing, however, is the fact that fractals *look* like natural surfaces [9–11]. This is important information because the natural appearance of fractals is strong evidence that they capture much of the perceptually relevant shape structure of natural surfaces.

In this paper we will use the fractal model to develop a method of estimating shape that applies to either smooth or rough surfaces, and that treats both shading and texture in a unified manner. The result will be a technique that resembles both previous shape-from-shading algorithms and previous shape-from-texture algorithms, but with some important changes that we believe may have a significant effect on the robustness of the estimation process. In particular, we suggest methods of addressing the problems of scale, and of anisotropic, “stretched” textures.

2. Background: The Fractal Model

During the last twenty years, Mandelbrot has developed and popularized a relatively novel class of mathematical functions known as *fractals* [9, 10]. The defining characteristic of a fractal is that it has a *fractional dimension*, from which we get the word “fractal.” The notion of fractal dimension is very close to our intuitive definition of roughness. In our previous work [7, 11], for instance, we have shown that the fractal dimension of a surface correlates nearly perfectly with our perception of roughness in many situations. Thus, if we were to generate a series of scenes with the same 3-D relief but with increasing fractal dimension D , we would obtain a sequence of surfaces with linearly increasing perceptual roughness (see Fig. 1).

Fractals are found extensively in nature [9, 10, 12]; Mandelbrot, for instance, shows that fractal surfaces are produced by many basic physical processes. One general characterization of naturally occurring fractals is that they are the end result of any physical processes that randomly modifies shape through local action. After innumerable repetitions, such processes will typically produce a fractal surface shape. Thus clouds, mountains, turbulent water, lightning and even music have all been shown to have a fractal form [9].

Such fractals are a generalization of Brownian motion, i.e., the path of a random walk. Such naturally occurring shapes have two important properties: (i) each segment is statistically similar to all others; (ii) segments at different scales are statistically indistinguishable, i.e., as we examine such a surface at greater or lesser imaging resolution its statistics (curvature, etc.) remain the same. Because of these invariances, the most important *variable* in the description of such a shape is how it varies with scale; in essence, how many

large features there are relative to the number of middle-sized and smaller-sized features. When we consider a natural surface such as a mountain, there will be a certain number of large features (e.g., bumps large relative to our field of view) within a particular patch of the surface, many more intermediate-sized features, and still more smaller features. For fractal shapes, and thus for many real shapes, the ratio of the number of features of one size to the number of features of the next larger size is a constant—a surprising fact that derives from the scale invariance of random walks.

If we move closer to such a surface, the intermediate-sized features now appear larger relative to our field of view, the small become intermediate-sized, and what was previously seen as surface “texture” becomes visible as small individual features. The ratio between the number of features of one size to the number of features of the next larger size, however, remains the same constant value as before. The statistics of the surface are constant across changes in scale, i.e., as we move closer to or farther away from the surface.

We have used this relationship to develop a new method of constructing a fractal shape, one that explicitly represents structures at all scales. This construction is illustrated in Fig. 2(a). We first pick a *fractal scaling parameter* r , $0 \leq r \leq 1$, a ratio that determines the fractal dimension of the surface (i.e., the fractal dimension D of a surface is determined by $D = T + r$, where T is the topological dimension of the surface) and randomly place n^2 large bumps on a plane, giving the bumps a Gaussian distribution of altitude (with variance σ^2), as seen in Fig. 2(a). We then add to that $4n^2$ bumps of half the size, and altitude variance $\sigma^2 r^2$, as shown in Fig. 2(b). We continue with $16n^2$ bumps of one quarter the size, and altitude $\sigma^2 r^4$, then $64n^2$ bumps one eighth size, and altitude $\sigma^2 r^6$ and so forth. The final result, shown in Fig. 2(c), is a true Brownian fractal shape. This construction does not depend on the particular shape of the bumps employed;¹ the only constraint is that the sum must fill out the Fourier domain. Figures 2(d) and 2(e) illustrate the power and generality of this construction; all of the forms and surfaces in these images can be constructed in this manner.

In practice, we often produce a fractal surface by filling a pyramid structure with Gaussian-distributed random noise of variance σ^2 at the top level of the pyramid, $\sigma^2 r^2$ at the next level, and so forth. Because of the expansion between successive levels of the pyramid, entries at the top of the pyramid produce bumps twice as large as entries at the next level, four times as large as the third level, and so forth. Thus when we successively expand and sum the levels of the pyramid, the end result is a surface that is a close approximation to a true fractal surface.

When the placement and size of these bumps is random, we obtain the

¹Differently shaped bumps will, however, give different appearance or texture to the resulting fractal surface; this is an important and as yet relatively uninvestigated aspect of the fractal model.




FIG. 2. (a)–(c) show the construction of a fractal shape by successive addition of smaller and smaller features with number of features and amplitudes described by the ratio $1/r$. All of the forms and surfaces shown in (d) and (e) (which are images by Voss and Mandelbrot, see [9]) can be generated in this manner.

classical Brownian fractal surface that has been the subject of our previous research. When some components of this sum are matched to a particular object, however, we obtain a description of that object that is exact for those details encompassed by the specified components. The advantage of this construction over other techniques (e.g., [9, 13]) is that it makes it possible to

exactly specify a shape at one range of scales while retaining a qualitative, statistical description at other scales.

To describe a complex natural form such as a cloud or mountain, for instance, we might specify the "bumps" down to the desired level of detail by fixing the larger elements of this sum, and then we specify only the fractal statistics of the smaller bumps in order to fix the qualitative appearance of the surface. Figure 1 illustrates an example of such description. The overall shape is that of a sphere; to this specified large-scale shape, smaller bumps were added randomly. The smaller bumps were added with six different choices of r (i.e., six different choices of fractal statistics) resulting in six qualitatively different surfaces—each with the same basic spherical shape.

