FRACTAL-BASED DESCRIPTION

Alex P. Pentland
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 such a description from image data. In order 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 good a choice for modeling 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. This characterization of image regions has been shown to be stable over transformations of scale and linear transforms of intensity.

Much work has been accomplished that is relevant to computing 3-D information from the image data, and the computation of a 3-D fractal-based representation from actual image data has been demonstrated using an image of a mountain. This example shows the potential of a fractal-based representation for efficiently computing good 3-D representations of natural shapes, including such seemingly-difficult cases as mountains, clumps of leaves and clouds.

I INTRODUCTION

This paper addresses two related problems: (1) representing natural shapes such as mountains, trees and clouds, and (2) computing such a description from image data. The first step towards solving these problems, it appears, is to obtain a model of natural surface shapes. The task of finding such a model is extremely important to computer vision because we face problems that seem impossible to address with standard descriptive techniques. How, for instance, should we describe the shape of leaves on a tree? Or grass? Or clouds? When we attempt to describe such common, natural shapes using standard shape-primitive representations, the result is an unrealistically complicated model of something that, viewed introspectively, seems very simple.

Furthermore, how can we extract 3-D information from the image of a textured surface when we have no models that describe natural surfaces and how they evidence themselves in the image? The lack of such a 3-D model has generally restricted image texture descriptions to being *ad hoc* statistical measures of the image intensity surface. 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. The ability to relate image to surface can provide the necessary leverage for dealing appropriately with the problems of finding a good representation for natural surfaces and computing such a description from the image data.

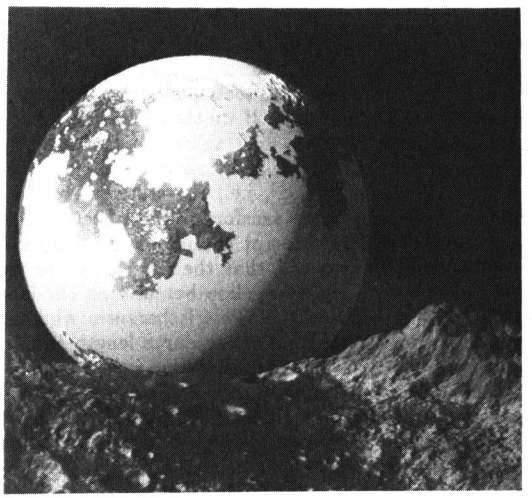

Figure 1. Fractal-based models of natural shapes, by Mandelbrot and Voss [4].

Even shape-from-shading and surface-interpolation methods are limited by the lack of a 3-D model of natural surfaces. Currently all such methods employ the heuristic of "smoothness" to relate neighboring points on the surface. Such heuristics are applicable to many man-made surfaces, of course, but are demonstrably untrue of most natural surfaces. In order to apply such techniques to natural surfaces, therefore, we must find a heuristic that is true of natural surfaces. Finding such a heuristic requires recourse to a 3-D model of natural surfaces.

Fractal functions seem to provide such a model of natural surface shapes. Fractals are a novel class of naturally-arising functions, discovered primarily by Benoit Mandelbrot. Mandelbrot and others [1,2,4] have shown that fractals are found widely in nature and that a number of basic physical processes, such as erosion and aggregation, produce fractal surfaces. Because fractals look *natural* to human beings, much recent computer graphics research has focused on using fractal processes to simulate natural shapes and textures (see Figure 1), including mountains, clouds, water, plants, trees, and primitive animals [3,4,5,6,7]. Additionally, we have recently conducted a survey of natural imagery and found that a fractal model of imaged 3-D surfaces furnishes an accurate description of both textured and shaded image regions, thus providing validation of this physics-derived model for both image texture and shading.

* The research reported herein was supported by the Defense Advanced Research Projects Agency under Contract No. MDA 903-83-C-0027; this contract is monitored by the U. S. Army Engineer Topographic Laboratory. Approved for public release, distribution unlimited.

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* [1,4]. Fractals are found widely in nature [1,2,4]. Mandelbrot shows that a number of basic physical processes, ranging from the aggregation of galaxies to the curdling of cheese, produce fractal surfaces. One general characterization is that any process that acts locally to produce a permanent change in shape will, after innumerable repetitions, result in a fractal surface. Examples are erosion, turbulent flow (e.g., of rivers or lava) and aggregation (e.g., galaxy formation, meteorite accretion, and snowflake growth). Fractals have also been widely and successfully used to generate realistic scenes (see Figure 1), including mountains, clouds, water, plants, trees, and primitive animals [3,4,5,6,7].

Perhaps the most familiar examples of naturally occurring fractal curves are 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 *all* scales of examination [2], i.e., the statistics of the curve are invariant with respect to transformations of scale. This fact causes problems when we attempt to measure the length of the coastline, because it turns out that the length we are measuring depends not only on the coastline but also on the length of the measurement tool itself [2]! This is because, whatever the size measuring tool selected, all of the curve length attributable to features smaller than the size of the measuring tool will be missed. Mandelbrot pointed out that, if we generalize the notion of dimension to include *fractional* dimensions (from which we get the word "fractal"), we can obtain a consistent measurement of the coastline's length.

The definition. A fractal is defined as a set for which the Hausdorff-Besicovich dimension is strictly larger than the topological dimension. Topological dimension corresponds to the standard, intuitive definition of "dimension." Hausdorff-Besicovich dimension D, also referred to as the *fractal dimension*, may be illustrated (and roughly defined) by the examples (1) of measuring the length of an island's coastline, and (2) measuring the area of the island.

