SHADING INTO TEXTURE

Alex P. Pentland
Artificial Intelligence Center, SRI International
333 Ravenswood Ave., Menlo Park, California 94025

ABSTRACT

Shape-from-shading and shape-from-texture methods have the serious drawback that they 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 [Pentland 83] 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 — thus allowing the reliable application of shape estimation techniques that assume smoothness. This model of surface shape has been used to derive a technique for 3-D shape estimation that treats shading and texture in a unified manner.

I. 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-texture [4,5] methods make similar assumptions: their models are concerned either with markings on a smooth surface, or discard three-dimensional notions entirely and deal only with ad hoc measurements of the image. 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).

* 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)

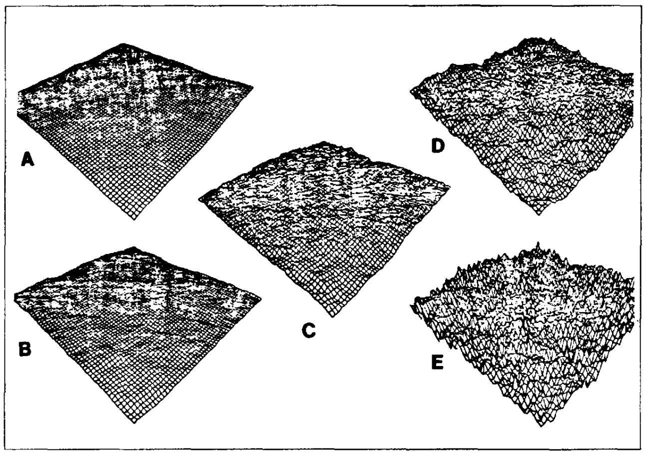

Figure 1. Surfaces of Increasing Fractal Dimension.

The fractal model of surface shape [6,7] appears to possess the required properties. Evidence for this comes from recently conducted surveys of natural imagery [6,8]. These survey found that the fractal model of imaged 3-D surfaces furnishes an accurate description of most textured and shaded image regions. Perhaps even more convincing, however, is the fact that fractals *look* like natural surfaces [9,10,11]. This is important information for workers in computer vision, because the natural appearance of fractals is strong evidence that they capture all of the perceptually relevant shape structure of natural surfaces.

II. FRACTALS AND THE FRACTAL MODEL

During the last twenty years, Benoit B. Mandelbrot has developed and popularized a relatively novel class of mathematical functions known as *fractals* [9,10]. Fractals are found extensively in nature [9,10,12]. Mandelbrot, for instance, shows that fractal surfaces are produced by many basic physical processes. The defining characteristic of a fractal is that it has a ***fractional dimension***, from which we get the word "fractal." One general characterization of fractals is that they are the end result of physical processes that modify shape through local action. After innumerable repetitions, such processes will typically produce a fractal surface shape.

The fractal dimension of a surface corresponds quite closely to our intuitive notion of roughness. 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, as is shown in Figure 1: (a) shows a flat plane ($D \approx 2$), (b) rolling countryside ($D \approx 2.1$), (c) an old, worn mountain range ($D \approx 2.3$), (d) a young, rugged mountain range ($D \approx 2.5$), and, finally (e), a stalagmite-covered plane ($D \approx 2.8$).

EXPERIMENTAL NOTE: ***Ten naive subjects (natural-***

language researchers) were shown sets of fifteen 1-D curves and 2-D surfaces with varying fractal dimension but constant range (e.g., see Figure 1), and asked to estimate roughness on a scale of one (smoothest) to ten (roughest). The mean of the subject's estimates of roughness had a nearly perfect 0.98 correlation (i.e., 96% of the variance was accounted for) ($p < 0.001$) with the curve's or surfaces's fractal dimension. The fractal measure of perceptual roughness is therefore almost twice as accurate as any other reported to date, e.g., [13].

Fractal Brownian Functions. Virtually all fractals encountered in physical models have two additional properties: (1) each segment is statistically similar to all others; (2) they are statistically invariant over wide transformations of scale. The path of a particle exhibiting Brownian motion is the canonical example of this type of fractal; 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

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

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

The fractal dimension of these functions can be measured either directly from $I(x)$ by using* of Equation 1, or from $I(x)$'s Fourier power spectrum** $P(f)$, as the spectral density of a fractal Brownian function is proportional† to f^{-2H-1}.

