SRI International

FRACTAL-BASED DESCRIPTION OF NATURAL SCENES

SRI Technical Note No. 280

March 20, 1984

By: Alex P. Pentland, Computer Scientist

Artificial Intelligence Center
Computer Science and Technology Division

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

SRI International

333 Ravenswood Ave. • Menlo Park, CA 94025
(415) 326-6200 • TWX: 910-373-2046 • Telex: 334-486

FRACTAL-BASED DESCRIPTION OF NATURAL SCENES

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

ABSTRACT

This paper addresses the problems of (1) representing natural shapes such as mountains, trees and clouds, and (2) computing their description from image data. To solve these problems, we must be able to relate natural surfaces to their images; this requires a good model of natural surface shapes. Fractal functions are a good choice for modeling 3-D natural surfaces because (1) many physical processes produce a fractal surface shape, (2) fractals are widely used as a graphics tool for generating natural-looking shapes, and (3) a survey of natural imagery has shown that the 3-D fractal surface model, transformed by the image formation process, furnishes an accurate description of both textured and shaded image regions.

The 3-D fractal model provides a characterization of 3-D surfaces and their images for which the appropriateness of the model is verifiable. Furthermore, this characterization is stable over transformations of scale and linear transforms of intensity.

The 3-D fractal model has been successfully applied to the problems of (1) texture segmentation and classification, (2) estimation of 3-D shape information, and (3) distinguishing between perceptually "smooth" and perceptually "textured" surfaces in the scene.

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

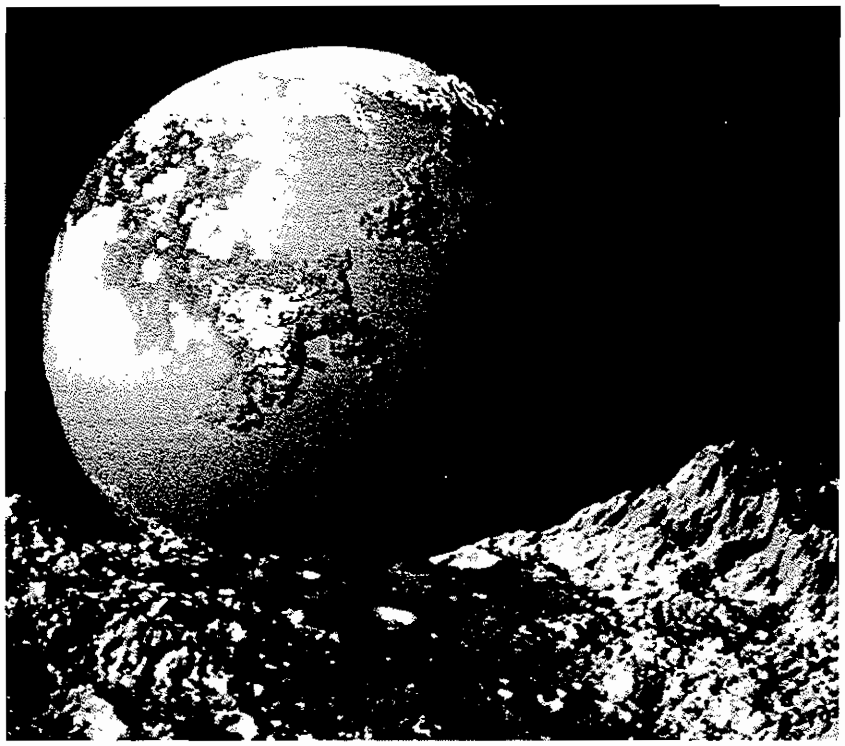

Figure 1. Fractal-Based Models of Natural Shapes (by Mandelbrot and Voss [10]).

I. INTRODUCTION

The world that surrounds us, except for man-made environments, is typically formed of complex, rough and jumbled surfaces (e.g., Figure 1). If we are to develop machines competent to deal with the natural world, therefore, we need a representational framework that is able to describe such shapes succinctly. The problem, then, is how shall we describe the shape of a crumpled newspaper? A clump of leaves? A jagged mountain?

Current representational schemes employ Plato's notion of Ideal Forms — e.g., spheres, cylinders and cubes — to describe three-dimensional shapes. Such shape-primitive representations function well in man-made, carpentered environments. When we attempt to describe the crenulated, crumpled surfaces typical of natural objects, however, the result is usually implausibly complex. Such awkwardness makes these shape-primitive representations difficult to envisage as the basis for human-performance-level capabilities.

Furthermore, how can we expect to extract 3-D information from the image of a rough or crumpled surface when all of our models refer to smooth surfaces only? We have no models that describe either the shape of such complex surfaces or how they evidence themselves in an image. 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.

Shape-from-shading [2,3] and surface interpolation methods [4], for instance, all employ the heuristic of "smoothness" to relate neighboring points on a surface. Such heuristics are applicable to many man-made surfaces, of course, but are demonstrably untrue of most natural surfaces. Texture descriptors, similarly, have dealt only with patterns assumed to lie on a smooth surface [5,6], or have discarded 3-D notions entirely and worked only with *ad hoc* statistical measures of the image intensity surface. Before we can 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.

In either case, we must have recourse to a 3-D model competent to describe naturally occurring surfaces. A good model of natural surfaces, together with the physics of image formation, would provide the analytical tools necessary for relating natural surfaces to their images. A formulation able to relate image to surface can provide the necessary leverage for usefully representing natural surfaces, as well as the computation of such descriptions from the image data[1] .

This paper, therefore, addresses two related problems: (1) finding a representation of shape capable of describing succinctly the surfaces of such natural objects as mountains, trees, and clouds, and (2) determining how such a description might be computed, given only raw image data. The first step towards solving these problems, of course, is to obtain a model of natural surface shapes.

Fractal functions appear to provide such a model, in part because many basic physical processes produce fractal surfaces (and thus fractals are quite common in nature), but perhaps even more importantly because fractals *look* like natural surfaces. This natural appearance has spurred recent computer graphics research to focus on using fractal processes[2] for simulating natural shapes and textures (as in Figure 1). Mountains, clouds, water, plants, trees, and even primitive animals [7-13] are all among the objects that have been realistically portrayed by use of fractal functions. 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.

Additional support for the fractal model comes from a recently conducted survey of natural imagery [14]. This survey found that the fractal model of imaged 3-D surfaces furnishes an accurate description of most textured and shaded image regions, thereby further validating this physics-derived model for both image texture and shading.

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* [7,10]. Fractals are found extensively in nature [7,8,10]. Mandelbrot, for instance, shows that fractal surfaces are produced by a number of basic physical processes, ranging from the aggregation of galaxies to the curdling of cheese.

[1]The word *representation* will be used to refer to the formal language in which *descriptions* of particular objects are couched.

[2]Most computer graphics techniques actually employ a stochastic approximation of true fractal functions [9]; however, this distinction is not important for our purposes.

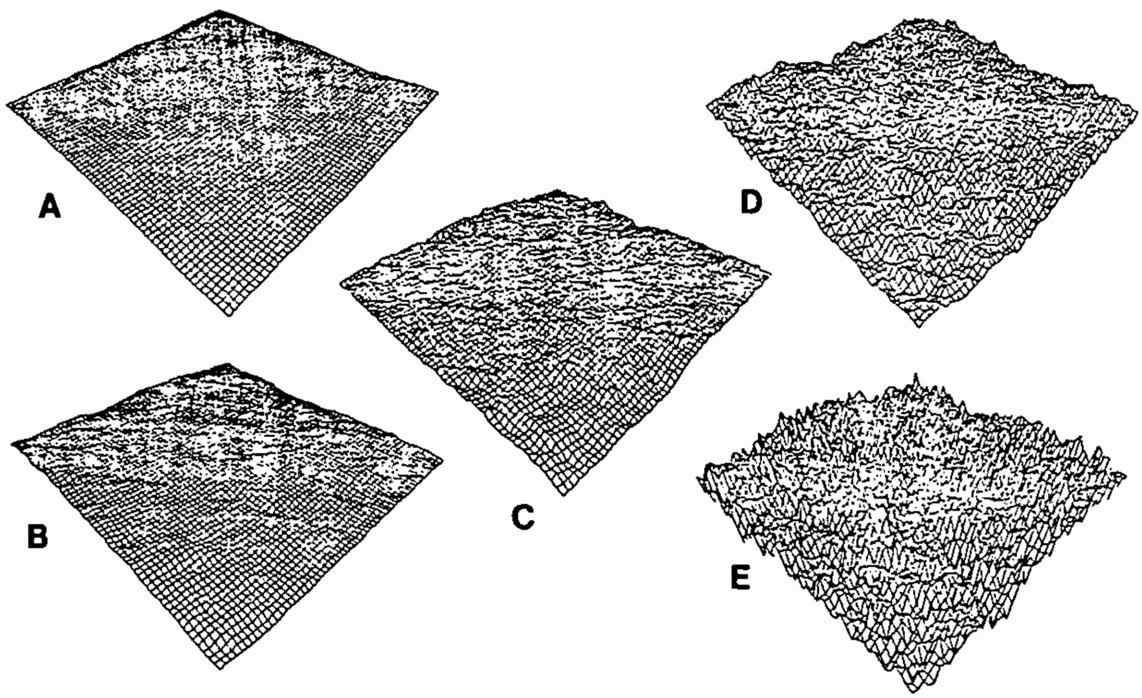

Figure 2. Surfaces of Increasing Fractal Dimension; the fractal dimension corresponds closely to our intuitive notion of roughness.