The fractal scaling parameter r is the ratio between the number of features of one size to the number of features at another size, and thus describes how the surface varies across different scales (resolutions, spatial frequency channels, etc.). This ratio, therefore, summarizes how complex the surface is; how many features of one size there are for each larger feature. It is an intrinsic property of the surface;² surfaces formed by different processes typically have different fractal scaling parameters. Thus this parameter allows us to crudely classify the surface in terms of the process that formed it.³ We have found that by measuring the fractal scaling parameter r in patches of the image we can infer the fractal scaling parameter r of the imaged 3-D surface [6, 7]; in experiments this has allowed us to closely predict people's perception of surface roughness [11]; we can speculate, therefore, that the demonstrated ability of people to preattentively segment an image on the basis of this scaling parameter gives them a method of segmenting the scene into regions that were separately formed.

2.1. The mathematics of fractal Brownian functions

The path of a particle exhibiting Brownian motion is the canonical example of most naturally occurring fractals; the discussion that follows, therefore, will be devoted exclusively to fractal Brownian functions, which are a mathematical generalization of Brownian motion.

A random function $I(x)$ is a fractal Brownian function if for all x and Δx

$$\Pr\left(\frac{I(x + \Delta x) - I(x)}{\|\Delta x\|^{1-r}} < y\right) = F(y) \quad (1)$$

²This ratio primarily depends on the spatial autocorrelation of the process that formed the surface.

³Because formative processes tend to act over a range of scales, real surfaces normally have a constant scaling parameter over fairly wide (e.g., 1:8) ranges of scale, although it is rare for it to be constant over several decades of scale. Thus if we observe that r at large scales is much different than at small scales (as in Fig. 1), we can reliably infer that two different processes were involved in forming the surface, and that they acted over different ranges of scale.

where $F(y)$ is a cumulative distribution function [7], and the variable r is the ratio r of the previous section. Note that x and $I(x)$ can be interpreted as vector quantities, thus providing an extension to two or more topological dimensions. If the topological dimension of $I(x)$ is T , the fractal dimension D of the graph described by $I(x)$ is $D = T + r$. If $r = \frac{1}{2}$ and $F(y)$ comes from a zero-mean Gaussian with unit variance, then $I(x)$ is the classical Brownian function.

It is important to note that (1) describes change over scale, but does *not* talk about pattern. Very highly patterned surfaces can be fractal; all that is required is that they scale appropriately. The description of pattern is *orthogonal* to the fractal description of (1).

2.1.1. The fractal model and imaging

Before we can use a fractal model of natural surfaces to help us understand images, we must determine how the imaging process maps a fractal surface shape into an image intensity surface. The first step is to define our terms carefully:

Definition 2.1. A *fractal Brownian surface* is a continuous function that obeys the statistical description given by (1), with x as a two-dimensional vector at all scales (i.e., values of Δx) between some smallest (Δx_{\min}) and largest (Δx_{\max}) scales.

Definition 2.2. A *spatially isotropic fractal Brownian surface* is a surface in which the components of the surface normal $N = (N_x, N_y, N_z)$ are themselves fractal Brownian surfaces of identical fractal dimension.

Note that “isotropic” in this usage does *not* mean that the surface cannot be stretched, like tree bark. What it means is that, up to a multiplicative (stretching) factor, the statistical process describing the surface is identical in all directions. Tree bark, for instance, appears to be well modeled by such an isotropic process.

The fractal Brownian surface is as constructed in Fig. 2. A spatially isotropic Brownian surface may be constructed by adding smaller bumps such that their height (axis of symmetry) is aligned with the normal to the existing surface, rather than always being aligned in the vertical direction.

Our previous papers [6, 7] have presented evidence showing that most (approximately 90%) natural surfaces are spatially isotropic fractals over scale ranges of at least 1:8 (i.e., $\Delta x_{\max}/\Delta x_{\min} > 8$). This finding has since been confirmed by others [8]. These findings imply that real surfaces often obey the fractal scaling law.

It is also interesting to note that practical fractal generation techniques, such as those used in computer graphics, have had to constrain the fractal generating

function to produce spatially isotropic fractal Brownian surfaces in order to obtain realistic imagery [13]. Thus, it appears that many real 3-D surfaces are spatially isotropic fractals, at least over a wide range of scales.

It must be emphasized that these findings make no implication about patterning of the surfaces—finding a surface to be “fractal” tells how different scales (spatial frequency channels, etc.) relate to one another, but not about the appearance at any one scale. What this finding *does* tell us is that the statistics of homogeneous surfaces normally vary over scale in such a way (i.e., exponentially) as to appear the same at each scale or resolution. This explains why it is so difficult to measure perspective gradients in real imagery: although things appear smaller with increasing distance, the (textural) statistics of the surface typically remain constant over wide ranges of scale.

2.1.2. *Estimation within homogeneous patches*

With these definitions in hand, we can now address the problem of how homogeneous patches of a 3-D fractal surface appear in the 2-D image.

Proposition 2.3. *A 3-D surface with a spatially isotropic fractal Brownian shape produces an image whose intensity surface is fractal Brownian and whose fractal dimension is identical to that of the components of the surface normal, given a Lambertian surface reflectance function and constant illumination and albedo.*

This proposition (proved in [7]) demonstrates that the fractal dimension of the surface normal dictates the fractal dimension of the image intensity surface and, of course, the dimension of the physical surface. Simulation of the imaging process with a variety of imaging geometries and reflectance functions indicates that this proposition will hold quite generally; the “roughness” of the surface seems to dictate the “roughness” of the image. Further, this proposition has proven to be an excellent predictor of people’s perception of 3-D surface roughness. Thus if we know that the surface is homogeneous (perhaps by examining color information), then we can estimate the fractal dimension of the surface by measuring the fractal dimension of the image data. What we have developed is a method for inferring a basic property of the 3-D surface—its fractal dimension—from the image data.⁴

⁴Experimental note: Fifteen naive subjects (mostly language researchers) were shown digitized images of eight natural textured surfaces drawn from [14]. They were asked “if you were to draw your finger horizontally along the surface pictured here, how rough or smooth would the surface feel?”—i.e., they were asked to estimate the 3-D roughness/smoothness of the viewed surfaces. A scale of one (smoothest) to ten (roughest) was used to indicate 3-D roughness/smoothness. The mean of the subject’s estimates of 3-D roughness had an excellent 0.91 correlation (i.e., 83% of the variance was accounted for $p < 0.001$) with roughnesses predicted by use of the image’s 2-D fractal dimension and Proposition 2.3. This result supports the general validity of Proposition 2.3.