Properties of Fractal Brownian Functions. Fractal functions must be stable over common transformations if they are to be useful as a descriptive tool. Previous reports [6,7] have shown that the fractal dimension of a surface is invariant with respect to linear transformations of the data and to transformations of scale. Estimates of fractal dimension, therefore, may be expected to remain stable over smooth, monotonic transformations of the image data and over changes of scale.

A. The Fractal Surface Model And The Imaging Process

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: A **fractal Brownian surface** is a continuous function that obeys the statistical description given by Equation (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: A **spatially isotropic fractal Brownian surface** is a surface in which the components of the surface normal $\mathbf{N} = (\mathbf{N}_x, \mathbf{N}_y, \mathbf{N}_z)$ are themselves fractal Brownian surfaces of identical fractal dimension.

Our previous papers [6,7] have presented evidence showing that most natural surfaces are spatially isotropic fractals, with Δx_{min} and Δx_{max} being the size of the projected pixel and the size of the examined surface patch, respectively. This finding has since been confirmed by others [8]. Furthermore, it is 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 [11]. Thus, it appears that many real 3-D surfaces are spatially isotropic fractals, at least over a wide range of scales* .

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

Proposition 1. 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. If we know that the surface is homogeneous,** we can estimate the fractal dimension of the surface by measuring the fractal dimension of the image data. What we have developed, then, is a method for inferring a basic property of the 3-D surface — i.e., its fractal dimension — from the image data. The fact that fractal dimension has also been shown to correspond closely to our intuitive notion of roughness confirms the fundamental importance of the measurement.

EXPERIMENTAL NOTE: *Fifteen naive subjects (mostly language researchers) were shown digitized images of eight natural textured surfaces drawn from Brodatz [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 accounted was for) ($p < 0.001$) with roughnesses predicted by use of the image's 2-D fractal dimension and Proposition 1. This result supports the general validity of Proposition 1.*

B. Identification of Shading Versus Texture

Fractal functions with $H \approx 0$ do not change their statistics as a function of scale. Such surfaces are planar except for random variations described by the function $F(y)$ in Equation (1). If the variance of $F(y)$ is small people judge these surfaces to be "smooth"; thus, the fractal model with small values of H is appropriate for modeling

*We rewrite Equation (1) to obtain the following description of the manner in which the second-order statistics of the image change with scale: $E(|\Delta I_{\Delta x}|)\|\Delta x\|^{-H} = E(|\Delta I_{\Delta x=1}|)$ where $E(|\Delta I_{\Delta x}|)$ is the expected value of the change in intensity over distance Δx. To estimate H, 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).

**That is, since the power spectrum $P(f)$ is proportional to f^{-2H-1}, we may use a linear regression on the log of the observed power spectrum as a function of f (e.g., a regression using $\log(P(f)) = -(2H+1)\log(f)+k$ for various values of f) to determine the power H and thus the fractal dimension.

†Discussion of the rather technical proof of this proportionality may be found in Mandelbrot [10].

*This does not mean that the surfaces are completely isotropic, mearly that their fractal (metric) properties are isotropic.

**Perhaps determined by the use of imaged color.

smooth, shaded regions of the image. If the surface has significant local fluctuations, i.e., if $F(y)$ is large, the surface is seen as being smooth but textured, in the sense that markings or some other 2-D effect is modifing the appearance of the underlying smooth surface. In contrast, fractals with $H > 0$ are not perceived as smooth, but rather as being rough or three-dimensionally textured.

The fractal model can therefore encompass shading, 2-D texture, and 3-D 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.

The ability to differentiate between "smooth" and "rough" surfaces is critical to the performance of current shape-from-shading and shape-from-texture techniques. For surfaces that, from a perceptual standpoint, are smooth ($H \approx 0$) and not 2-D textured ($Var(F(y))$ small), it seems appropriate to apply shading techniques.* For surfaces that have 2-D texture it is more appropriate to apply available texture measures. Thus, use of the fractal surface model to infer qualitative 3-D shape (namely, smoothness/roughness), has the potential of significantly improving the utility of many other machine vision methods.