To measure the length of the coastline we might select a measuring stick of length λ and determine that n such measuring sticks could be placed end to end along the coastline. The length of the coastline is then intuitively $n\lambda$. If we were measuring the area of the island, we could use a square of area λ^2 to derive an area of $m\lambda^2$, where m is the number of squares it takes to cover the island. If we actually did this, we would find that both of these measurements vary with λ, the length of the measuring instrument — an undesirable result.

In these two examples the length λ is raised to a particular power: the power of one to measure length, the power of two to measure area. These are two examples of the general rule of raising λ to a power that is the *dimension* of the object being measured. In the case of the island, raising λ to the topological dimension does not yield consistent results. If, however, we were to use the power 1.2 instead of 1.0 to measure the length, and 2.1 instead of 2.0 to measure the area, we would find that the measured length and area remained constant regardless of the size of the measuring instrument chosen.* The positive real number D that yields such a consistent measurement is the *fractal dimension*. D is always greater than or equal to the topological dimension.

*This example is discussed at greater length in Mandelbrot's book, "Fractals: Form, Chance and Dimension." The empirical data are from Richardson 1961.

The most important lesson the work of Mandelbrot and others teaches 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 the fractal dimension. Fractal dimension, therefore, is a *necessary* part of any consistent description of such shapes.

This result, which could almost be stated as a theorem, demonstrates the *fundamental* importance of knowing the fractal dimension of a surface. It implies that *any* description of a natural shape that does not include the fractal dimension cannot be relied upon to be correct at more than one scale of examination.

Fractal Brownian functions. Virtually all the 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. Motion of a particle undergoing Brownian motion is the canonical example of this type of fractal. The discussion that follows will be devoted exclusively to fractal Brownian functions, a generalization of Brownian motion.

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

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

where $F(y)$ is a cumulative distribution function [1]. The fractal dimension D of the graph described by $B(x)$ is

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

If $H = 1/2$ and $F(y)$ is a zero-mean Gaussian with unit variance, then $B(x)$ is the classical Brownian function. This definition has obvious extensions to two or more topological dimensions. The fractal dimension of a fractal Brownian function can also be measured from its Fourier power spectrum, as the spectral density of a fractal Brownian function is proportional to f^{-2H-1}. Discussion of the rather technical proof of this fact may be found in [1].

The fractal dimension of a surface corresponds roughly to our intuitive notion of jaggedness. Thus, if we were to generate a series of scenes with the same 3-D relief but increasing fractal dimension D, we would obtain the following sequence: first, a flat plane ($D \approx 2$), then rolling countryside ($D \approx 2.1$), a worn, old mountain range ($D \approx 2.3$), a young, rugged mountain range ($D \approx 2.5$), and finally a stalagmite-covered plane ($D \approx 2.8$).

The fractal dimension of a surface is invariant with respect to transformations of scale, as Δx is independent of H and $F(y)$. The fractal dimension is also invariant with respect to linear transformations of the data and thus it remains stable over smooth, monotonic transformations.

A. Fractals And The Imaging Process

Before we can use a fractal model of natural surfaces to help us understand images, however, we must determine how the imaging process maps a fractal surface shape into an image intensity surface. The mathematics of this problem is difficult and no complete solution has as yet been achieved. Nonetheless, simulation of the imaging process with a variety of fractal surface models can provide us with an empirical answer — i.e., that images of fractal surfaces are themselves fractal as long as the fractal-generating function is spatially isotropic. It is worth noting that practical fractal-generation techniques, such as those used in computer graphics, have had to constrain the fractal generating function to be isotropic so that realistic imagery could be obtained [3].

Real images do not, of course, appear fractal over all possible scales of examination. The overall size of the imaged surface places an upper limit on the range of scales for which the surface shape appears to be fractal, and a lower limit is set by the size of the surface's constituent particles. In between these limits, however, we may use Equation (1) to obtain a useful description of the surface.

Simulation shows that the fractal dimension of the physical surface dictates the fractal dimension of the image intensity surface; it appears that the fractal dimension of the image is a logarithmic function of the fractal dimension of the surface. If we assume that the surface is homogeneous, therefore, we can estimate the fractal dimension of the surface by measuring the fractal dimension of the image data. Even if the surface is not homogeneous, we can still infer the fractal dimension of the surface from imaged surface contours and bounding contours, by use of Mandelbrot's results.

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. The fact that the fractal dimension corresponds closely to our intuitive notion of roughness shows the 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, what kind of cloth the surface was cut from. The fact that the fractal dimension also describes the basic metric properties of the imaged surface is further indication that it is a critical element in any consistent representation of natural surfaces.

B. Applicability Of The Fractal Model

An implication of the fractal surface model is that the image intensity surface is itself fractal — and *vice versa*. This is because image intensity is primarily a function of the angle between the surface normal and the incident illumination; thus, if the image intensities satisfy Equation (1), then (for a homogeneous surface) the angle between surface normal and illuminant must also and, integrating, we find that the 3-D surface is a spatially isotropic fractal.

A method of evaluating the usefulness of the fractal surface model, therefore, is to determine whether or not images of natural surfaces are well described by a fractal function. To evaluate the applicability of the fractal model, 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:

$$E(|dI_{\Delta x}|)\|\Delta x\|^{-H} = E(|dI_1|) \tag{3}$$

where $E(|dI_{\Delta x}|)$ is the expected value of the absolute value of the change in intensity over distance Δx. Equation (3) is a hypothesized relation among the image intensities; a hypothesis that we may test statistically. If we find that Equation (3) is true of the image intensity surface and the viewed surface is homogeneous and continuous then we may conclude that the 3-D surface is itself fractal. 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 a fractal model for natural textures, the homogeneous regions from each of six images of natural scenes were densely sampled. In addition, twelve textures taken from Brodatz [8] were digitized and examined (see Figure 6). The intensity values within each of these regions were then approximated by a fractal Brownian function and the approximation error observed.