2.1.3. Estimation from imaged contour

When a surface is not homogeneous it is better to use imaged contours to infer the fractal dimension of the surface. Such contours can also be used as independent confirmation of the estimate provided by Proposition 2.3. The following proposition addresses this issue.

Proposition 2.4. *The fractal scaling parameter r of an imaged contour is identical to that of the imaged 3-D surface given that (1) the surface is a spatially isotropic fractal, and (2) the 3-D shape of the contour generator is described by the intersection of a smooth function (fractal dimension = topological dimension) with the imaged surface.*

The fractal dimension of the locus of intersection between a fractal surface and a smooth surface is determined entirely by the properties of the fractal surface. Further, the fractal properties of a curve are unaffected by projection. Thus the fractal statistics of an imaged contour produced by cast shadow, occlusion, and so forth, are reliably indicative of the fractal statistics of the underlying surface. The requirement of isotropy does not exclude stretched surfaces such as tree bark; the isotropy applies only to the fractal scaling parameter of the surface excluding, e.g., a developable surface (generalized cylinder) with a fractal cross-section.

2.1.4. Practical measurement of the fractal dimension

The fractal dimension of a function can be measured either directly from the second-order statistics (dipole statics) of $I(x)$ by use of (1), or from $I(x)$'s Fourier power spectrum $P(f)$, as the spectral density of a fractal Brownian function is proportional to f^{2r-3} .

To measure the fractal dimension from the dipole statistics, we rewrite (1) to obtain the following description of the manner in which the second-order statistics of the image change with scale:

$$E(|\Delta I_{\Delta x}|) |\Delta x|^{r-1} = E(|\Delta I_{\Delta x=1}|),$$

where $E(|\Delta I_{\Delta x}|)$ is the expected value of the change in intensity over distance Δx . To estimate r , and thus D , we calculate the quantities $E(|\Delta I_{\Delta x}|)$ for various Δx , and use a least-squares regression on the log of our rewritten equation (1).

We may also measure the fractal dimension from the Fourier power spectrum. Since the power spectrum $P(f)$ is proportional to f^{2r-3} , we may use a linear regression on the log of the observed power spectrum as a function of f (e.g., a regression using $\log P(f) = (2r - 3) \log f + k$ for various values of f) to determine the power r and thus the fractal dimension $D = T + r$.

This power spectrum method suggests a method of measuring fractal dimen-

sion using physiologically plausible filters. The integral of the squared response of band pass filters (e.g., Laplacian or center-surround filters) is proportional to the amount of power (energy) in the Fourier spectrum that lies within the filter's sensitive region. Thus, if we take two filters that have equal volume in the Fourier domain but different center frequencies, then the log of the ratio of the squared response of such filters will be linearly related to the fractal dimension.

3. Shape Estimation and the Fractal Model

One of the major problems in using current shape-from-shading [2, 3], or shape-from-texture [4, 5] operators for recovering estimates of 3-D shape is that these techniques have been developed only for the case of smooth surfaces. They cannot, in general, be reliably applied to the rough, natural surfaces. Moreover, it has not even been possible to discriminate "smooth" surfaces from "rough" ones, or "shaded" surfaces from "textured" ones, so that we have not been able to determine what technique to apply, or been able to tell when we can expect a reliable answer.

By providing a unified mathematical description of both rough and smooth surfaces, the fractal model offers us a path out of these problems. The simplest application of the fractal model to this problem is to improve previously developed shape-from-shading and shape-from-texture methods by providing a method of checking some of their assumptions.

We have shown that people judge⁵ fractal functions with $r \approx 0$ to be "smooth," with gentle random undulations described by the function $F(y)$ in (1). In contrast, fractals with $r > 0$ are not perceived as smooth, but rather as being rough or three-dimensionally textured. The extreme case of $r \approx 1$ produces a nearly planar surface completely covered with jagged spikes; this is the salt-and-pepper Gaussian noise image standardly used as a model of imaged textures.

The fractal model, therefore, encompasses smooth shaded surfaces, three-dimensionally textured surfaces, and surfaces that are flat but covered with salt-and-pepper texture, with shading as a limiting case in the spectrum of 3-D texture granularity. The fractal model thus allows us to make a reasonable, rigorous and perceptually plausible definition of the categories "textured" versus "shaded," "rough" versus "smooth," in terms that can be measured by using the image data.

This can allow us to determine when to apply shading techniques, and when to apply 2-D texture techniques. For surfaces that are perceptually smooth ($r \approx 0$), it is appropriate to apply shading techniques; while for surfaces that have 2-D texture ($r \approx 1$) it is more appropriate to apply available texture measures.

For surfaces with $0 < r < 1$, however, the assumptions of the standard shape-from- x techniques are violated: the surface is neither smooth nor planar

⁵See earlier note about psychophysical evidence.

with Gaussian noise. We can use the fractal surface model to extend current techniques so that they address these three-dimensionally textured cases too. Further, because we have a single model that encompasses shading, 2-D texture, and 3-D textures, we can derive a shape estimation procedure that treats shaded, two-dimensionally textured, and three-dimensionally textured surfaces in a single, unified manner. The remainder of this paper will be addressed to this derivation.

3.1. Measuring properties intrinsic to the surface

Shape evidences itself within a small neighborhood by foreshortening, the apparent compression of surface features in the image that occurs along the direction of surface tilt (the image-plane component of surface orientation, i.e., the image direction in which the projected surface normal points). To form a local estimate of surface shape, therefore, we must be able to measure this apparent foreshortening for some feature intrinsic to the surface. For homogeneous, unmarked surfaces the only features that are intrinsic to the surface may all be expressed in term of the surface normal and its derivatives. To form an estimate of shape for such surfaces, therefore, we must be able to observe foreshortening of the surface normal, its derivatives, or functions of the surface normal and derivatives.