III. Shape Estimates From Texture And Shading

The fractal surface model allows us to do quite a bit better than simply identifying smooth *versus* textured surfaces and applying previously discovered techniques. Because we have a unified model of shading, 2-D texture and 3-D texture, we can derive a shape estimation procedure that treats shaded, two-dimensionally textured, and three-dimensionally textured surfaces in a single, unified manner.

A. Development of a Robust Texture Measure

Let us assume that: (1) albedo and illumination are constant in the neighborhood being examined, and (2) the surface reflects light isotropically (Lambert's law). We are then led to this simple model of image formation:

$$I = \rho\lambda(\mathbf{N} \cdot \mathbf{L}) \tag{2}$$

where ρ is surface albedo, λ is incident flux, $\mathbf{N}$ is the [three-dimensional] unit surface normal, and $\mathbf{L}$ is a [three-dimensional] unit vector pointing toward the illuminant. The first assumption means that the model holds only within homogeneous regions of the image, e.g., regions without self-shadowing. The second assumption is an idealization of matte, diffusely reflecting surfaces and of shiny surfaces in regions that are distant from highlights and specularities [3].

In Equation (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.,

$$d^2I = \rho\lambda(d^2\mathbf{N} \cdot \mathbf{L}) \tag{3}$$

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

The fractal model taken together with previous results [15], implies that on average $d^2\mathbf{N}$ is parallel to $\mathbf{N}$. Consequently, if we divide Equation (2) by Equation (3) we will *on average* obtain the following relationship:

$$E\left(\left|\frac{d^2I}{I}\right|\right) = E\left(\left|\frac{\rho\lambda(d^2\mathbf{N}\cdot\mathbf{L})}{\rho\lambda(\mathbf{N}\cdot\mathbf{L})}\right|\right) \approx E(\|d^2\mathbf{N}\|) \tag{4}$$

where $E(x)$ denotes the expected value [mean] of x. That is, we can estimate how crumpled and textured the surface is (i.e., the average magnitude of the surface normal's second derivative) by observing $E(|d^2I/I|)$.

Equation (4) provides us with a measure of 3-D texture that is (on average and under the above assumptions) independent of illuminant effects. This measure is affected by foreshortening, however, which acts to increase the apparent frequency of variations in the surface, e.g., the average magnitude of $d^2\mathbf{N}$. We can, therefore, obtain an estimate of surface orientation by employing the approach adopted in other texture work [5]: if we assume that the 3-D surface texture is isotropic, the surface tilt* is simply the direction of maximum $E(|d^2I/I|)$ and the surface slant** can be derived from the ratio between $\max_\theta E(|d^2I/I|)$ and $\min_\theta E(|d^2I/I|)$, where θ designates the [implicit] direction along which the texture measure is evaluated. Specifically, the surface slant is the arc cosine of z_N, the z-component of the surface normal, and for isotropic textures z_N is equal to the square root of this ratio. The square-root factor is necessitated by the use of second-derivative terms.

One of the advantages of this shape-from-texture technique is that not only can it be applied to the 2-D textures addressed by other researchers [4,5] (by simply using this texture frequency measure in place of theirs†), but it can also be applied to surfaces that are three-dimensionally textured — and in exactly the same manner. This texture measure, therefore, allows us to extend existing shape-from-texture methods beyond 2-D textures to encompass 3-D textures as well.

B. Development of a Robust Shape Estimator

These shape-from-texture techniques are critically dependent upon the assumption of isotropy: when the textures are anisotopic (stretched), the error is substantial. Estimates of the fractal dimension 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 dimension is not affected by the orientation of the surface.†† Thus, if we measure the fractal dimension 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 dimension between the x and y directions. When the surface texture is variable, however, this indicator of anisotropy becomes less useful. Thus, local variation in the surface texture remains a major source of error in our estimation techniques; it is therefore important to develop a method of estimating surface orientation that is robust with respect to local variation in the surface texture.

*Indeed, it is only in these cases that measurement noise can be reduced (by averaging) to the levels required by shape-from-shading techniques without simultaneously destroying evidence of surface shape.

*The image-plane component of the surface normal, i.e., the direction the surface normal would face if projected onto the image plane.

**The depth component of the surface normal.

†This measure includes edge information, i.e., the frequency of Marr-Hildreth zero-crossings as we move in a given direction appears to be proportional to $E(|d^2I/I|)$ along that direction; consider that Marr-Hildreth zero-crossings are also zero-crossings of d^2I/I.