The defining characteristic of a fractal is that it has a *fractional dimension*, from which we get the word "fractal." Technically, a fractal is defined as a set for which the Hausdorff-Besicovich dimension is strictly larger than the topological dimension, i.e.. a set for which the only consistent description of its metric properties requires a "dimension" value larger than our standard, intuitive definition of the set's "dimension."

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 2: (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 ($D \approx 2.8$).

One general characterization of fractals is that they are the end result of physical processes that modify shape through local action. Such processes will, after innumerable repetitions, typically produce a fractal surface shape. Examples are erosion, aggregation (e.g., galaxy formation, meteorite accretion, and snowflake growth) and turbulent flow (e.g., of rivers or lava).

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 2), 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* ($p < 0.001$) *with the curve's fractal dimension — i.e., fractal dimension accounted for* 96% *of the variance in the roughness estimates. The fractal measure of perceptual roughness is therefore almost twice as accurate as any other reported to date, e.g., [29]*

An Illustration of Fractal Dimension. One familiar example of naturally occurring fractal curves is coastlines. When we examine a coastline (as in Figure 1), we see a familiar scalloped curve formed by innumerable bays and peninsulas. If we then examine a finer-scale map of the same region, we shall again see the same type of curve. It turns out that this characteristic scalloping is present at almost all scales of examination [8], i.e., the statistics of the curve are invariant with respect to transformations of scale.

To illustrate the importance of fractal dimension, let us suppose that we wish to measure the area of an island or the length of its coastline. Metric properties are, in general, estimated by taking a measuring instrument of size λ, determining that n such instruments will "cover" the curve or area to be measured, and applying the formula

$$M = n\lambda^D \quad ,$$

where M is the metric property to be measured (e.g., length, area), and D the topological dimension of the measuring instrument.

The fact that the coastline is scalloped at *all* scales causes a problem when we attempt to measure it because all of the curve's features that are smaller than the size of the measuring tool will be missed, whatever the size of the measuring tool selected. When we attempt to estimate the length of such a curve, therefore, the measurement we obtain depends not only on the coastline but also on the length of the measurement tool itself [8]!

Mandelbrot pointed out that, in order to obtain a consistent measurement of the coastline's length, we must generalize the notion of dimension to include *fractional* dimensions. The use of a fractional power in our mensuration formula compensates, in effect, for the length or area lost because of details smaller than λ. The unique fractional power that yields consistent estimates of a set's metric properties is called that set's *fractal dimension*. Because it provides the correct adjustment factor for all those details smaller than λ, it may also be viewed as a measurement of the shape's *roughness*.

One of the more important lessons such examples teach us is the following: standard notions of length and area do NOT produce consistent measurements for many natural shapes. The basic metric properties of these shapes vary as a function of their fractal dimension. Fractal dimension, therefore, is a *necessary* part of any consistent description of the metric properties of such shapes, for any description that lacks it will not be correct at more than one scale of examination.

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, 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 extension to two or more topological dimensions. If $I(x)$ is scalar, then the fractal dimension D of the graph described by $I(x)$ is

$$D = 2 - H \tag{2}$$

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 use[3] of Equation 1, or from $I(x)$'s Fourier power spectrum $P(f)$, as the spectral density of a fractal Brownian function is proportional[4] to f^{-2H-1}.

A. Fractals 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. Real images and surfaces can not, of course, be true mathematical fractals, because the latter are defined to exist at *all* scales. Physical surfaces, in contrast, have an overall size that places an upper limit on the range of applicable scales. A lower limit is set by the size of the surfaces' constituent particles. Fractals, in common with all mathematical abstractions, can only approximate physical objects over a range of physical parameters.

Because it is unreasonable to expect a physical surface to be fractal over *all* scales, the only physically reasonable definition of a "fractal surface" is a surface that may be accurately approximated by a single fractal function over a *range* of scales. We shall say, therefore, that a surface is *fractal* if the fractal dimension is stable over a wide range of scales, the implication being that it can be accurately approximated over that range of scales by a single fractal function.

These considerations prompt the following two definitions, the first applicable to a two-dimensional function such as the image intensity surface, the second applicable to a topologically two-dimensional surface embedded in three dimensions, such as the surface of a mountain.

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.

[3] See the beginning portions of Section III, and IVc.

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

In the next section I will present evidence showing that many 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. Further, 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 [9]. 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.

Proof. Under the Lambertian and constancy assumptions, the image intensity I at a point P is a function of the surface normal $\mathbf{N}$ at the surface point that projects to P:

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

where ρ is the albedo of the surface, λ is the illuminant intensity, and $\mathbf{L} = (l_x, l_y, l_z)$ is the illuminant direction. Variations in I, therefore, are dependent only upon variations in $\mathbf{N}$.

The proposition claims that the image intensity I will obey the rule

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

To show this, we let $\mathbf{N}_1$ be the normal at point (x, y) and $\mathbf{N}_2$ be the normal at point $(x + \Delta x, y)$. Then we expand Equation (3), yielding

$$Pr\left(\frac{\rho\lambda(\mathbf{N}_1 \cdot \mathbf{L}) - \rho\lambda(\mathbf{N}_2 \cdot \mathbf{L})}{\|\Delta x\|^H} < y\right) = F(y)$$

Expanding the dot products, we obtain

$$Pr\left(\frac{\rho\lambda(\mathbf{N}_{1x}l_x + \mathbf{N}_{1y}l_y + \mathbf{N}_{1z}l_z) - \rho\lambda(\mathbf{N}_{2x}l_x + \mathbf{N}_{2y}l_y + \mathbf{N}_{2z}l_z)}{\|\Delta x\|^H} < y\right) = F(y)$$

As $\mathbf{N}_x$, $\mathbf{N}_y$ and $\mathbf{N}_z$ are all fractal Brownian functions, by virtue of the surface being assumed a spatially isotropic fractal Brownian function and, as ρ, λ and $\mathbf{L}$ are constant, then $\rho\lambda\mathbf{N}_x l_x$, $\rho\lambda\mathbf{N}_y l_y$ and $\rho\lambda\mathbf{N}_z l_z$ are also fractal Brownian (see Proposition 2 in the following section); thus

$$I = \rho\lambda(\mathbf{N} \cdot \mathbf{L}) = \rho\lambda(\mathbf{N}_x L_x + \mathbf{N}_y L_y + \mathbf{N}_z l_z)$$

must also be. Note that this proof may be generalized to include all cases in which the reflectance function is an affine transformation of $\mathbf{N}$. The dimension of I is the same as that of the components of **N**, since multiplication does not affect fractal dimension (see Proposition 2) •.

This proposition 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[5] 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[6] , therefore, 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 — its fractal dimension — from the image data. That fractal dimension is required to obtain a scale-invariant description of a surface's metric properties is an indication of its usefulness. That fractal dimension has also been shown to correspond closely to our intuitive notion of roughness shows the fundamental importance of the measurement: we can now discover from the image data whether the 3-D surface is rough or smooth, isotropic or anisotropic. We can know, in effect, from what kind of cloth the surface was cut.

EXPERIMENTAL NOTE:*Fifteen naive subjects (mostly language researchers) were shown digitized images of eight natural textured surfaces drawn from Brodatz [15]. These are shown in Figure 8. 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. This procedure was then repeated for the vertical direction, yielding a total of sixteen roughness estimates for each subject. A scale of one (smoothest) to ten (roughest) was used to indicate 3-D roughness/smoothness. The fractal dimension of the 2-D image was then computed along the horizontal and vertical directions by the use of Equation (5), as described in the following section, and the viewed surface's 3-D fractal dimension was estimated by the use of Proposition 1. The mean of the subject's estimates of 3-D roughness had an excellent* 0.91 *correlation* ($p < 0.001$) *with roughnesses predicted by use of the image's 2-D fractal dimension and Proposition 1, i.e., the 3-D fractal dimension predicted by use of the measured 2-D image's fractal dimension accounted for* 83% *of the variance in the subject's estimates of 3-D roughness. This result, therefore, supports the general validity of Proposition 1.*

Properties of Fractal Brownian Functions. Fractal functions must be stable over common transformations if they are to be useful as a descriptive tool. The following propositions prove that the fractal dimension of a surface is invariant with respect to linear transformations of the data and to transformation 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.

[5]The surface normal is a function of the first derivative of depth; thus, we can construct an integration procedure that converts surface normals into surface shape (as a depth map).

[6]Rubin and Richards [28] describe a scheme whereby the homogeneity of a surface may be determined from its imaged color.