3.1.1. *A model of image formation*

To understand how we may observe foreshortening of the surface normal or derivatives, we employ a simple model of image formation that expresses image intensity in terms of the surface normal. This model makes two assumption: (i) albedo and illumination are constant in the neighborhood being examined, and (ii) the surface reflects light isotropically (Lambert's law, an idealization of matte, diffusely reflecting surfaces and of shiny surfaces in regions that are distant from highlights and specularities [3]). Figures 1 and 2(a)–(c) were generated by use of this model.

We are then led to this simple linear equation relating image intensity to the surface normal:

$$I = \rho\lambda(N \cdot L) \quad (2)$$

where ρ is surface albedo, λ is incident flux, N is the mean (three-dimensional) unit surface normal within the neighborhood being examined, and L is a (three-dimensional) unit vector pointing toward the weighted mean of all the illumination sources.

In (2), image intensity is dependent upon the surface normal, as all other variables have been assumed constant. Similarly, the second derivative of image intensity is dependent upon the second derivative of the surface normal, i.e.,

$$d^2I = \rho\lambda(d^2N \cdot L). \quad (3)$$

(*Notation.* We will write d^2I and d^2N to indicate the second derivative quantities computed along some image direction (dx, dy) —this direction to be indicated implicitly by the context.)

3.1.2. A texture measure

The problems in observing N directly are well known; perhaps even worse problems apply to observing dN , the first derivative of N (see [16]). For d^2N (the second derivative of the surface normal), however, a statistical regularity—that the vector d^2N is isotropically distributed—allows us to measure the intrinsic surface property d^2N , estimate the magnitude and direction of foreshortening, and thus estimate surface shape.

For surfaces of the type addressed in this paper the distribution of d^2N is isotropic in space. Thus the expected value of the vectors d^2N (measured for some particular direction (dx, dy)) will be the zero vector and therefore (by (3)) the mean of d^2I will be zero.

If we define the vector d^2N^+ to be

$$d^2N^+ = \begin{cases} d^2N, & \text{if } d^2N \cdot L \geq 0, \\ -d^2N, & \text{if } d^2N \cdot L < 0, \end{cases} \quad (4)$$

then we can⁶ describe the distribution of $|d^2I|$ by

$$|d^2I| = |\rho\lambda d^2N \cdot L| = \rho\lambda d^2N^+ \cdot L. \quad (5)$$

As d^2N is isotropically distributed the expected value of these vectors will then be

$$E_\theta(d^2N^+) = (0, 0, kE_\theta(\|d^2N\|)) = kE_\theta(\|d^2N\|)L, \quad (6)$$

where $\|d^2N\|$ denotes the magnitude of d^2N , and $E_\theta(d^2N)$ the expectation of d^2N , where the derivatives are taken along the direction θ . In this expression, k is a positive constant dependent upon the variance of the angle between d^2N and L .

Consequently, if we knew the intensity I_0 of a point with normal N_0 such that $N_0 = L$, that is, a point which faces the mean illuminant direction, then

$$E_\theta\left(\left|\frac{d^2I}{I_0}\right|\right) = \frac{E_\theta(\rho\lambda(d^2N^+ \cdot L))}{\rho\lambda(N_0 \cdot L)} \approx \frac{E_\theta(d^2N^+) \cdot L}{L \cdot L} = kE_\theta(\|d^2N\|). \quad (7)$$

This relation is not exact because of the differential effect of the illuminant across the angular spread of the distribution of d^2N .

⁶Assuming that the distribution of illumination is constant across the surface.

Because this is a measure of an intrinsic surface property, it may be used as a texture measure. It is affected by foreshortening which acts to increase the apparent frequency of variations in the surface, i.e., to increase the average magnitude of d^2N . We may, therefore, use this texture measure to estimate surface orientation.

3.2. Estimating the surface tilt

We start by picking a space constant σ which defines the resolution of our estimates of surface orientation. That is, if $z(x, y)$ is the original imaged surface, we will estimate the orientation of the $z^*(x, y) = G(\sigma, x, y) \otimes z(x, y)$ (where $G(\sigma, x, y)$ is a two-dimensional Gaussian of variance σ^2) by estimating foreshortening's contribution to the weighted mean value of our texture measure over the support of $G(\sigma, x, y)$. Thus expectations are estimated using convolutions with $G(\sigma, x, y)$, and we assume negligible change in the orientation and curvature of z^* within its support.

If we let (u, v) be a coordinate system tangent to the surface z^* , oriented so that u is the direction of the surface tilt, and let (x^*, y^*) be an image-plane coordinate system in which u projects only onto x^* , and v only onto y^* (i.e., (x^*, y^*) is a rotated version of the viewer's coordinates (x, y)), then

$$E\left(\frac{dN}{dx^*}\right) = E\left(\frac{dN}{du} \frac{du}{dx^*} + \frac{dN}{dv} \frac{dv}{dx^*}\right) = E\left(\frac{dN}{du}\right) \frac{du}{dx^*} \quad (8)$$

as $dv/dx^* = 0$. Similarly,

$$E\left(\frac{d^2N}{dx^{*2}}\right) = E\left(\frac{d^2N}{du^2} \left(\frac{du}{dx^*}\right)^2 + \frac{dN}{du} \frac{d^2u}{dx^{*2}} + \frac{dN}{dv} \frac{d^2v}{dx^{*2}}\right). \quad (9)$$

Equivalent equations hold for the y^* and v directions, except that $dv/dy^* = 1$, so that finally,

$$\begin{aligned} E(\nabla^2 N) = & E\left(\frac{d^2N}{du^2} \left(\frac{du}{dx^*}\right)^2 + \frac{d^2N}{dv^2}\right) \\ & + \left(\frac{dN}{du} + \frac{dN}{dv}\right) \left(\frac{d^2u}{dx^{*2}} + \frac{d^2v}{dx^{*2}} + \frac{d^2u}{dy^{*2}} + \frac{d^2v}{dy^{*2}}\right). \end{aligned} \quad (10)$$

Note that all of the terms of this expression are independently distributed. We can use these relations to obtain an estimate of the surface tilt (which is the image-plane component of the surface normal, i.e., the direction the surface normal would face if projected onto the image plane) by finding the gradient direction of $E(|\nabla^2 I/I_0|)$.