Figure 2 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.

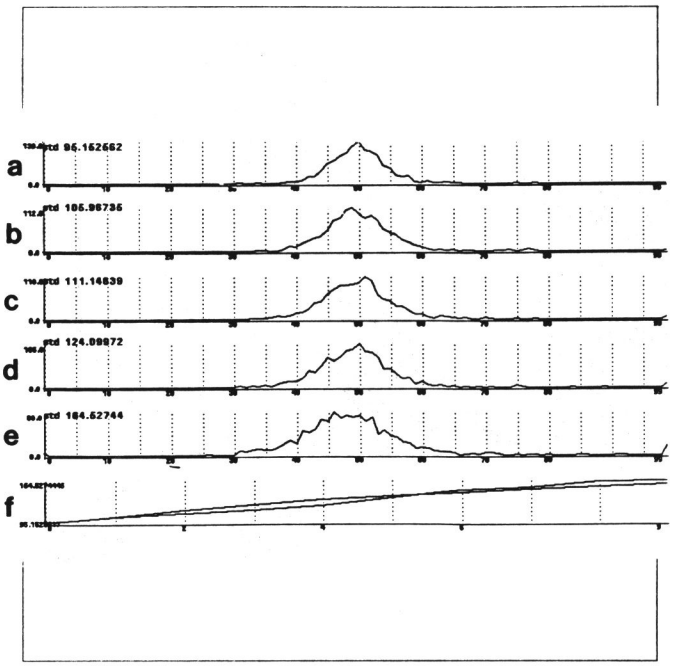

Figure 2. Results For a Typical Textured Patch.

Figure 2(f) shows a plot of the variance of the transition distribution as a function of distance between the pixels. Overlaid on this graph is a least-squares fit of a 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 model described the image data accurately. 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 fractal dimension was equal to the topological dimension and thus the data's dimension is technically not "fractional." The fit was poor in only 8% of the regions examined and, in many of these cases, it appeared that the image digitization had become saturated.

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 provide a useful description of natural textures and their images. Further, in the cases where 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 "uses up" all of the second-order difference information.

Fractal Brownian functions do *not*, of course, account for large-scale spatial structure such as seen in the image of a brick wall or of a tiled floor. Such structure must be accounted for by other means. Similarly, the fact that the fractal Brownian function appears to account for all of the second-order difference information does *not* mean that there are aren't other texture "dimensions" which contain useful texture information.

C. Texture Edges

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 to use the model. Conversely, when we discover an image region which can not be fit using the fractal model, we can infer that the image region was not caused by a homogeneous fractal surface.

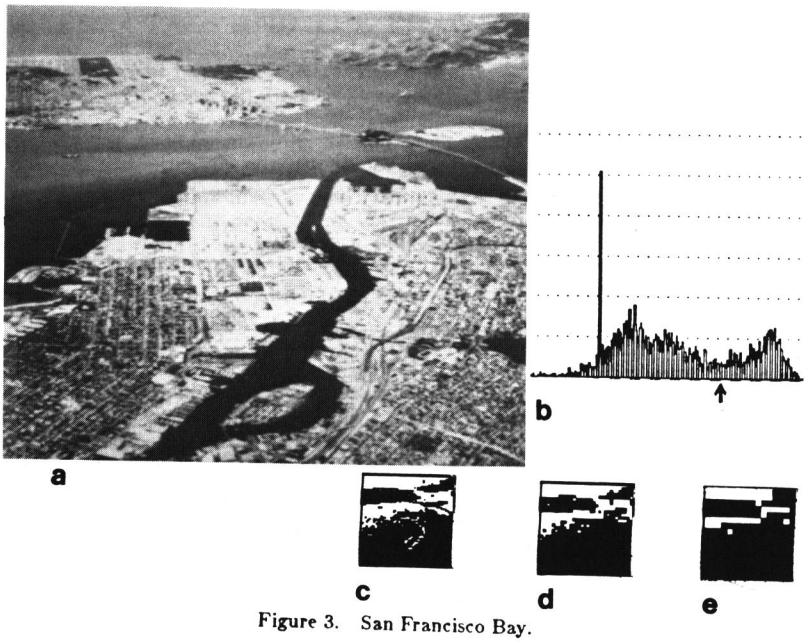

Figure 3. San Francisco Bay.

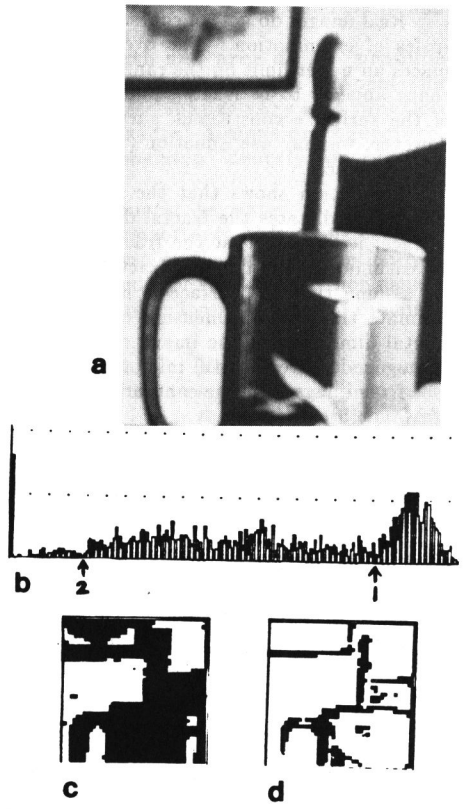

Figure 4. A Picture Of a Mug.