Proposition 2. A linear transformation of a fractal Brownian function is a fractal Brownian function with the same fractal dimension.

Proof. The proposition claims that if $I(x)$ is a fractal Brownian function, i.e., obeys Equation (1) then

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

will be true of $AI(x)+B$. This second expression may be rewritten as

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

or

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

thus proving the proposition; linear transforms merely scale the distribution $F(y)$ •.

Proposition 3. The fractal dimension of a fractal Brownian function is invariant over transformations of scale.

Proof. The proposition claims that, if $I(x)$ is a fractal Brownian function, i.e., it obeys Equation (1), then

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

will be true of $I(x)$. This is trivially true; we need only set $\Delta x^{*} = k\Delta x$, and then the second expression may be rewritten as

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

thus proving the proposition •.

B. Contours And The Imaging Process

We have described a method whereby the fractal dimension of the surface can be inferred for homogeneous, uniformly lit surfaces. Even if the surface is not homogeneous or uniformly illuminated, however, we can still hope to infer the fractal dimension of the surface from imaged surface contours and bounding contours.

Contour shape is often primarily a function of surface shape; this is especially true for contours that lie mostly within a plane intersecting the surface. Common examples of such approximately-planar contours are bounding contours and contours that are "drawn" on the surface, e.g., cast shadows. The imaged projection of such planar contours is simply a linear transform of the 3-D contour; recalling that linear transforms do not alter the fractal dimension of a function, we see that the fractal dimension of these imaged contours is the same as that of the 3-D contour.

Thus we may use the fractal dimension of imaged contours to directly infer that of the 3-D surface (the surfaces' dimension is simply one plus the contours' dimension). Consequently, the estimate of fractal dimension obtained from contours can be used to corroborate the one derived from image intensities.

III. Applicability Of The Fractal Model

Proposition 1 proved that a fractal surface implies that the image intensity surface is itself fractal. The reverse is also true, as is proved in the following proposition. This proposition, therefore, gives us a method of evaluating the usefulness of the fractal surface model for particular image data: to determine whether or not a 3-D surface is fractal, all we need to do is to determine whether its image is fractal (given that we have first determined that the surface is homogeneous, perhaps by use of color information [28]).

Proposition 4. If an image intensity surface is a two-dimensional fractal Brownian then the imaged 3-D surface must be spatially-isotropic fractal Brownian, given that the surface is Lambertian and the illumination and albedo are constant.

Proof. For a Lambertian surface, the image intensity I is a linear function of the components of the surface normal $\mathbf{N}_x$, $\mathbf{N}_y$, and $\mathbf{N}_z$, e.g.,

$$I = \rho\lambda(\mathbf{N} \cdot \mathbf{L}) = \rho\lambda(\mathbf{N}_x L_x + \mathbf{N}_y L_y + \mathbf{N}_z l_z) \tag{4}$$

To prove that the a fractal Brownian image intensity surface necessarily entails a 3-D surface that is spatially-isotropic fractal Brownian, it suffices (by definition) to prove that a fractal image implies that the components of the surface normal are fractal.

We first note that Proposition 2 proves that linear transforms do not affect the fractal nature of a function nor its dimension. Thus Equation (4) shows that the fractal nature (and dimension) of the image intensity surface is determined by the sum of $\mathbf{N}_x$, $\mathbf{N}_y$, and $\mathbf{N}_z$. As the finite sum of non-fractal functions is non-fractal, we see that if the image intensity surface is fractal Brownian then so must at least one of the components of the surface normal. If the image is two-dimensionally fractal (e.g., fractal in both the x and y image directions), then at least two independent components of the surface normal must be fractal. Thus, as the surface normal has only two degrees of freedom by virtue of being constrained to have unit magnitude, a two-dimensionally fractal Brownian image intensity surface implies that all of the surface normals' components must be fractal Brownian and the surface is therefore spatially-isotropic fractal Brownian. •

To evaluate the applicability of the fractal model for a particular surface and its image data, then, we need only verify the homogeneity of the surface and the fractal nature of the image intensity surface. It appears that verification of surface homogeneity can be done by use of color information [28]; in order to verify the fractalness of the image we first rewrite Equation (1) to obtain the following description of the manner in which the second-order statistics of the image change with scale:

$$E(|\Delta I_{\Delta x}|)\|\Delta x\|^{-H} = E(|\Delta I_{\Delta x=1}|) \tag{5}$$

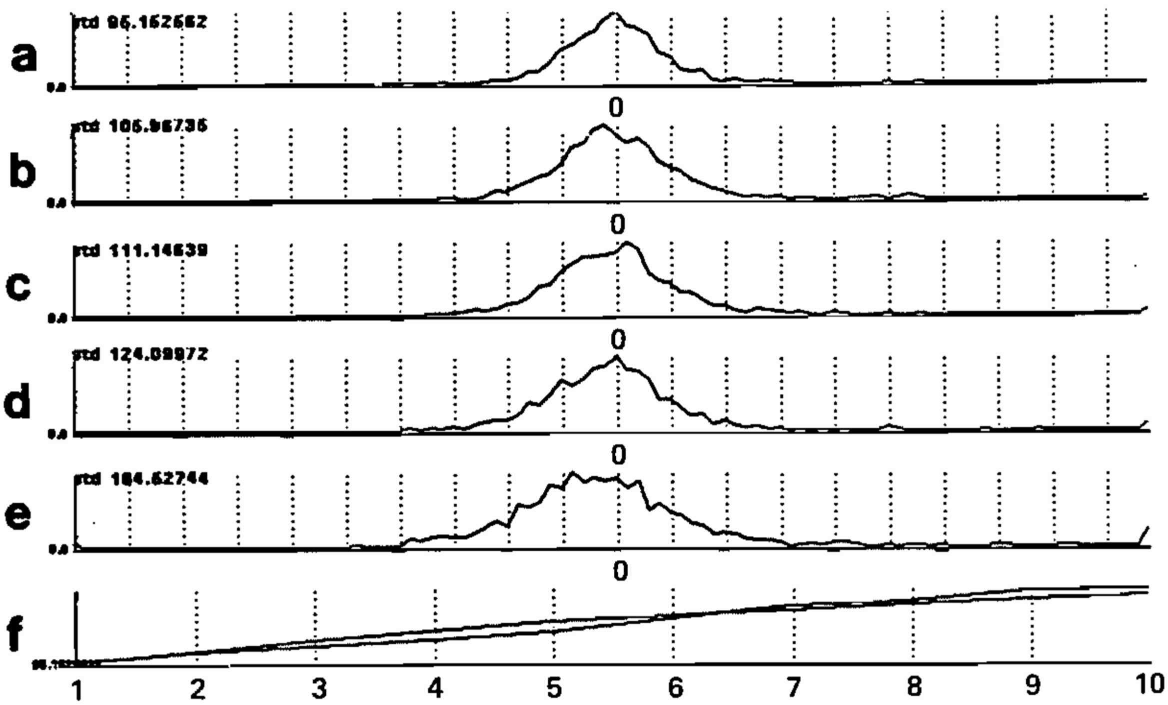

Figure 3. Results For a Typical Textured Patch.

where $E(|\Delta I_{\Delta x}|)$ is the expected value of the change in intensity over distance Δx. Equation (5) is a hypothesized relation among the image intensities; a hypothesis that we may test statistically. If we find that Equation (5) is true of the image intensity surface within a homogeneous image region[7] , then Proposition 4 tells us that the viewed surface must be a 3-D fractal Brownian surface, and thus the fractal model is appropriate. It is an important characteristic of the fractal model that we can determine its appropriateness for particular image data because it means that we can know when (and when *not*) to use the model.

To evaluate the suitability of the fractal model for natural surfaces, the homogeneous regions from each of six images of natural scenes were densely sampled. In addition, detailed images of twelve textured surfaces (see Brodatz [15]) were digitized and examined (see Figure 8). The intensity values within each of these regions were then approximated by a fractal Brownian function and the approximation error observed.

Figure 3 shows the results for a typical textured patch. The graphs (a) to (e) show the distribution of intensity differences (i.e., the second-order difference statistics) at one, two, three, five and ten pixel distances; the distributions are approximately Gaussian. Figure 3(f) shows a plot of the standard deviation of these distributions as a function of scale (i.e., $E(|\Delta I_{\Delta x}|)$ as a function of Δx in pixels). Overlaid on this graph is a least-squares fit of a

[7] I.e., we calculate the quantities $E(|\Delta I_{\Delta x}|)$ for various Δx, use a least-squares regression (using the log of Equation (5)) to estimate H, and examine the residuals.

fractal rule. As can be seen, the fit is quite good -- implying that the intensity surface in this region is actually a fractal Brownian function, at least over the 10 : 1 range of scales measured.

For the majority of the textures examined (77%), the fit was as good or better than the example shown. In 15% of the cases the region was constant except for random, zero-mean perturbations; consequently, the fractal function correctly approximates the image data, although the estimated fractal dimension is equal to the topological dimension[8] . The fit was poor in only 8% of the regions examined. In some of these cases it appeared that the image digitization had become saturated, and thus the poor fit may have been artifactual.