The surface tilt lies along the gradient direction under the conditions that (i)

the texture is homogeneous within the averaging patch⁷—although it may be stretched like wood bark, (ii) the third and crossed second derivatives of the averaged surface z^* are small,⁸ and (iii) the surface is not planar. Under these assumptions we may combine (7) and (10) to derive

$$\begin{aligned} \frac{\partial}{\partial x^*} E \left(\left| \frac{\nabla^2 I}{I_0} \right| \right) &\approx \frac{\partial}{\partial x^*} kE(\|\nabla^2 N\|) = 2kE \left(\left\| \frac{d^2 N}{du^2} \right\| \right) \frac{du}{dx^*} \frac{d^2 u}{dx^{*2}}, \\ \frac{\partial}{\partial y^*} E \left(\left| \frac{\nabla^2 I}{I_0} \right| \right) &= 0. \end{aligned} \quad (11)$$

Thus the gradient direction will be along either the direction x^* or $-x^*$ (depending upon the sign of $d^2 u/dx^{*2}$), which by definition is parallel to the direction of surface tilt.

The first assumption is a requirement of any texture analysis method, and thus presents no unusual restriction. Similarly, the second assumption seems only slightly restrictive. That the mean third derivatives are small follows from the properties of physical surfaces that exhibit fractal scaling behavior.⁹ That the mean crossed second derivative $d^2 u/dx dy$ is typically small is a result of the estimated surface z^* (and thus u, v) being defined by convolution with a smooth function of x and y , rather than of u and v . It must be noted, however, that these observations about the second assumption are only statistical in nature; it is straightforward to construct specific examples in which the local surface structure violates the assumption.

The final assumption—that the surface is not planar—is somewhat restrictive; however, it is clear that people suffer from a similar restriction. People tend to see planar anisotropic (stretched) textures as facing directly towards them unless there is a significant perspective gradient. This is part of the phenomena known as “regression toward the frontal plane.” It seems, therefore, that estimating the surface tilt by the gradient of $E(|\nabla^2 I/I_0|)$ is a reasonable procedure—even when the surface is “stretched,” like wood bark.

3.2.1. Albedo and illumination variations

This formulation still has two undesirable traits; one, it requires finding I_0 , something that is not always possible, and two, it makes very strong use of the assumptions of constant albedo and illumination. In natural scenes it is especially important to avoid dependence on assumptions of constant albedo

⁷I.e., $dE(d^2 N/du^2)/dx = dE(d^2 N/dv^2)/dx = dE(dN/du)/dx = dE(dN/dv)/dx = 0$ and similarly for dy .

⁸I.e., $d^2 u/dx dy = 0, d^3 u/dx^3 = d^3 u/dy^3 = \dots = 0$.

⁹Almost all natural physical surfaces have a fractal scaling parameter $0.5 > r \geq 0$, which carries the implication that the mean magnitude of the first derivatives is greater than that of the second derivatives, which in turn is greater than that of the third derivatives. This follows from the autocorrelation function of these surfaces.

and illumination, for such assumptions are very rarely correct and can introduce large errors. It has, therefore, been found advantageous to use the mean of the local intensities,¹⁰ rather than I_0 , as the divisor,¹¹ e.g.,

$$E_{\theta} \left(\left| \frac{d^2 I}{I} \right| \right) = E_{\theta} \left(\left| \frac{\rho \lambda (d^2 N \cdot L)}{\rho \lambda (N \cdot L)} \right| \right) \approx E_{\theta} (\|d^2 N\|). \quad (12)$$

This use of I rather than I_0 introduces a bias of I/I_0 when the albedo and illumination are actually constant. This is of no consequence for local estimation of surface orientation, as that involves ratios between measurements with the same bias, and thus the effects will divide out.

This bias will, however, affect later calculations that make comparisons between widely separated points. The motivation for accepting this bias is that by using the local mean of intensity we remove all dependence on the albedo and intensity of illumination; further, previous analyses have shown that the practical effect of this bias is often not significant [16, 19].

3.3. An alternate development

We derive our result in the following alternate, independent manner. First we note that the zero-crossing density of the Laplacian of an image may be used to estimate surface orientation, i.e., for isotropic, locally planar surfaces with surface markings that generate zero-crossings the tilt of the surface is the direction along which the maximum density of zero-crossing contours occurs [5].

Secondly, we note that

$$\lambda^2 \approx \text{Var} \left(\frac{\partial}{\partial \theta} \nabla^2 G \otimes I(x, y) \right), \quad (13)$$

where λ is the zero-crossing density along direction θ for the Laplacian of a zero-mean Gaussian process [15]. Using the fact that the Brownian fractal functions discussed here are zero-mean Gaussian processes we may combine these two results to obtain an independent derivation of our shape estimator.

First we show that our texture measure is in fact an estimator of zero-crossing density. For a Gaussian process $E(|x|) = 0.67 \text{Std}(x)$ where $\text{Std}(x)$ is the standard deviation of the zero-mean random variable x . Thus

$$\lambda^2 \approx (0.67)^2 E^2 \left(\left| \frac{\partial}{\partial \theta} \nabla^2 G \otimes I(x, y) \right| \right).$$

Noting that

$$\frac{\partial}{\partial \theta} \left| \nabla^2 G \otimes I(x, y) \right| = \pm \left| \frac{\partial}{\partial \theta} \nabla^2 G \otimes I(x, y) \right|$$

¹⁰I.e., $G(\sigma, x, y) \otimes I(x, y)$, where the space constant σ of G can be chosen to be the same as that of the filter $\nabla^2 G(\sigma, x, y)$, or (perhaps better) substantially larger.