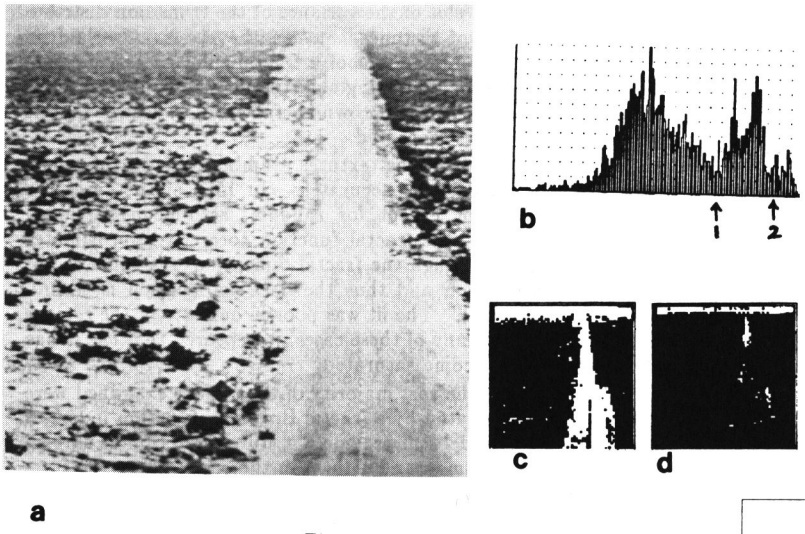

Figure 5. A Desert Scene.

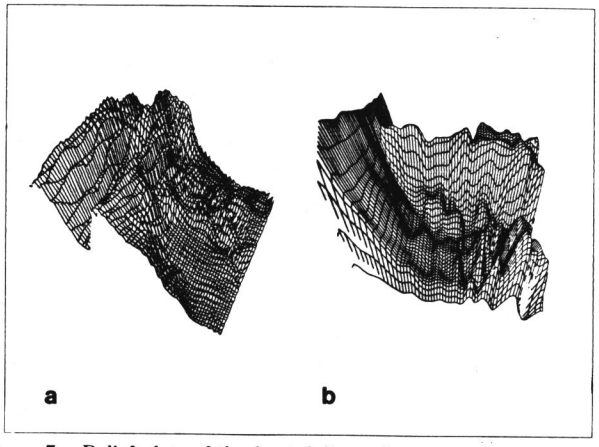

Figure 7. Relief plots of the fractal dimension computed for Figures 3 and 5.

Boundaries between homogeneous regions are one example of a physical configuration which can not be accounted for by the fractal model. Thus, when we examine points which 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 non-fractal intensity surface provides a method of detecting image points which are likely to be texture 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* for image data covering a boundary between two homogeneous areas, we normally calculate a fractal dimension which is less than the topological dimension. This is a physical impossibility, implying that the assumptions of the fractal model are inappropriate for that specific image data. Thus, a fractal dimension which appears to be less than the topological dimension is a sufficient condition for concluding that the fractal model is inappropriate; it may also be a necessary condition.

III 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 texture segmentation, classification and shape-from-texture. The first four headings of this section describe the research that has been performed in this area, and indicate likely directions for further research.

Fractal functions with $H \approx 0$ can be used to model smooth surfaces and their reflectance properties. For the first time, therefore, we can offer a single model encompassing both image shading and texture, with shading as a limiting case in the spectrum of texture granularity. The fractal model thus allows us to make a reasonable and rigorous definition of the categories "texture" and "shading," thus enabling us to discover similarities and differences between them. The final heading of this section briefly discusses this result.

A. Some Examples Of Texture Segmentation

In the following examples, the fractal dimension was measured from the Fourier power spectrum by using an 8×8 block of pixels, after which the measured fractal dimension was histogramed. 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. No attempt was made to incorporate orientational information into the measurement of the local fractal dimension, i.e., differences in dimension between various image directions at a point were collapsed into one average measurement.

Figure 3(a) is an aerial view of San Francisco Bay. Figure 3(b) shows a histogram of the fractal dimensions computed over the whole image. Figure 3(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 was achieved — one that cannot be obtained by thresholding on image intensity. The values to the left of the large spike in (b) have a computed fractal dimension which is less than the topological dimension; these points all occur along the water-land boundary and partially delineate that boundary.

This image was then averaged down, from 512×512 pixels into 256×256 and 128×128 pixel images, and the fractal dimension recomputed for each of the reduced images. Figures 3(d) and (e) illustrate the segmentations that result from using the *same* cut point as employed in the original full-resolution segmentation, demonstrating stability across wide (4 : 1) variations in scale.

*Using the Fourier technique.

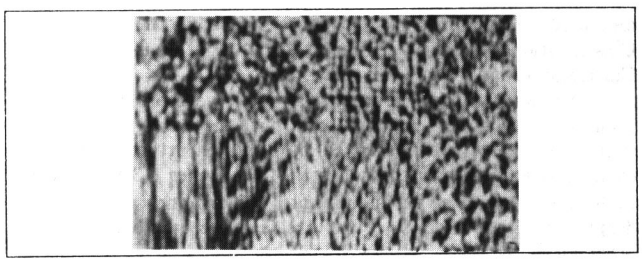

Figure 6. The Brodatz textures used for comparison.