The fact that the vast majority of the regions examined were quite well approximated by a fractal Brownian function indicates that the fractal model will often provide a useful description of natural surfaces and their images. In those cases for which the fractal description is appropriate, the only statistical structure that remains unaccounted for by the fractal Brownian function is zero-mean unit-variance Gaussian noise --- indicating that the fractal description effectively exhausts all of the second-order difference information within the image.

Following initial report of this work [14] a similar investigation was conducted by Peleg *et al* [27]. In their work high-resolution images drawn from Brodatz [15] were examined over much larger ranges of scale, and their data show that the images' fractal dimension was not constant over *all* scales but rather only over *ranges* of scale. This data might be naively interpreted as indicating that these textures are not fractal, however such an interpretation is incorrect.

As observed earlier, physical processes do not typically act at all possible scales but rather only over a *range* of scales. Thus we should expect that a physical surface (and thus its image) will change its fractal characteristics when we pass from a range of scales dominated by one formative process to a range of scales that was shaped in a different manner. It is this realization that real surfaces will be fractal over ranges of scale, rather than fractal over all scales, that prompted the careful inclusion of *limited* ranges of scale in this papers' definition of 2-D and 3-D fractal surfaces. The ranges of constant fractal dimension observed in the Peleg *et al* data, therefore, are consistent with (and provide independent confirmation of) the fractal surface model.

A. The Relationship Between Fractals and Regular Patterns

Fractal Brownian functions do *not*, of course, describe regular or large-scale spatial structures such as are seen in the image of a brick wall or a tiled floor. Such structures must be accounted for by other means. It is important to realize, however, that while fractal Brownian surfaces are required to have particular second-order statistics, this does *not* mean that they cannot be regularly patterned.

To understand this, consider that the probability of a random number generator producing the string "1010..." is exactly the same as the probability of any other particular string with half 1's, half 0's. Both strings have the same statistics, and thus the same probability of occurrence, although one is regularly patterned and the other isn't. Similarly, a

[8]In these cases the data's dimensionality is technically not "fractional," but this distinction need not concern us here.

surface such as a brick wall can be a perfectly good Brownian fractal: the overall distribution of second-order statistics is correct; it simply contains position-dependent patterns.

The fact that fractal Brownian functions can exhibit regularities allows us to smoothly pass from random, chaotic surfaces to regular, patterned ones within the same conceptual framework [16]. Regular surfaces, for instance, can be generated by adding constraints (patterning) to the random-number generator used in conjunction with computer graphics techniques for recursively generating fractal Brownian functions [9].

B. Detection Of Edge Points

It is an important characteristic of the fractal model that we can determine its appropriateness for particular image data, because this allows us to know when, and when *not*, to use the model. If we discover an image region that does not fit the fractal model, Proposition 4 allows us to infer that we are not viewing a homogeneous fractal surface.

Boundaries between homogeneous regions are one example of a physical configuration that does not fit well into the fractal model. Thus, when we examine points that lie on the boundary between two image regions we find that the fit between the fractal model and the image data is normally poor. The fact that boundaries seem to be the most common event giving rise to a nonfractal intensity surface provides a method of detecting image points that are likely to be edges.

One simple way to find such points is examination of the computed fractal dimension. It turns out that when we compute the fractal dimension of a region covering a boundary between two homogeneous areas, by using the regions's Fourier power spectrum[9] , we normally calculate a fractal dimension that is less than the topological dimension. As this is a physical impossibility, the implication is that the assumptions of the fractal model are inappropriate for that specific image data. When we observe a measured fractal dimension that is less than the topological dimension, therefore, we can reasonably expect that we have found a texture edge. Examples of this will be shown in the following sections.

IV. INFERRING SURFACE PROPERTIES

Fractal functions appear to provide a good description of natural surface textures and their images; thus, it is natural to use the fractal model for image segmentation, texture classification, shape-from-texture and the estimation of 3-D roughness from image data. It is also natural to inquire into the relationship between the fractal surface model and the various other models of shape and texture that have previously been reported. This section, consequently, describes the research performed in these areas.

A. Examples of Image Segmentation

Proposition 1 tells us that, within a homogeneous region, the fractal dimension in

[9] 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.

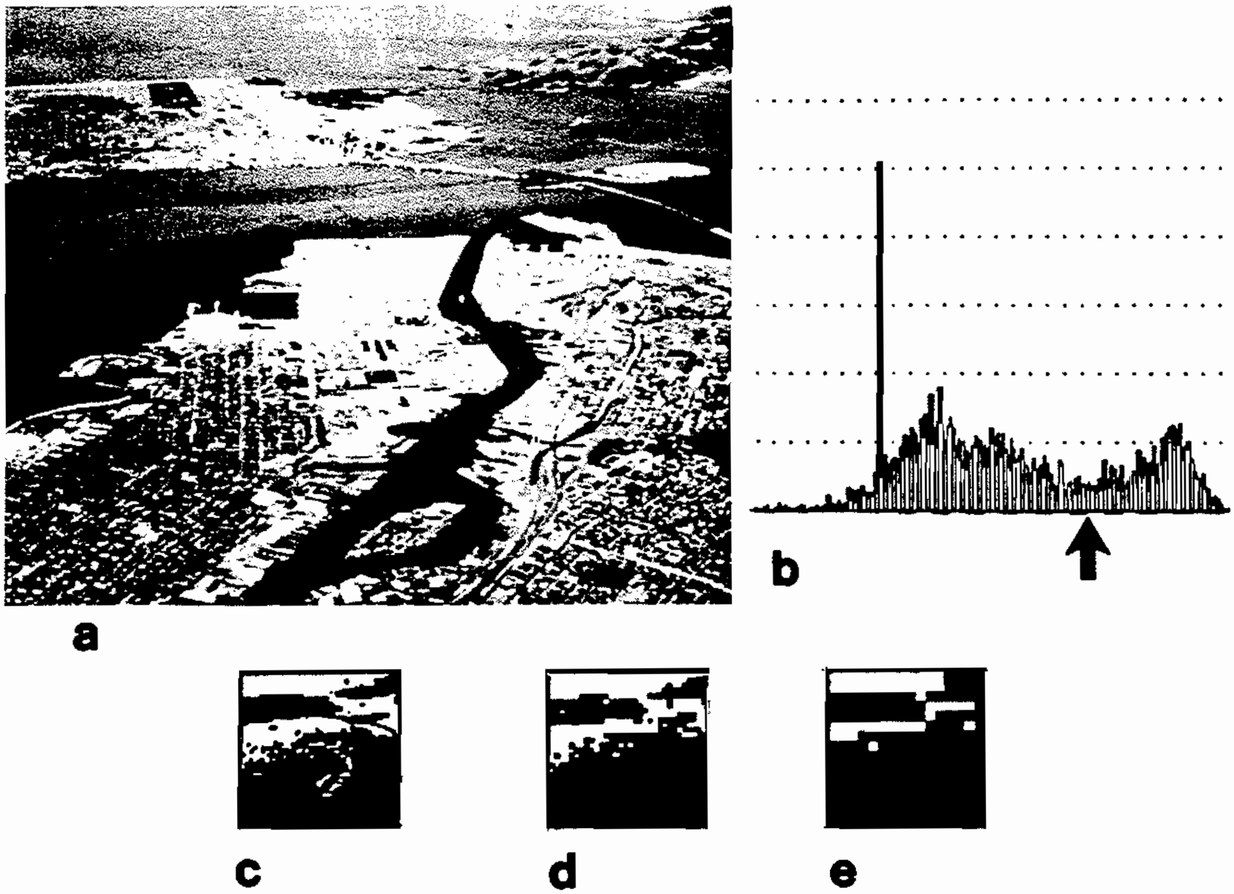

Figure 4. San Francisco Bay, and Its Segmentations.

the image is dependent upon that of the 3-D surface, thus giving us a technique for inferring a 3-D property of the viewed surface that closely corresponds to people's concept of roughness/smoothness. This suggests that measurement of the fractal dimension in the image will be useful in segmenting natural imagery.

Figure 4(a) shows an aerial view of San Francisco Bay. This image was digitized and the fractal dimension computed for each 8×8 block of pixels by means of the Fourier technique — i.e., the parameter H was estimated by a least-squares regression of the Fourier-domain fractal definition onto the power spectrum of the block of pixels. Orientational information was not incorporated into measurement of the local fractal dimension — i.e., differences in dimension among various image directions at a point were collapsed into one average measurement. Figure 4(b) shows a histogram of the fractal dimensions computed over the whole image[10] .

This histogram of fractal dimension was then broken at the "valleys" between the modes of the histogram, and the image segmented into pixel neighborhoods belonging to one mode or another. Figure 4(c) shows the segmentation obtained by thresholding at the breakpoint indicated by the arrow under (b); each pixel in (c) corresponds to an 8×8 block of pixels in the original image. As can be seen, a good segmentation into water and land

[10]The values to the left of the large spike in (b) have a computed fractal dimension that is less than the topological dimension; thus, these points are likely caused by patches that cross distinct regional boundaries; in fact, they all occur along the water-land boundary and delineate that boundary.