¹¹Experiments on natural imagery show that $|d^2 I|$ is uncorrelated with I , thus $E(|d^2 I|/I) = E(|d^2 I|)/E(I)$. This also follows from Markov process models of imagery.

except at zero-crossings of $\nabla^2 G \otimes I(x, y)$, we see that

$$\frac{\partial}{\partial \theta} \left| \frac{\nabla^2 G \otimes I(x, y)}{I_0(x, y)} \right| \approx \frac{\pm 1}{I_0(x, y)} \left| \frac{\partial}{\partial \theta} \nabla^2 G \otimes I(x, y) \right| \quad (14)$$

and thus, as convolution (used to estimate expectation) and differentiation commute, we have that

$$\frac{\partial}{\partial \theta} E \left(\left| \frac{\nabla^2 G \otimes I(x, y)}{I_0(x, y)} \right| \right) \approx k \text{Std} \left(\frac{\partial}{\partial \theta} \nabla^2 G \otimes I(x, y) \right) \approx k\lambda, \quad (15)$$

where k is $\pm 0.67/I_0(x, y)$. Thus by (13) the gradient direction of our texture measure—i.e., the gradient direction of $E(|\nabla^2 I(x, y)/I_0(x, y)|)$ —is approximately parallel to the direction of maximum zero-crossing density.

Using the result that for locally planar surfaces with isotropic surface markings (i.e., markings that are not “stretched”) the direction of maximum frequency of zero-crossing density is also the direction of surface tilt, we see that the gradient direction of $E(|\nabla^2 I(x, y)/I_0(x, y)|)$ is approximately parallel to the surface tilt.

We may generalize this to rough surfaces by noting that the positions of zero-crossings in an image of a rough, 3-D textured surface are localized to within about $\frac{1}{2}\sigma_0$ of the true position of the surface’s inflection point, where σ_0 defines the finest resolution at which we view the surface $z(x, y)$. Thus, if σ (the resolution of the estimated surface $z^*(x, y)$) is much larger than σ_0 , then we may treat the zero-crossings produced by the 3-D texture as being fixed on the surface. Thus when the surface $z(x, y)$ is not anisotropically “stretched,” and $\sigma \gg \sigma_0$, we may apply the above result to correctly estimate the direction of surface tilt. As in the previous development, we may substitute the assumption that $z(x, y)$ is not stretched for the assumptions that the texture is homogeneous and that the third and crossed second derivatives are small.

3.4. Development of a robust slant estimator

Surface slant (the depth component of the surface normal) is much more difficult to determine, because it requires not only that we estimate the direction of foreshortening (which yields the tilt) but also the amplitude of the foreshortening.

For example, to estimate the foreshortening amplitude we may compare the value of our texture measure along the tilt direction to the value along the perpendicular direction in order to estimate the amount of foreshortening. This technique is, however, critically dependent upon the isotropy of the texture. Or we may look at the magnitude of the texture gradient to estimate the rate of change of the foreshortening. This method, however, is critically dependent upon knowing the curvature of the averaged surface.

There seems to be no really good way to locally estimate surface slant.

because estimation of foreshortening amplitude *requires* comparison to some base line estimate of the unforeshortened texture. Thus each method is critically dependent upon comparing texture density estimates at different points or along different directions, and so local variation in the surface texture is a major source of error in each slant estimation technique.¹² It is important, therefore, to develop a method of estimating surface slant that is robust with respect to local variation in the surface texture.

3.4.1. *Using regional constraints*

Progress towards such robustness can be obtained by applying regional, rather than purely local, constraints. Natural textures are often "homogeneous" over substantial regions of the image, although there may be significant local variation within the texture, because the processes that act to create a texture typically affect regions rather than points on a surface. This fact is the basis for interest in texture segmentation techniques. Current shape-from-texture techniques do not make use of the regional nature of textures, relying instead on point-by-point estimates. By capitalizing on the regional nature of textures we can derive a substantial additional constraint on our shape estimation procedure.

The most direct improvement provided by this additional constraint comes from using the statistical fact that a better estimate of the texture parameters (i.e., one independent of local variations) may be made from many samples than can be made from one sample.

We may illustrate this by an example. Let us assume that we are viewing a textured planar surface whose orientation is a 30° slant and a vertical tilt. Let us further suppose that the surface texture varies randomly from being isotropic to being anisotropic (stretched) up to an aspect ratio of 3:1, with the direction of this anisotropy also varying randomly. Such a surface, covered with small crosses, is shown in Fig. 3(a); for comparison, the same surface, minus variations, is shown in Fig. 3(b).

If we apply standard shape estimation techniques—i.e., estimating the amount of foreshortening (and thus surface orientation) by the ratio of some texture measure along the (apparently) unforeshortened and (apparently)

¹²Estimates of the fractal scaling parameter of the viewed surface [6, 7], by virtue of their independence with respect to multiplicative transforms, offer a partial solution to this problem. Because foreshortening is a multiplicative effect, the computed fractal scaling parameter is not affected by the orientation of the surface (at least not until self-occlusion effects have become dominant in the appearance of the surface). Thus, if we measure the fractal scaling parameter of an isotropically textured surface along the *x*- and *y*-directions, the measurements must be identical. If, however, we find that they are unequal, we then have *prima facie* evidence of anisotropy in the surface. This method of identifying anisotropic textures is most effective when each point on the surface has the same direction and magnitude of anisotropy, for in these cases we can accurately discriminate changes in fractal scaling parameter between the *x*- and *y*-directions. When the surface texture is variable, however, this indicator of anisotropy becomes less useful.




FIG. 3. (a) Variation in local texture; (b) same texture without variation.

maximally foreshortened directions—our estimates of the foreshortening magnitude will vary widely, with a mean error of 65% and an rms error of 81%. If, however, we estimate the value α of the unforeshortened texture measure by examining the entire region, and then compare this regional estimate to the texture measure along the (apparently) maximally foreshortened direction, then our mean error is reduced to 40% and the rms error to 49%. By estimating the texture parameter over the entire region, a better estimate is made than when estimating it separately at each point.

A second, equally important improvement in shape estimation provided by regional constraints comes from the requirement of consistency: at a bounding contour the principal direction of apparent foreshortening in the texture pattern must be orthogonal to the contour. If this regional consistency constraint is not met, then the pattern is anisotropic, and often the general direction of the anisotropic stretching can be determined.