Images of man-made surfaces can also be usefully segmented, as shown in Figure 4. Figure 4(a) shows a picture of a mug and, just behind it, a chairback. This image was digitized into 256×256 pixels and the local fractal dimension computed; (b) is a histogram of the fractal dimensions with the two breakpoints indicated by arrows. Figure 4(c) shows the image segmented into two classes at the point indicated by arrow number one. A good partial segmentation results. Figure 4(d) illustrates the points whose computed fractal dimension is less than the topological dimension (the points to the left of arrow number two in (b)); these are clearly edge points.

One final example is the desert scene shown in Figure 5(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. In each case a good segmentation was achieved. The computed fractal dimension, and thus the segmentation, was found to be stable over at least 4 : 1 variations in scale; most 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 natural surfaces that has been shown to be invariant with respect to transformations of scale [2].

The fact that the fractal description of texture is stable with respect to scale is a critically important property. After all, 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?** This example of texture property measurement reiterates what we observed earlier, i.e., the fact that the fractal dimension of the surface is *necessary* to any consistent description of a natural surface.

B. A Comparison With Other Segmentation Techniques

To obtain an objective comparison with previously established texture segmentation techniques, a mosaic of eight natural textures taken from Brodatz [8] was redigitized. The digitized texture mosaic, shown in Figure 6, was constructed by Laws [10] for the purpose of comparing various texture segmentation procedures. The textures 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 for 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 this data.

For this comparison simple orientational information was incorporated into the fractal description; the fractal dimension was calculated separately for the x and y directions. The sample size used to compute the fractal dimension was 16×16 pixels, and the dimension was computed directly, using Equation (1). Five subregions were extracted from each of the eight textures

and used to establish the parameters of a quadratic classifier. Classification probabilities were then computed using a Monte Carlo simulation.

The two-parameter fractal segmenter yielded a classification accuracy of 84.4%. This compares quite favorably with correlation techniques [11,12] reported by Laws as attaining 65% accuracy, as well as with co-occurrence techniques [13,14] reported to be 72% accurate.* This superior performance was achieved despite the large number of texture features employed by these other methods.

The simple two-parameter fractal segmenter even compares well with Laws' own texture energy statistics; even though Laws' segmentation procedure included more than a dozen texture statistics that were optimized for the test data, theoretical segmentation accuracy was only 3% better.** Thus, the results of these comparisons indicate that fractal-based texture segmentation will likely prove to be a general and powerful technique.

C. Relationship To Texture Models

The fact that the fractal dimension of the image data can be measured by using either co-occurrence statistics in conjunction with Equation (1), or by means of the Fourier power spectrum, suggests one interesting aspect of the fractal model: it highlights a formal link between co-occurrence texture measures [13,14] and Fourier techniques [15,16,17]. The mathematical results Mandelbrot derives for fractal Brownian functions show that the way interpixel differences change with distance determines the rate at which the Fourier power spectrum falls off as frequency is increased, and vice versa.

Thus, it appears that the fractal model offers potential for unifying and simplifying the co-occurrence and Fourier texture descriptions. If we believe that natural surface textures and their images are fractal (as seems to be indicated by the previous results), then the fractal dimension is the most relevant parameter in differentiating among textures. In this case we would expect both the Fourier and co-occurrence techniques to provide reasonable texture segmentations, as both yield sufficient information to determine the fractal dimension. The advantage of the fractal model would be that it captures a simple physical relationship underlying the texture structure — a relationship lost with either of the other two characterizations of texture.

D. Shape From Texture

There are two ways surface shape is reflected in image texture: (1) projection foreshortening, a function of the angle between the viewer and the surface normal, and (2) the perspective texture gradient that is due to increasing distance between the viewer and the surface. These two phenomena can be independently observed, therefore 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 texture gradient of the proper magnitude and same direction [17,18]. We may be confident our estimate is correct when such independent confirmation is found.

The fractal dimension found in the image appears to be nearly independent of the orientation of the surface (by virtue of independence with respect to scale); therefore fractal dimension cannot be used to measure surface orientation. Projection foreshortening does, however, affect the variance of the distribution $F(y)$ associated with the fractal dimension (see Equation (1)). Foreshortening affects $Var(F(y))$ in exactly the manner it affects the distribution of tangent direction. Thus, to estimate surface orientation, we might assume that the surface texture is isotropic and estimate surface orientation on the basis of previously derived results [18].

If we apply this technique to the image in Figure 3 we estimate a vertical tilt direction, and a slant of 45°; for the upper half of the image in Figure 5 we obtain a vertical tilt estimate and a slant of 47°. While this technique for estimating surface orientation often performs well, the necessity of assuming isotropy is a serious shortcoming of the technique. An important new result, therefore, is that we may in part 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 texture, because fractal dimension is unaffected by projection. It may even be possible to correct the surface orientation estimate using this information.

However a foreshortening-derived estimate of surface orientation is produced, we may still seek *confirmation* of it by measuring the perspective texture gradient; if confirmation is found, we may be confident of our estimate. The images in both Figures 3 and 5 also show a perspective texture gradient: the houses in Figure 3 and the bushes in Figure 5 dwindle in size with increasing distance from the viewer. If this texture gradient can be measured in these images, then the foreshortening-derived estimates of surface orientation can be confirmed.

Figures 7(a) and (b) show relief plots of the local fractal dimension computed from Figures 3(a) and 5(a), respectively. The computed dimensions are viewed from the left-hand side of the original image. In both (a) and (b) there is a gradual rise in the computed fractal dimension with increasing distance, a result of the pixel size becoming large with respect to the largest-scale variations in surface texture. This gradient in the computed fractal dimension confirms the foreshortening-derived estimates of surface tilt in this imagery.* These two examples demonstrate that perspective texture gradients can be inferred from the locally computed fractal dimension.