††At least not until self-occlusion effects have become dominant in the appearance of the surface.

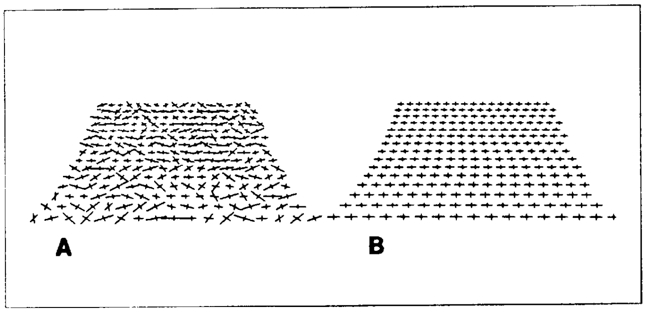

Figure 2. Variation in Local Texture (a) Compared with No Variation (b).

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.

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 Figure 2(a); for comparison, the same surface, minus anisotropies, is shown in Figure 2(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] unforshortened and [apparently] 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 unforshortened 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 combining this notion of regional estimation with the texture measure developed above, i.e., $E(|d^2 I/I|)$, we can construct the following shape-from-texture algorithm that is able to deal with both smooth two-dimensionally textured surfaces and rough, three-dimensionally textured surfaces, and that is robust with respect to local variations in the surface texture.

C. A Shape Estimation Algorithm

We may construct a rather elegant and efficient shape estimation algorithm based on the notion of regional estimation and on the texture measure introduced above by employing the fact that

$$\nabla^2 I = \frac{d^2 I}{du^2} + \frac{d^2 I}{dv^2} \tag{5}$$

for any orthogonal u, v. This identity will allow us to estimate the surface slant immediately rather than having to search all orientations for the directions along which we obtain the maximum and minimum values of $E(|d^2 I/I|)$.

Let us assume that we have already determined $\alpha = \min_\theta E(|d^2 I/I|)$, which is the regional estimate of unforeshortened $E(|d^2\mathbf{N}|)$. When the estimate of α is exact, Equation (5) gives us the

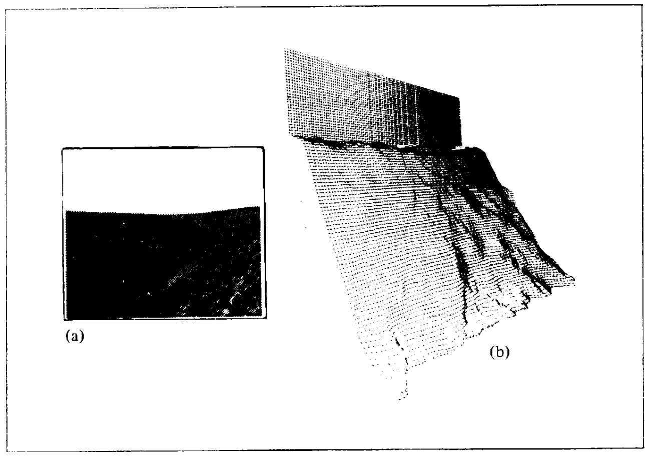

Figure 3. Tuckerman's Ravine.

result that

$$E\left(\left|\frac{\nabla^2 I}{I}\right|\right) - \alpha = \max_\theta E\left(\left|\frac{d^2 I}{I}\right|\right) \tag{6}$$

as the directions of maximum and minimum $E\left(\left|\frac{d^2 I}{I}\right|\right)$ are orthogonal.

We may therefore estimate z_N, the z component of the surface normal, by

$$z_N = \left(\frac{\beta - \alpha}{\alpha}\right)^{-1/2} \tag{7}$$

where $\beta = E(|\nabla^2 I/I|)$ and α is the regional estimate of the unforeshortened value of $E(|d^2 I/I|)$. The constant α can be estimated either by the median of the local [apparently] unforeshortened texture-measure values, or by use of the constraint that $0 \leq z_N \leq 1$ within the region. The direction of surface tilt can then be estimated by the gradient of the resulting slant field — e.g., the local gradient of the z_N values — or (as in other methods) by examining each image direction to find the one with the largest-value of the texture frequency measure. In actual practice we have found that the gradient method is more stable.

D. 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. 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 identical to the local shape-from-shading technique previously developed by the author [15]. That is, we have developed a shape-from-x technique that applies equally to 2-D texture, 3-D texture and shading.