3.4.2. Estimation of surface slant

We may construct an interesting slant estimation algorithm based on this notion of regional estimation and on the texture measure introduced above by employing the fact that

$$\nabla^2 I = \frac{d^2 I}{du^2} + \frac{d^2 I}{dv^2} \quad (16)$$

for any orthogonal u, v .

We will assume that (i) we have a texture region with a single, overall anisotropy, e.g., that the texture is "stretched" in approximately the same direction throughout the entire region, like, for instance, bark on a tree, and (ii)

the mean second derivatives of the averaged surface z^* are small relative to the mean first derivatives (a statistical property of surfaces that have fractal scaling parameter $0.5 > r > 0$).

We start by finding the direction θ_0 along which the texture region as a whole is most compressed, i.e., the direction that appears to be most foreshortened. If the mean normal throughout the region is not toward the viewer, this will affect the apparent average "stretch" of the texture. This overall bias can often be detected by the regional consistency constraint, and some adjustment to θ_0 made.

We may then form an estimate of the value of the texture measure along the most compressed (apparently most foreshortened) direction averaged throughout the region R ; for instance, one method of forming this estimate is as follows:

$$\alpha_1 = E_{R, \theta_0}(|d^2I/I|). \quad (17)$$

Similarly, we may form an estimate of the average value of the texture measure along the least compressed (apparently least foreshortened) direction within the overall region R ; e.g.,

$$\alpha_2 = E_{R, \theta_0 - \pi/2}(|d^2I/I|). \quad (18)$$

Other methods of estimating α_1 or α_2 are also available; for instance, when smooth occluding contours bound a texture region then the constraint that $N \cdot V \geq 0$ may be used to determine their values (see [16]).

Thus, letting θ be the angular difference between the tilt direction at a particular point and the direction of texture anisotropy within the entire region, we find that for an unforeshortened planar patch our texture measure will be

$$E_{\theta} \left(\left| \frac{d^2I}{I} \right| \right) \approx \alpha_1 \cos \theta + \alpha_2 \sin \theta. \quad (19)$$

Then by combining (10), (16) and (19) under the above assumptions, we obtain the result that

$$\frac{E(|\nabla^2 I/I|) - (-\alpha_1 \sin \theta + \alpha_2 \cos \theta)}{(\alpha_1 \cos \theta + \alpha_2 \sin \theta)} \approx \left(\frac{du}{dx} \right)^2 \quad (20)$$

from which we can estimate the surface slant σ using

$$\sigma \approx \cos^{-1} \left(\frac{du}{dx} \right)^{-1}. \quad (21)$$

When the texture is isotropic, then $\alpha_1 = \alpha_2$, and so

$$\sigma \approx \cos^{-1} \left(\frac{E(|\nabla^2 I/I|) - \alpha}{\alpha} \right)^{1/2} \quad (22)$$

where α is the regional estimate of the unforeshortened value of $E(|d^2 I/I|)$.

4. A Unified Treatment of Shading and Texture

The fractal surface model captures the intuitive notion that, if we examine a series of surfaces with successively less three-dimensional texture, eventually the surfaces will appear shaded rather than textured. This is illustrated by Fig. 1. Because the shape-from-texture technique developed here was built on the fractal model, we might expect that it too would degrade gracefully into a shape-from-shading method. This is in fact the case: this shape-from-texture technique is a generalization of the shape-from-shading technique previously developed by the author [16]. That is, we have developed a shape-from- x technique that applies to both texture and shading. (The shape-from-shading technique described in [16] also assumed that surface curvature was isotropic, so that $\alpha_1 = \alpha_2$. In [16] local averaging of the image measurements was not explicitly part of the shape estimation procedure; it was, however, introduced in order to produce depth estimates by integration of local surface orientation estimates. Thus the procedure used to derive the displayed shape estimates is identical to that developed here.)

4.1. Shading versus texture

This technique does not, however, apply *equally* to both texture and shading. Estimates from textured surfaces will be more reliable than those from shaded surfaces. A small patch of a smooth, shaded surface (such as top row of Fig. 1) has essentially only one surface orientation, Gaussian curvature, etc., because the parameters of the surface vary only slowly. In contrast, a convoluted 3-D surface, such as appears in the bottom row of Fig. 1, has a great many unrelated surface patches because the parameters of the surface vary rapidly. For the convoluted surface, therefore, the (on average, statistical) truth that for a 3-D fractal surface the mean second derivative vector $d^2 N$ is parallel to the viewer direction V may be relied upon because we know that with many independent observations our estimated mean will be close to the true mean.

For the smooth, shaded surface, however, we have only one observation and can have no such statistical assurance. Thus when we use the relationship between mean $d^2 N$ and V to estimate the shape of a smooth surface we are making the strong assumption that this relationship will hold not just "on average," but that it will hold *at each and every point*. For smooth surfaces, therefore, we will often be in error—although our assumption (which is equivalent to assuming that we are viewing a spherical surface) seems to be the

best a priori estimate we can make, and in fact often gives a reasonable estimate of shape (see [16, 17]).

The reliability of this shape estimator increases gradually as we proceed from completely smooth surfaces, to smoother "rolling" surfaces such as the middle row of Fig. 1, and finally reaching maximum reliability on rough, convoluted surfaces such as the bottom row of Fig. 1. This gradient of reliability is seen in human perception too; compare your perception of shape over these same images, for instance. This gradient of reliability has recently been demonstrated in controlled psychophysical testing [18].

This gradient of performance as we move from rough surfaces to smooth is observed when we apply this shape estimator to natural imagery. Figure 4(a)




FIG. 4. (a) The image of a log, together with the relief map generated from the shape algorithm's estimates of surface orientation, (b) the image of a rock, together with the relief map generated from the shape algorithm's estimates of surface orientation..

shows the image of a log, together with the relief map generated from the shape algorithm's estimates of surface orientation. Note that both the surface and the texture in this figure are not isotropic. Figure 4(b) shows the image of a rock, together with the relief map generated from the shape algorithm's estimates of surface orientation. The relief maps in Figs. 4(a) and 4(b) correspond closely to the actual shapes of these two objects. This is as expected for such rough, complex surfaces.