This two new results, i.e., the ability to obtain evidence of surface texture anisotropy and the measurement of the perspective texture gradient, are extremely important because they offer a way to make shape-from-unfamiliar-texture techniques sufficiently reliable so as to be useful.

E. Shading And Texture

Fractal functions with $H \approx 0$ can be used to model smooth surfaces and their reflectance properties accurately. When $H \approx 0$, the surface is locally planar, except for small, random variations described by the function $F(y)$ in Equation (1). 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 Lambertian.

The fractal model, therefore, is a single model that can account for both image shading and texture, with shading corresponding to the limiting value of H. The fractal model thus allows us to make a reasonable and rigorous definition of the categories "texture" and "shading," in terms that can be measured

*These comparisons are between the actual classification accuracies of the co-occurrence and correlation techniques and the theoretical accuracy of the fractal segmenter; however, Laws results indicate that the theoretical accuracy closely reflects the accuracies actually obtained.

**Laws' classification probabilities were achieved using linear classifier; See [10], page 148.

*Note, however, that in Figure 7(b) the portion of the image closest to the viewer does not exhibit a smooth texture gradient. In this near area the image data used to compute fractal dimension are often *boundaries* between textured areas, and thus 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. Thus, most of the "texture gradients" in the near portion of Figure 7(b) can be identified as artifacts.

by using the image data. One important goal of future research will be to discover similarities or differences between these two categories; initial results indicate that local shape-from-shading results [9] can be generalized to include shape-from-texture.

IV COMPUTING A DESCRIPTION

Current methods for representing the three-dimensional world suffer from a certain awkwardness and inflexibility that makes them difficult to envisage as the basis for human-performance-level capabilities. They have encountered problems in dealing with partial knowledge or uncertain information, and they become implausibly complex when confronted with the problem of representing a crumpled newspaper, a clump of leaves or a puffy cloud. Furthermore, they seem ill-suited to solving the problem of representing a *class* of objects, or determining that a particular object is a member of that class.

What is wrong with conventional shape representations? One major problem is that they make too much information explicit. Experiments in human perception [20] lead one to believe that our representation of a crumpled newspaper (for instance) is not accurate enough to recover every z value; rather, it seems that we remember the general "crumpledness" and a few of the major features, such as the general outline. The rest of the newspaper's detailed structure is ignored; it is unimportant, *random*.

From the point of view of constructing a representation, the only important constraints on shape are the crumpledness and general outline. What we would like to do is somehow capture the notion of **constrained chance**, that is, the intuition that "a crumpled newspaper has x, y and z structural regularities and the rest is just variable detail," thus allowing us to avoid dealing with inconsequential (random) variations and to reason instead only about the structural regularities.

A. The Process Of Computing A Description

How shall we go about computing such a "constrained chance" description?* Let us consider the problem formally and see where that leads us. The process of computing a shape description (given some sensory data) seems best characterized as attempting to confirm or deny such hypotheses as "shape x is consistent with these sense data." Computation of a shape description, therefore, seems to be a problem in induction [19].

If, naively, we try to use an inductive method, we start with the set of all possible shape hypotheses; we then attempt to winnow the set down to a small number of hypotheses that are confirmed by the sensory data. The "set of all shape hypotheses," however, is much too large to work with. Consequently, we must take a slightly different tack.

Using the notion of constrained chance. Rather than attempting to enumerate "all shape hypotheses" explicitly, let us instead construct a *shape generator* that uses a random number generator to produce a surface shape description (I shall shortly describe how to do this). If we were to run this shape generator for an infinite period, it would eventually produce instances of every shape within a large class of shapes. If the generator were so constructed that the class of shapes produced was exactly the set of "all hypotheses" about shape, then the program for the shape generator, together with a the program for the random number generator, would comprise a description of the set of all shape hypotheses.

The shape generator illustrates how the notion of constrained chance may be used to obtain a compact description of an infinite set of shapes. By changing the constraints that determine how the output of the random number generator is translated into shape, we can change the set of shapes described; specifically, we can introduce constraints that rule out some classes of shape and thus restrict the set of shapes that are described. The ability to progressively restrict the set of shapes described allows us to use the constrained-chance shape generator as the basis for induction, rather than being forced to use the explicitly enumerated set of all shape hypotheses.

The process of computing a "constrained chance description" is straightforward. We use image data to infer (using knowledge of the physics of image formation) constraints on the shape, and then introduce those constraints into the shape generator. The end result will be a programlike description that is capable of producing all the shapes that are consistent with the image data; i.e., we shall have a description of the shapes confirmed by the image data. This, then, is the type of description we wanted: a description of shape that contains the important structural regularities that can be inferred from the image (e.g., crumpledness, outline), but one that leaves everything else as variable, random.

Some people are already doing this. Something very much like this constrained-chance representation is already being widely utilized in the computer graphics community. Natural-looking shapes are produced by a simple fractal program that recursively subdivides the region to be filled, introducing random jaggedness of appropriate magnitude at each step [3,5]. The jaggedness is determined by specifying the fractal dimension. The shapes that can be produced in this manner range from planar surfaces to mountainlike shapes, depending on the fractal dimension. Current graphics technology often employs fractal shape generators in a more constrained mode; often the overall, general shape or the boundary conditions are specified beforehand. Thus, a scene is often constructed by first specifying initial constraints on the general shape, and then using a fractal shape generator to fill in the surface with appropriately jagged (or smooth) details. The description employed in such graphics systems, therefore, is exactly a constrained-chance description: important details are specified, and everything else is left unspecified except in a qualitative manner.