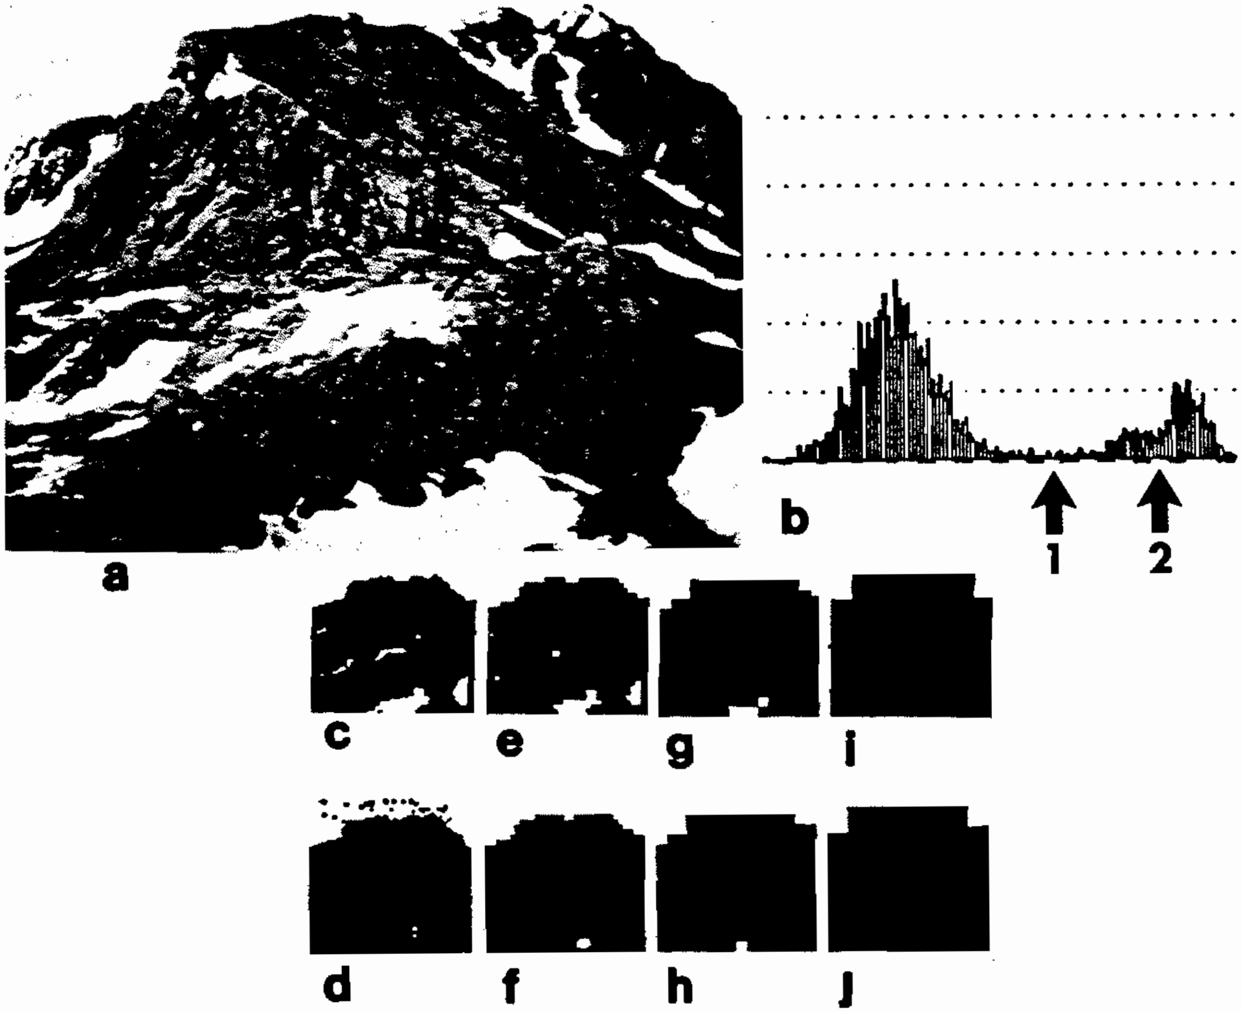

Figure 5. Mount Dawn, and Its Segmentations.

was achieved — one that cannot be obtained by thresholding on image intensity.

Proposition 3 indicates that this segmentation should be stable over transformations of scale. To test this prediction, the image was averaged down, from 512 $\times$ 512 pixels to 256 $\times$ 256 and 128 $\times$ 128 pixel images, and the fractal dimension recomputed for each of the reduced images. Figures 4(d) and (e) illustrate the segmentations produced by using the *same* breakpoint as had been employed in the original full-resolution segmentation. These results, therefore, demonstrate the stability of the fractal dimension measure across wide (4 : 1) variations in scale, as predicted by Proposition 3.

Figure 5(a) shows a view of Mount Dawn. This image was digitized into 512 $\times$ 512 pixels and the fractal dimension computed as before; (b) shows a histogram of the computed fractal dimension. Arrows at the bottom of (b) show where the distribution of fractal dimension was broken to produce a segmentation of the image. Figure 5(c) shows the image segmented into two classes (land; snow-and-sky) at the first histogram breakpoint. Figure 5(d) shows the sky separated from the land and snow by the second histogram breakpoint.

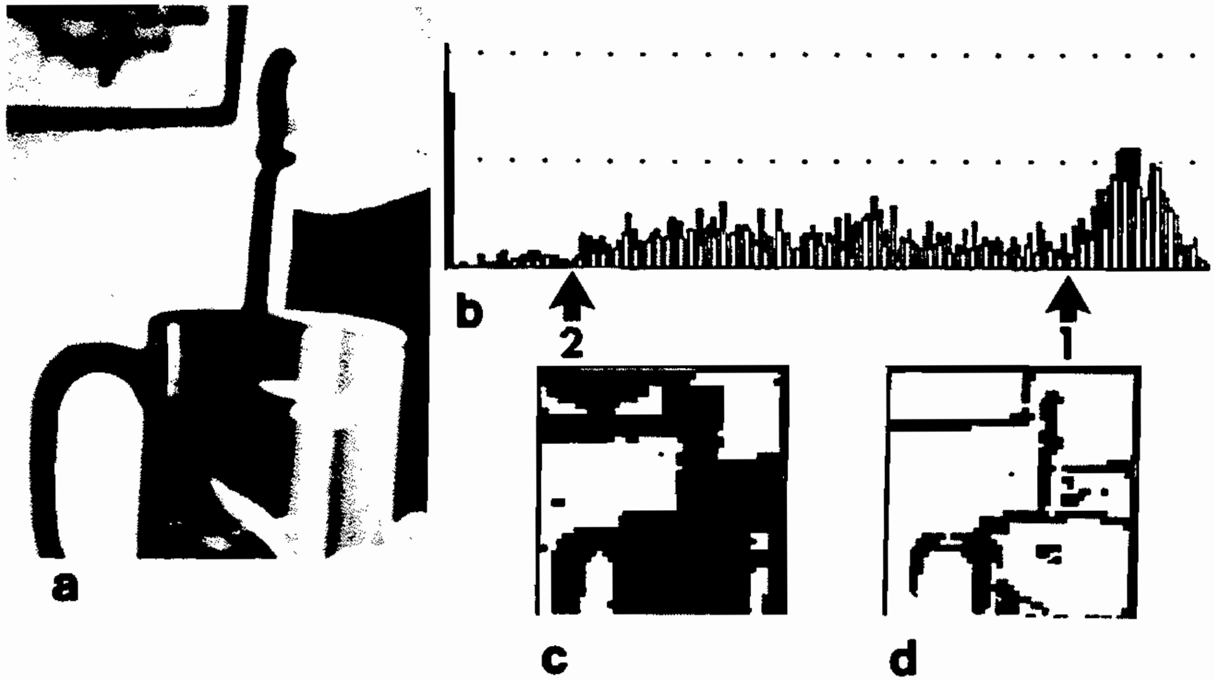

Figure 6. A Picture Of a Mug, and Its Segmentation.

Taken together, (c) and (d) demonstrate a good segmentation into mountain, snow, and sky. Note that the distinction between snow and sky is very subtle; it is impressive that this fine of a separation can be made by use of a simple image-wide histogram of roughnesses.

This image was also averaged down to 256 × 256, 128 × 128 and 64 × 64 pixel images, and the fractal dimension recomputed for each of the reduced images. Figures 5(e) — 6(j) illustrate the segmentations that result from using the *same* cut points as were employed in the original, full-resolution segmentation; it can be seen that the segmentations in these figures are quite similar, again demonstrating the stability of the fractal description across wide (8 : 1) variations in scale.