In contrast, low accuracy is often obtained on smooth surfaces such the digitized picture of a face shown in Fig. 5(a). Figure 5(b) shows a relief map of the surface slant estimated for that image (eye and eyebrow regions were masked out by hand). No relief map of the estimated surface shape could be obtained from this surface because of difficulties in integrating the slant and tilt




FIG. 5. (a) The digitized picture of a small portion of a face, (b) a relief map of the surface slant estimated for that image (eye and eyebrow regions were masked out by hand).

estimates. In this slant map representation regions with higher relief face toward the viewer, while lower relief regions face away. Note, however, that many important details of the surface shape are still apparent; for instance, the structure of the nose, the cheeks, and the eyebrow ridges is plainly visible.

4.2. The problem of scale

We have seen that this estimator is most reliable when the estimated surface $z^*(x, y)$ is "locally planar" or "smoothly varying" and the full-resolution surface is rough and crenulated. This, then, raises the problem of scale: given that we know this performance characteristic of the estimator, at what resolution (scale) should we attempt to estimate the surface shape?

The solution to this scale problem can be seen if we state the reliability conditions more formally, in terms of the fractal scaling parameter r . The reliability conditions are: the texture measure is most reliable when $r \approx 0$ for the resolution at which we estimate shape, and when $r \gg 0$ at finer resolutions.

We can, therefore, determine the best scale for shape estimation by examining the surface's scaling parameter r over the entire range of available resolutions (see Section 2.1.4 on measurement of the fractal scaling parameter). That resolution which best satisfies the above criterion for reliability is the best scale for shape estimation. For example in Fig. 1 the best scale at which to estimate shape is the finest scale at which the surface still looks like a sphere: the finest scale where r is small. At finer scales $r > 0$, and so is best considered as "texture" to be used for estimating larger-scale "shape."

Note that this "best" resolution is normally not constant over the entire image. If we view, for example, the rolling hills covered with grass then the best scale for shape estimation will be the minimum resolution at which we can no longer see the individual blades of grass: the resolution at which we transition from the $r \gg 0$ of the grass, to the $r \approx 0$ of the rolling hills.

5. Summary

Shape-from-shading and texture methods have suffered from the lack of a representation for complex, natural scenes: they have been applicable only to smooth surfaces, while real surfaces are often rough and crumpled. We have argued that the fractal model can help remedy these problems because it seems to be a good model of these natural surface shapes: many basic physical processes produce fractal surfaces, and moreover fractal surfaces also *look* like natural surfaces.

We have used the fractal model to extend previous shape-from- x methods to complex, natural surfaces. The fractal model's ability to distinguish successfully between perceptually "smooth" and perceptually "rough" surfaces, for instance, may permit the reliable application of shape estimation techniques that assume smoothness. More importantly, however, the fact that the model

describes both smooth and rough surfaces within the same framework has allowed us to construct a method of estimating 3-D shape that treats shading and texture in a unified manner.

REFERENCES

1. Barrow, H.G. and Tenenbaum, J.M., Recovering intrinsic scene characteristics from images, in: A. Hanson and E. Riseman (Eds.), *Computer Vision Systems* (Academic Press, New York, 1978).
2. Horn, B.K.P., Shape from shading: a method for obtaining the shape of a smooth opaque object from one view. AI Tech. Rept. 79, Project MAC, MIT, Cambridge, MA, 1970.
3. Ikeuchi, K. and Horn, B.K.P., Numerical shape from shading and occluding boundaries, *Artificial Intelligence* 17 (1981) 141–184 (Special Volume on Computer Vision).
4. Kender, J.R., Shape from texture: an aggregation transform that maps a class of textures into surface orientation, in: *Proceedings Sixth International Joint Conference on Artificial Intelligence*, Tokyo, Japan (1979).
5. Witkin, A.P., Recovering surface shape and orientation from texture, *Artificial Intelligence* 17 (1981) 17–47 (Special Volume on Computer Vision).
6. Pentland, A., Fractal-based description, in: *Proceedings Eighth International Joint Conference on Artificial Intelligence*, Karlsruhe, F.R.G. (1981) 973–981.
7. Pentland, A., Fractal-based description of natural scenes, *IEEE Trans. Pattern Anal. Machine Intelligence* 6 (1984) 661–674.
8. Medioni, G. and Yasumoto, Y., A note on using the fractal dimension for segmentation, *IEEE Computer Vision Workshop*, Annapolis, MD, 1984.
9. Mandelbrot, B.B., *Fractals: Form, Chance and Dimension* (Freeman, San Francisco, CA, 1977).
10. Mandelbrot, B.B., *The Fractal Geometry of Nature* (Freeman, San Francisco, CA, 1982).
11. Pentland, A., Perception of three-dimensional textures, *Investigative Ophthalmology and Visual Science* 25 (1984) 201.
12. Richardson, L.F., The problem of contiguity: an appendix of statistics of deadly quarrels, in: *General Systems Yearbook* 6 (1961) 139–187.
13. Fournier, A., Fussel, D. and Carpenter, L., Computer rendering of stochastic models, *Commun. ACM* 25 (1982) 371–384.
14. Brodatz, P., *Textures: A Photographic Album for Artists and Designers* (Dover, New York, 1966).
15. Kass, M. and Witkin, A., Analyzing oriented patterns, in: *Proceedings Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA (1985) 944–952.
16. Pentland, A.P., Local shape analysis, *IEEE Trans. Pattern Anal. Machine Intelligence* 6 (1984) 170–187.
17. Yu, Sheng Hsuan, Implementation of shape-from-shading algorithms, Tech. Rept. Image Understanding Research, Intelligent Systems Group, UCLA, Los Angeles, CA. DARPA order No. 3119, 1983.
18. Todd, J.T. and Mingolla, E., The perception of surface curvature and direction of illumination from patterns of shading, *J. Experimental Psychol. Human Perception Performance* 9 (1983) 583–595.

Received August 1985; revised version received January 1986.