This type of description bears a close relationship to surface interpolation methods. Typically, such schemes fit a smooth surface that satisfies whatever boundary conditions are available. The initial boundary conditions, together with the interpolation function, constitute a precise description of the surface shape. Such schemes are limited to smooth surfaces, however, and therefore are incapable of dealing with most natural shapes. In contrast, a fractal-based representation allows either rough or smooth surfaces to be fit to the initial boundary conditions, depending upon the fractal dimension. This method of description, therefore, is quite capable of describing most natural surfaces — and that is why the graphics community is turning to the use of fractal-based descriptions for natural surfaces.

In order to make use of this type of description it is necessary to be able to specify the surface shape in a *qualitative* manner, i.e., how rugged is the topography? This specification of qualitative shape can be accomplished by fixing the fractal dimension. The fact that we have recently developed a method of inferring the fractal dimension of the 3-D surface directly from the image data means that we are now able, for the first time, to actually compute a fractal or constrained-chance description of a real scene from its image.

Not only terrestrial topography has been modeled by use of a constrained-chance representation, but also clouds, ponds, riverbeds, snowflakes, ocean surf and stars, just to name a few examples [1,3,4,5,6,7]. Researchers have also used constrained-chance generators to produce plant shapes [1,4,6]. A very natural-

*The term "representation" will be used to refer to the scheme for representing shapes, while the term "description" will be reserved for specific instances. Thus, one can compute a description of some object; it will be a member of the class of shapes that can be accounted for within the representation.

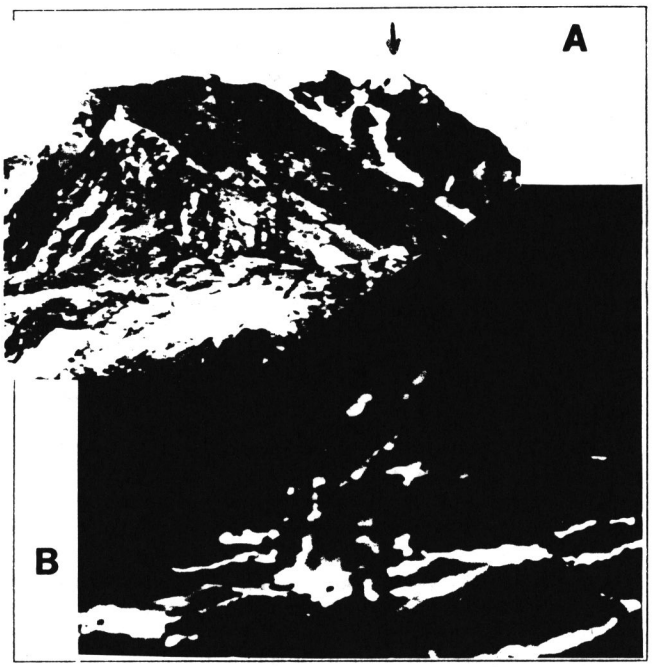

Figure 8. An example of computing a constrained-chance description.

looking tree can be produced by recursively applying a random number generator and simple constraints on branching geometry. In each case a random number generator plus a surprisingly small number of constraints can be used to produce very good models of apparently complex natural phenomena. Thus, there is hope for extending this approach well beyond the domain of land topography.

B. An Example Of Computing A Description

Figure 8 illustrates an actual example of computing such a description. Figure 8(a) is an image of a real mountain. Let us suppose that we wished to use the image data to construct a three-dimensional model of the rightmost peak (arrow), perhaps for the purpose of predicting whether or not we could climb it. I will take the standard fractal technology used in the computer graphics community as the unconstrained "primal" shape generator, as it provides an apparently accurate model of a wide range of natural surfaces.

All that is necessary to construct a description of this mountain peak is to extract shape constraints from the image and insert them into the primal shape generator. The fractal dimension of the 3-D surface is the principal parameter (constraint) required by our fractal shape generator; roughly speaking, it determines the ruggedness of the surface. The fractal dimension of the 3-D surface in the region near the rightmost peak was inferred from the fractal dimension of the image intensity surface in that area by use of the empirically-derived relationship between surface and image. Constraint on the general outline of this peak was derived from the estimated z-values of distinguished points (those with high curvature) along the boundary between sky and mountain. These two constraints, together with the shape generator, *are* a 3-D representation of this peak; the question is: how good a representation? A view of a 3-D model derived from this representation is shown in Figure 8(b). It appears that these simple constraints are sufficient for computing a good* 3-D representation of the peak.

*Rather primitive ray tracing, etc., was used to generate this image; better code is being implemented.

C. What Do We Accomplish With This Approach?

Let's consider the problems cited above:

(1) The problem of representing a complex shape, such as a crumpled newspaper. The problem with a shape-primitive representation such as surface normals, voxcels or generalized cylinders is that the resulting description seems hopelessly complex. Because the constrained-chance representation allows us to deal only with the structural regularities and to ignore inconsequential details, the problem can become much simpler. Thus, for instance, the graphics community has found that constrained-chance fractal descriptions of complex objects (e.g., a mountain) are quite compact and easy to manipulate. It also turns out that many previously simple things, such as describing a smooth plane, remain simple.