As an example of the application of this shape-from-texture-and-shading technique,* Figure 3 shows (a) the digitized image of Tuckerman's ravine (a skiing region on Mt. Washington in New Hampshire), and (b) a relief map giving a side view of the estimated surface shape, obtained by integrating the slant and tilt estimates.**

*This example was originally reported in Pentland [15] as the output of a local shape-from-shading technique followed by averaging and integration. This algorithm is identical to the shape-from-texture technique described here; in fact, investigation of the shape-from-texture properties of this method was motivated by the consternation caused by this successful application of a *shading* technique to a *textured* surface.

This relief map may be compared directly with a topographic map of the area; when we compare the estimated shape with the actual shape, we find that the roll-off at the top of Figure 3(b) and the steepness of the estimated surface are correct for this surface; the slope of this area of the ravine averages 60°.

IV. Summary

Shape-from-shading and texture methods have had the serious drawback that they are applicable only to smooth surfaces, while real surfaces are often rough and crumpled. We have extended these methods to real surfaces using the fractal surface model [6,7]. The fractal model's ability to distinguish successfully between perceptually "smooth" and perceptually "rough" surfaces allows reliable application of shape estimation techniques that assume smoothness. Furthermore, we have used the fractal surface model to construct a method of estimating 3-D shape that treats shading and texture in a unified manner.

REFERENCES

[1] H.G. Barrow and J.M. Tenenbaum, "Recovering Intrinsic Scene Characteristics From Images," in A. Hanson and E. Riseman, Eds., *Computer Vision Systems,* Academic Press, New York, New York 1978

[2] B. K. P. H. Horn, "Shape From Shading: A Method for Obtaining the Shape of a Smooth Opaque Object from One View," A.I. Technical Report 79, Project MAC, M.I.T. (1970).

[3] B. K. P. H. Horn and K. Ikeuchi, "Numerical Shape from Shading and Occluding Boundaries," Artificial Intelligence, 15, Special Issue on Computer Vision, pp. 141-184 (1981).

[4] J. R. Kender, "Shape From Texture: An Aggregation Transform that Maps a Class of Textures Into Surface Orientation," Proceedings of the Sixth International Joint Conference on Artificial Intelligence, Tokyo, Japan (1979).

[5] A. P. Witkin, "Recovering Surface Shape and Orientation from Texture," Artificial Intelligence, 17, pp. 17-47 (1981).

[6] A. Pentland, "Fractal-Based Description," Proceedings of International Joint Conference on Artificial Intelligence (IJCAI) '83, Karlsruhe, Germany, August 1983.

[7] A. Pentland, "Fractal-Based Description Of Natural Scenes," IEEE Transactions on Pattern Analysis and Machine Intelligence, *to appear* September 1984.

[8] G. Medioni and Y. Yasumoto, "A Note on using the Fractal Dimension for Segmentation," IEEE Computer Vision Workshop, Annapolis, MD, April 30- May 3, 1984.

[9] B. B. Mandelbrot, "Fractals: Form, Chance and Dimension," W. H. Freeman and Co.,San Francisco, California, 1977.

[10] B. B. Mandelbrot, "The Fractal Geometry of Nature," W. H. Freeman, San Francisco, 1982.

[11] A. Fournier, D. Fussel and L. Carpenter, "Computer Rendering of Stochastic Models," Communications of the ACM, vol. 25, 6, pp. 371-384, 1982.

[12] L. F. Richardson, "The Problem of Contiguity: an Appendix of Statistics of Deadly Quarrels," General Systems Yearbook, vol. 6, pp. 139-187, 1961.

[13] H. Tamura, S. Mori, and T. Yamawaki, "Textural Features Corresponding to Visual Perception," IEEE Trans. on Sys., Man and Cyber., Vol. SMC-8, No. 6, pp.460-473, June 1978

[14] P. Brodatz, "Textures: A Photographic Album for Artists and Designers," Dover, New York, New York, 1966.

**The shape algorithm produces estimates of the surface orientation. For display purposes, these estimates were integrated to produce a relief map of the surface.

[15] Pentland, A. P. (1984), "Local Shape Analysis," IEEE Transactions on Pattern Analysis and Machine Intelligence, March 1984, pp. 170-187