Images of smooth, man-made surfaces can also be usefully segmented, as shown in Figure 6. Figure 6(a) shows a picture of a mug and, just behind it, a chairback. This image was digitized into 256 × 256 pixels and the fractal dimension computed; (b) is a histogram of the computed fractal dimensions with the breakpoint indicated by an arrow. Figure 6(c) shows the image segmented into two classes at the point indicated by arrow No. 1. A good partial segmentation results[11] . Figure 6(d) illustrates the points whose computed fractal dimension is less than the topological dimension (the points to the left of arrow No. 2 in (b)); as expected, these are edge points.

One final example is the desert scene shown in Figure 7(a). This scene was segmented into three classes based on the histogram shown in (b); the segmentations are shown in (c) (road-and-sky versus desert) and (d) (road and desert versus sky). As can be seen, there is a good segmentation into desert, road, and sky.

Several other images have been segmented in this manner and, in each case, a good segmentation was achieved. The computed fractal dimension (and thus the segmentation)

[11]Nor can this segmentation be achieved by thresholding on intensity values.

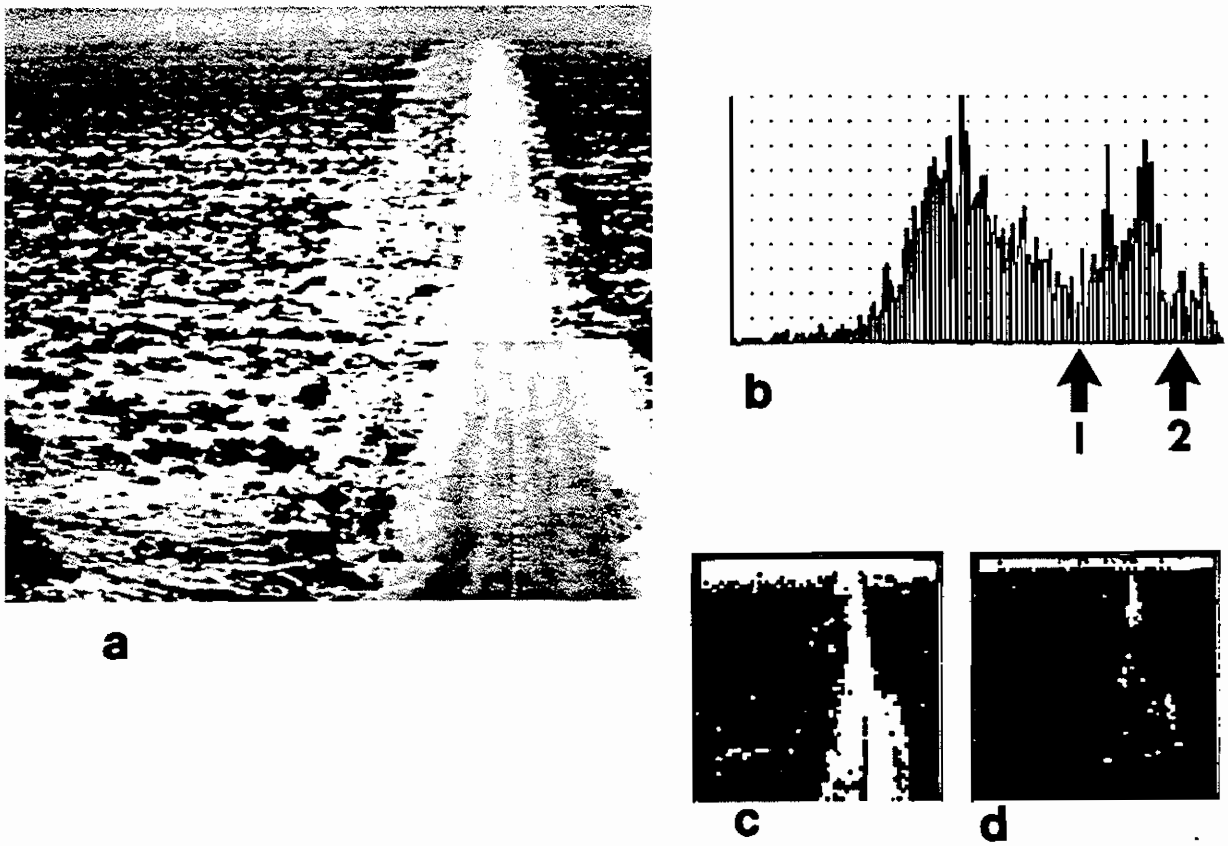

Figure 7. A Desert Scene, and Its Segmentation.

was always stable over at least 4 : 1 variations in scale; most segmentations were stable over a range of 8 : 1.

Stability of the fractal description is to be expected because the fractal dimension of the image is directly related to the fractal dimension of the viewed surface, which is a property of 3-D natural surfaces that is typically stable with respect to transformations of scale [8]. The fact that the fractal description is stable with respect to scale is a critically important property. After all, let us consider: *how can we hope to compute a stable, viewer-independent representation of the world if our information about the world is not stable with respect to scale?*

B. Comparison With Established Segmentation Techniques

To obtain an objective comparison with established segmentation techniques, a mosaic of eight images of natural surfaces [15] was digitized. The mosaic, shown in Figure 8, was constructed by Laws [17,18] for the purpose of comparing various texture segmentation procedures. The images that comprise this data set were chosen to be as visually similar as possible; gross statistical differences were removed by mean-value- and histogram-equalization.

Segmentation performance on these data exists for several techniques and, although differences in digitization complicate any comparisons we might wish to make, Laws's performance figures nevertheless serve as a useful yardstick for assessing performance on

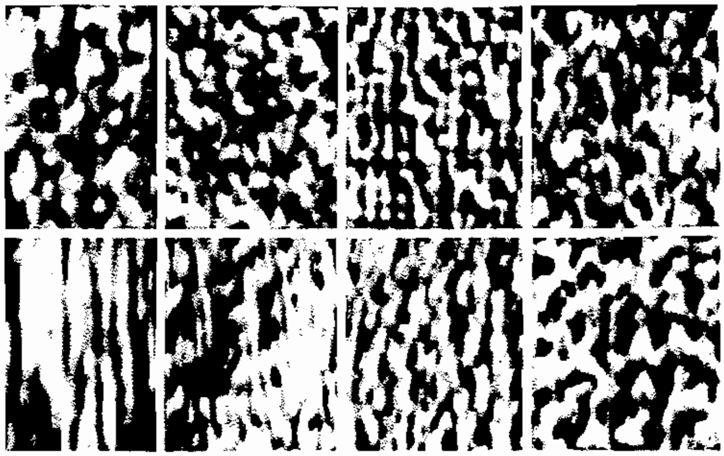

Figure 8. The Brodatz Textures Used For Comparison.

these data.

For this comparison, simple orientational information was incorporated into the fractal description; the fractal dimension was calculated separately along the x and y image directions. Fractal dimension was estimated, by using Equation (5), within five 16×16 pixel nonoverlapping subregions extracted from each of the eight regions. This data was next used to estimate the mean and variance of fractal dimension in the x and y directions, and theoretical classification probabilities were then computed.

The two-parameter fractal segmenter yielded a classification accuracy of 84.4%. This performance compares quite favorably with other segmentation techniques — despite the *much* larger number of texture features employed by these alternative methods. For example, Laws [17] reports accuracies of 65% for correlation statistics [19,20], 72% for co-occurrence statistics [21,22], and a theoretical accuracy of 87.4% for texture energy statistics[12] . The results of this comparison, therefore, indicate that fractal-based segmentation will likely prove a general and powerful technique[13] .

C. Shape Estimates

There are two ways surface shape is reflected in image patterning: (1) projection foreshortening, a function of the angle between the viewer and the surface normal, and (2) perspective gradients, which are due to increasing distance between the viewer and the surface. These two phenomena are independent in that they have separate causes. Thus, they can serve to confirm each other — i.e., if projection foreshortening is used to estimate surface tilt, that estimate is *independently confirmed* if there is a perspective gradieut of the proper magnitude and same direction [25,6]. We may be confident our estimate is correct when such independent confirmation is found.

[12]See [17], Page 148.

[13]On a data set of 12 Brodatz textures, in which the textures were somewhat less similar, the classification accuracy was 87.6%. Following initial report of this work [14], Peleg *et al.* [27] investigated other techniques for calculating fractal dimension and have reported essentially 100% accuracy on similar Brodatz textures, by using separate estimates of fractal dimension at many different scales. This allowed them to incorporate information about Δx_{min} and Δx_{max}, the limits at which the textures behave as a single fractal, and thus to improve their classification accuracy.

The fractal dimension found in the image, by virtue of its independence with respect to scale, appears to be nearly independent of the orientation of the surface. Fractal dimension, therefore, cannot be used to measure projection foreshortening. Projection foreshortening does, however, affect the variance of the distribution $F(y)$ associated with the fractal dimension (see Proposition 2) in the same manner in which it affects the distribution of tangent direction. Thus, to estimate surface orientation, we might assume that the surfaces' structure is isotropic and estimate surface orientation on the basis of previously derived shape-from-contour and shape-from-texture results [5,6].