How does this representation function when we want to compute a description of a *specific* mountain, bush or other entity from its image? Current "shape-from-x" research furnishes constraints on shape in a variety of forms: surface orientation (from texture, shading), relative depth (from motion, contour), and absolute depth (from stereo, egomotion). It appears to be fairly straightforward to mix each of the various flavors of constraint into the vanilla-flavor shape generator [3,5], although significant research remains to be done. As more shape constraints are obtained from the image, the description becomes more and more precise; i.e., there is less and less chance in the description. Eventually, only one shape satisfies all of the constraints.

How complex could such a description become? The constrained-chance representation would *at worst* be as complex as a two-dimensional array of z values representing the same surface, because we could always use it to actually generate such an array of z values. As mentioned previously, experiments in human perception indicate that our representations are usually not accurate enough to recover every z value. The representation of a particular object, therefore, is likely to be quite a bit simpler than a full depth map.

(2) The problem of representing *classes* of shapes, such as are referred to by the terms "a mountain," or "a bush." Again, the ability to specify important structural details and leave the rest only qualitatively constrained allows simplification of the problem. The definition of "a mountain," for instance, might reasonably consist entirely of a specification of the fractal dimension of the surface and a caveat concerning size. If we are to judge by the results reported in the computer graphics literature, the notion of representation by constrained chance thus allows us, using only a few lines of code, to produce an accurate description of the *class* of shapes we label "mountains," or "bush."

(3) The problem of determining the set of appropriate descriptions when the shape is underconstrained by the sense data. The problem with standard shape-primitive representations is that either we must generate all combinations of shape primitives consistent with the sense data (a very hard problem), or pick a prototype and specify error bounds. The problem with using prototypes plus error bounds is that we are forced to overcommit ourselves by choosing the prototype; e.g., there is something seriously wrong about describing a cube as "a sphere $\pm 0.4r$", even though the cube certainly fits within the specified volume.

Because the constrained-chance representation allows details to be left constrained but unspecified, it allows us to deal with insufficient sense data by simply adding in those constraints that can be deduced from the image data and committing ourselves no further. The result is a programlike description that can be analyzed and manipulated, does not overcommit itself as to object shape, and allows examples of shapes consistent with the image data to be generated and examined.

(4) The problem of determining that a specific description is a member of a more general class. Here the problem

with shape-primitive representations is that there is so much variability among the descriptions of the members of a class such as "mountain" that a description of the class as a whole seems extremely difficult, and determination of class membership even more so.

The problem of establishing class membership by using constrained-chance representations reduces to determining whether the constraints used to specify a particular description are a subset of those of the more general class. A determination regarding class membership is, therefore, exactly equivalent to determining whether one program's output is a subset of another program's output. While such automatic proof is a difficult problem, it is at least tractable and well-defined — unlike the equivalent problem can be when using a shape-primitive representation. Thus, a constrained-chance representation allows a clear and potentially useful definition of what it means to "recognize that x is an y."

Further, because we need only deal with the structural regularities, this problem can become much simpler than it might at first appear. Taking the class "a mountain" to be defined by fractal dimension and overall size (a definition that is actually sufficient to produce realistic mountain shapes) we can, for instance, easily determine that the description computed by us for the mountain peak is in fact a description of part of a mountain — a task that previously seemed to be nearly impossible.

V SUMMARY

Fractal functions seem to provide a good model of natural surface shapes. Many basic physical processes produce fractal surfaces. Fractal surfaces also *look* like natural surfaces, and so have come into widespread uses in the computer graphics community. Furthermore, we have conducted a survey of natural imagery and found that a fractal model of imaged 3-D surfaces furnishes an accurate description of both textured and shaded image regions.

Fractal functions, therefore, are useful for addressing the related problems of representing complex natural shapes such as mountains, and computing a description of such shapes from image data. The following describes the progress achieved toward the solution of these problems.

Computing a description. Characterization of image texture by means of a fractal surface model has shed considerable light on the physical basis for several of the 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 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 (unfamiliar) texture. Specifically, it seems that measurement of fractal dimension can provide (1) evidence of surface texture anisotropy, and (2) an estimate of the perspective texture gradient. Both capabilities are extremely important because they provide a way to obtain independent confirmation of the assumptions on which previously-reported shape-from-unfamiliar-texture techniques are based.

Representing natural shapes. A constrained-chance representation modeled after the fractal techniques used by the graphics community seems useful for representing complex natural shapes, such as a crumpled newspaper or a mountain. The problem encountered when using conventional shape-primitive representations to describe natural surfaces is that the resulting description is often hopelessly complex. Because the constrained-chance representation allows us to deal only with the structural regularities and to ignore inconsequential details, the problem can become much simpler. Thus, for instance, the graphics community has found that constrained-chance fractal descriptions of complex objects (e.g., a mountain) are quite compact and easy to manipulate. Similarly, the problem of representing *classes* of shapes, such as are referred to by the terms "a mountain," or "a bush," can also be significantly simplified.

The encouraging progress that has already been achieved on both of these problems augers well for this approach. It appears that a constrained-chance representation incorporating a fractal model of surface shape will provide an elegant solution for some of the most difficult problems encountered when attempting to progress from the image of a natural scene to its description.

REFERENCES

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

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

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

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

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

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

[7] L. C. Carpenter, "Vol Libre," Computer Generated Movie, 1980.

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

[9] A. P. Pentland, "Local Computation Of Shape," SRI Artificial Intelligence Center Technical Memo 272.

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

[11] 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.

[12] 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.

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

[14] 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.

[15] 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.

[16] 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.

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

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

[19] R. L. Gregory, "Eye And Brain: The Psychology of Seeing," New York, McGraw-Hill, 1972.

[20] D. A. Norman, "Memory and Attention," New York, Wiley, 1976.