This estimation technique often works; for instance, a vertical tilt and 45° slant was estimated for the image in Figure 4, and a vertical tilt and 47° slant was estimated for the upper portion of the image in Figure 7. The necessity of *assuming* isotropy, however, is a serious shortcoming of this technique — for, when the assumption is wrong, the estimate may be very much in error.

An important new result, therefore, is that we may partially cure this problem by observing the fractal dimensions in the x and y directions. If they are unequal we have *prima facie* evidence of anisotropy in the surface, because fractal dimension is largely unaffected by projection.

Regardless of how a foreshortening-derived estimate of surface orientation is produced, we may still seek *confirmation* of it by measuring the perspective gradient; if confirmation is found, we may be confident of our estimate. Such a gradient appears in Figure 4: the houses dwindle in size with increasing distance from the viewer.

Figures 9(a) and 9(b) show relief plots of the fractal dimension computed from Figures 4(a) and 7(a), respectively. In both 9(a) and (b) the x and y axes correspond to the horizontal and vertical directions in Figures 4(a) and 7(a); i.e., they are viewed as if from the left-hand side of the original images. The z axis shows the computed fractal dimension for each 8×8 block of pixels.

In both of these examples there is a gradual rise in the estimated fractal dimension with increasing distance. We can track this effect of perspective foreshortening and thus observe the perspective gradient.

It is at first somewhat puzzling to observe the fractal dimension changing with increasing distance, for fractal dimension is stable with respect to changes in scale. What we are in fact observing in these examples is interaction between our sampling rate and the range of scales over which the fractal approximation valid (see section IIa).

Real surfaces are not fractal at all scales; there are smallest and largest components to their shape (e.g., grain size and region size). These largest and smallest components define the limits between which the surface can be described with a single fractal function. When the projected size of a pixel becomes comparable to either of these limits, the fractal approximation can break down. If the pixel size large with respect to the largest shape components, we observe the familiar Nyquist sampling behavior: the surface appears to become smoother as the pixel size is increased. When the pixel size is less than the smallest shape components, we observe "texture" edges; i.e., inhomogeneities in the shape structure.

By observing the limits within which the fractal approximation holds[14] , we measure

[14] By use of Equation (5). Also see Peleg *et al.* [27], which describes an alternate fractal-dimension technique for discovery of these scale limits.

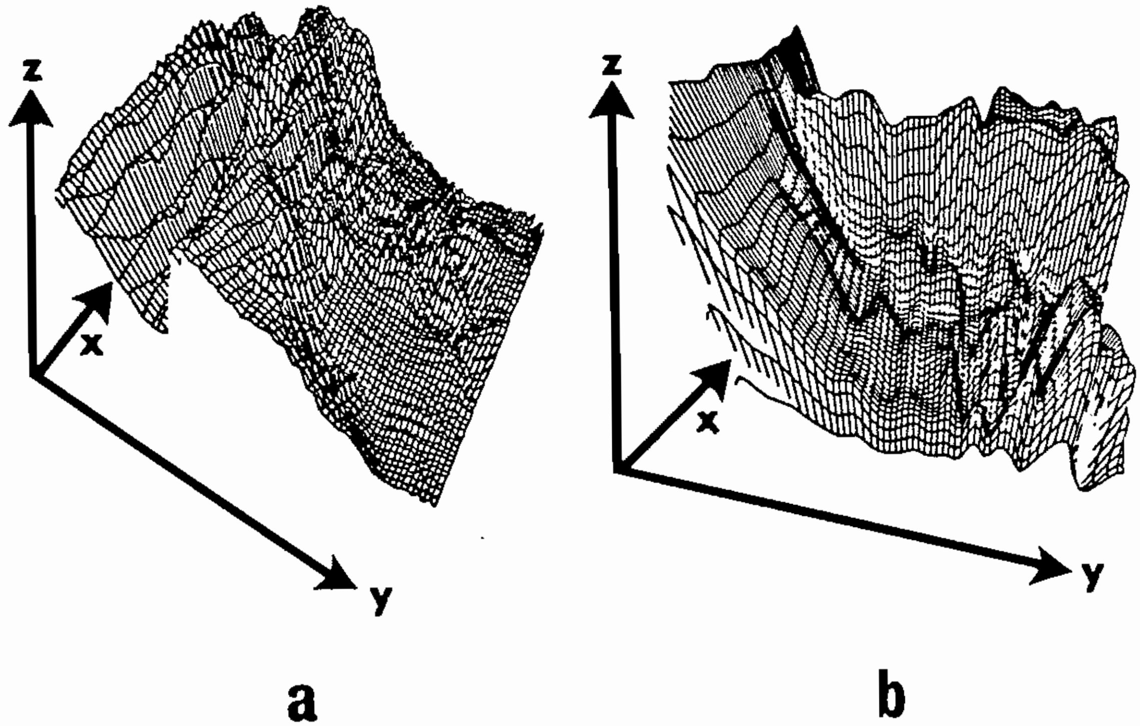

Figure 9. Relief Plots of The Computed Fractal Dimension For Figures 4(a) and 7(a).

a property intrinsic to the surface: the range of scales over which the surface obeys the fractal rule. Because such measurement is in terms of pixel size, it relates the size of the projected pixel relative to the (presumably invariant) size distribution of the surfaces' shape components. Measuring the largest scale at which a single fractal rule holds, therefore, gives us the ratio between pixel size and an aspect of the surfaces' intrinsic structure – and thus allows us to observe the perspective gradient.

In Figure 9 we are measuring the size of the surfaces' largest shape components by using the fact that the apparent smoothness of the viewed surface increases with increasing size of the projected pixel, given that the pixels' projected size is comparable to that of the surfaces' largest shape components. Measurements of fractal dimension may thus be used to measure the perspective gradient, thereby providing independent confirmation of the foreshortening-derived estimates of surface orientation. In imagery of large, planar regions it may also be possible to use the magnitude of the perspective gradient to estimate the surface's orientation.

Note, however, that in Figure 9(b) the portion of the image closest to the viewer does not exhibit a smooth gradient, because the size of the region used to compute the fractal dimension was small relative to the size of the rocks and bushes. In this near area the image data used to compute fractal dimension are often *boundaries* between homogeneous areas, and therefore do not fit the fractal model. As described previously, the appropriateness of the fractal model may be determined for the specific image data under consideration. In the case of the image in Figure 7(a), the fractal model is inappropriate for much of the data in the near portion of the image. Thus, most of the apparent "perspective gradients" in the near portion of Figure 9(b) can be identified as artifacts.

These two new results — the ability to obtain evidence of surface anisotropy and the measurement of the perspective gradient — represent significant advances in shape estimation, because they offer a way to substantially improve the reliability of shape-from-foreshortening [5,6] techniques.

D. Shading Into 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); e.g., they are stationary. Because these surfaces are judged by people[15] to be "smooth" the fractal model with small values of H is appropriate for modeling smooth, shaded regions of the image. In contrast, fractals with $H > 0$ are not perceived as smooth, but rather as rough or textured.

The fractal model can therefore encompass both image shading and texture, with shading as a limiting case in the spectrum of texture granularity[16] . The fractal model thus allows us to make a reasonable, rigorous and perceptually plausible definition of the categories "texture" and "shading" 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 such techniques as shape-from-shading [3,4,26], surface interpolation [4], and shape-from-texture [5,6] — to mention only the obvious cases. Thus, use of the fractal model to infer qualitative 3-D shape, i.e., smoothness/roughness, has the potential to significantly improve the utility of many other machine vision methods.

E. Relationship To 2-D Texture Models

One of the more interesting aspects of the fractal surface model is that it relates 2-D texture measures based on co-occurrence statistics [21,22], Fourier spectra [23,24,25], Markov processes [13], or autocorrelation [19,20] to each other and to 3-D surface structure.

We have seen that fractal Brownian functions may be defined in terms of either the way interpixel differences (second-order statistics) change with distance, or the rate at which the Fourier power spectrum falls off with increasing frequency. Similarly, fractal functions may be characterized by the way the autocorrelation function falls off [7,10], or by Markov processes [10,13]. Because the fractal image of a 3-D fractal surface may be described in any of these terms, it follows that for fractal images we may relate each of these texture measures to the other and to the 3-D fractal surface model. The fractal surface model, therefore, offers the potential of unifying and simplifying these various 2-D texture descriptions, as well as the possibility of interpreting them in terms of the 3-D structure of the world.

To say that the fractal model can be described in these other terms is not to say that the fractal model is *equivalent* to these other models: the fractal model is clearly a *3-D*

[15]The surface may, however, have significant local fluctuations: these are usually seen as "dust" or some other extraneous effect modifying the underlying smooth surface. It may also be that beyond some limiting value of the variance of $F(y)$ the surface is no longer perceived as smooth.

[16]If we assume that incident light is reflected at the angle of incidence and we make the variance of $F(y)$ small relative to the pixel size, the surface will be mirrorlike. If, on the other hand, the variance of $F(y)$ is large relative to the pixel size, the surface will become more isotropically reflecting. Thus, we can use the fractal model to capture the intuitive notion that reflectance functions are due to the structure of the microtexture.

model, whereas the texture models are only 2-D. Further, although it is true that a fractal function can be characterized in terms of Fourier spectra, co-occurrence matrices, etc., it is not true that *any* characterization of an image in these terms captures the 3-D properties of the viewed surface, as is the case with the fractal description. Characterization of an image in terms of radial slices of the Fourier domain (for instance) is completely orthogonal to the fractal description and, as a result, constrains the shape of the 3-D surface hardly at all[17]. As a consequence, we cannot expect the image segmentation performance of these texture techniques to be generally indicative of the performance of the fractal surface model.

One may still ask, since we can "translate" the fractal surface model into these various texture descriptions, why should one employ the fractal model rather than some other? The principal advantage of describing textures in terms of fractal surfaces, rather than in any of these other vocabularies, is that it allows us to capture a simple physical relationship that underlies the texture structure; a relationship that allows us to interpret the 2-D texture measurements in terms of the 3-D world. The fact that this physical interpretation can be lost with the 2-D characterizations of texture makes it seem advantageous to characterize texture problems in terms of the 3-D fractal surface model.

F. Relationship To Human Texture Perception

In light of the fact that the fractal surface model has been shown to predict peoples' perception of 3-D roughness (see section II), it is worth examining the relationship between the fractal surface model and models of human texture perception.

The most widely-known model of human texture perception is due to Julesz [30] who suggested that pre-attentive texture perception is dependant upon the global second-order statistics of the texture. Although this suggestion is now known to be wrong, Gagalowicz [31] has presented evidence that the problems with Julesz's conjecture are obviated by making texture discrimination dependent upon only the *local* second-order statistics. Others, such as Richards and Polit [32], have presented evidence that texture perception is mediated by spatial-frequency tuned channels.

It turns out that both the second-order statistics and the Fourier models fit well with the notion that people use the fractal dimension of the image (and thus of the 3-D surface) in pre-attentive discrimination of unpatterned textures. That the fractal surface model and these perceptual models complement each other is not so surprising, for we have already described how to measure the images' fractal dimension either by use of the local second-order statistics[8] or the Fourier power spectrum[10].

Let us look first at Gagalowicz' model of texture perception, and for the sake of argument let us assume that discrimination between unpatterned texture is based on the perceived roughness of the corresponding 3-D surface. In this case the fractal surface model agrees with Gagalowicz' claim that texture discrimination will be dependant the local second-order statistics — as these determine the images' roughness (i.e., its fractal dimension) and thus the roughness of the 3-D surface.

Similarly, because the images' fractal dimension can also be measured from its Fourier

[17]Illumination effects can account for most variation in such a description. In general a description in these other terms will constrain the 3-D interpretation only to the extent that the description allows recovery of the images fractal parameters.

power spectrum, the image and surface roughness (fractal dimension) can be determined by use of spatial frequency channels. Thus the suggestion that spatial frequency channels mediate human texture perception also agrees with the fractal surface model and the assumption that pre-attentive texture discrimination depends on perceived roughness.

V. SUMMARY

Fractal functions seem to provide a good model for describing the rough, crenulated and crumpled 3-D surfaces typical of natural scenes. The evidence in support of this assertion is the following:

(1) Many basic physical processes produce fractal surfaces.

(2) Fractal surfaces *look* like natural surfaces, and thus appear to capture all of the shape structure relevant to human perception.

(3) We have conducted a survey of natural imagery and found that a fractal model of imaged 3-D surfaces, when transformed by the image formation process, furnishes an accurate description of both textured and shaded regions in most natural imagery.

Fractal functions, therefore, are useful for describing the complex 3-D surfaces typical of natural objects. By transforming this 3-D surface model through the image formation process we can obtain a useful model of how such surfaces appear in the image data. One important aspect of this model is that it is easy to test its appropriateness for particular image data.

Characterization of image texture by means of a 3-D fractal surface model has shed considerable light on the physical basis for several of the 2-D texture techniques currently in use, and made it possible to describe image texture in a manner that is stable over transformations of scale and linear transforms of intensity. These properties of the fractal surface model allow it to serve as the basis for an accurate image segmentation procedure that is similarly stable over a wide range of scales.

Because fractal dimension is not affected by projection distortion, its measurement can significantly enhance our ability to estimate shape from foreshortening. Specifically, measurement of fractal dimension can provide (1) evidence of surface anisotropy, and (2) an estimate of the perspective gradient. Both capabilities are extremely important because they provide a way to obtain independent confirmation of the assumptions on which previously reported techniques are based.

One further important result is that measurement of the 2-D image fractal dimension enables estimation of the 3-D fractal dimension. Knowledge of the 3-D fractal dimension has been shown to be a nearly perfect predictor of people's perception of roughness. Thus, the 3-D fractal model allows us to determine which imaged regions are perceived as smooth, and which ones appear textured. This discrimination is of special importance to shape-from-shading, shape-from-texture, and surface interpolation methods as their performance relies on assumptions about the smoothness or roughness of the viewed surface.

The encouraging progress that has already been achieved in research on these problems augers well for the fractal-based approach. It appears that the 3-D fractal model of surface shape will constitute a significant aid in efforts to proceed from the image of a natural scene to its description.

Acknowledgments. I would like to thank Ken Laws for his expertise and advice, Marty Fischler and Andy Witkin for their feedback, and Andy Hanson for getting me interested in fractals in the first place.

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] W. E. L. Grimson, "Computing Shape Using A Theory Of Human Stereo Vision," Ph.D. Thesis, Dept. of Mathematics, M.I.T. (1980).

[5] 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).

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

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

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

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

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

[11] A. Norton, "Generation and Display of Geometric Fractals in 3-D," Computer Graphics, vol. 16, 3, pp. 61-67, 1982.

[12] Y. Kawaguchi, "A Morphological Study Of The Form Of Nature," Computer Graphics, vol. 16, 3, pp. 223-232, 1982.

[13] L. C. Carpenter, "Vol Libre," computer-generated movie, 1980.

[14] A. Pentland, "Fractal-Based Description Of Natural Scenes," Proceedings of IEEE Conference on Computer Vision and Pattern Recognition '83. July 1983, Arlington, Virginia.

see also

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

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

[16] T. Pavlidis, *Personal Communication*

[17] K. Laws, "Textured Image Segmentation," Report 940, USC Image Processing Institute, Los Angeles, California, 1980.

[18] D. H. Ballard and C. M. Brown, "Computer Vision," Prentice-Hall Inc., Englewood Cliffs, New Jersey, 1982.

[19] W. K. Pratt, O.D. Faugeras, and A. Gagalowicz, "Visual Discrimination of Stochastic Texture," IEEE Transactions on Systems, Man and Cybernetics, vol SMC-8, pp. 460-473, 1978.

[20] K. Deguchi and I. Morishita, "Texture Characterization and Texture-Based Image Partitioning Using Two-Dimensional Linear Estimation Techniques," IEEE Transactions on Computers, vol C-27, pp. 739-745, 1978.

[21] A. Rosenfeld and E. B. Troy, "Visual Texture Analysis," IEEE Conference on Feature Extraction and Analysis, pp. 115-124, Argonne, Illinois Oct 1970.

[22] R. M. Haralick, K. Shanmugam and J. Dinstein, "Textural Features for Image Classification," IEEE Transactions on Systems, Man and Cybernetics, vol. SMC-3. pp. 610-621, 1973.

[23] R. Bajacsy and L. Lieberman, "Computer Description of Real Outdoor Scenes," Proceedings of 2d International Joint Conference on Pattern Recognition, pp. 174-179, Copenhagen, Aug 1974.

[24] H. Maurer, "Texture Analysis With Fourier Series," Proceedings of the 9th International Symposium on Remote Sensing of the Environment, pp 1411-1420, Ann Arbor, Michigan, April 1974.

[25] R. Bajacsy and L. Lieberman, "Texture Gradient as a Depth Cue," Computer Graphics and Image Processing, vol. 5, 1, pp 52-67, 1976.

[26] A. P. Pentland, "Local Computation Of Shape," Proceedings of the National Conference on Artificial Intelligence, Pittsburgh, Pennsylvania (1982).

[27] S. Peleg *et al*, University of Maryland Computer Science TR-1306. July, 1983.

[28] J.M. Rubin and W.A. Richards, "Color Vision and Image Intensities: When Are Changes Material?," Bio. Cyber., 45, pp. 215-226, 1982.

[29] 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

[30] B. Julesz, "Texture and Visual Perception," Sci. Am., 212, pp. 38-48, 1965

[31] A. Gagalowicz, "A New Method for Texture Fields Synthesis: Some Applications to the Study of Human Vision," IEEE Trans. on Pat. Anal. and Mach. Intell., Vol 3, 5, pp. 520-533, Sept. 1981

[32] W. Richards and A. Polit, "Texture Matching," Kybernetik, 16, pp. 155-162, 1974.