

USING FRACTAL DIMENSION FOR TARGET DETECTION IN CLUTTER

The detection of targets in natural backgrounds requires that we be able to compute some characteristic of target that is distinct from background clutter. We assume that natural objects are fractals and that the irregularity or roughness of the natural objects can be characterized with fractal dimension estimates. Since man-made objects such as aircraft or ships are comparatively regular and smooth in shape, fractal dimension estimates may be used to distinguish natural from man-made objects.

INTRODUCTION

Image processing associated with weapons systems is often concerned with methods to distinguish natural objects from man-made objects. Infrared seekers in cluttered environments need to distinguish the clutter of clouds or solar sea glint from the signature of the intended target of the weapon. The discrimination of target from clutter falls into a category of methods generally called segmentation, which derives localized parameters (e.g., texture) from the observed image intensity in order to discriminate objects from background. Essentially, one wants these parameters to be insensitive, or invariant, to the kinds of variation that the objects and background might naturally undergo because of changes in how they are illuminated or the vantage point from which they are viewed.

Fractal dimension is one such parameter; if an object is fractal, it looks like a fractal from a variety of perspectives. The characterization is rooted in physical first principles, albeit conceptually, as opposed to ad hoc texture-based segmentation methods.

In this article, I will summarize an application of fractal dimension estimation to target detection. I begin with a short discussion of fractals and fractal dimension. Although the fractal geometry of nature leads us to the use of fractal dimension as a target discriminant, the application is a computational procedure that can be judged solely by its efficacy in detecting targets. For this reason and for the sake of brevity, I provide here only a short discussion of fractal geometry. (Interested readers should peruse the bibliography for texts on fractal mathematics, physics, graphics, and history.) Then I will summarize work that led me to apply a particular fractal dimension estimator and the results of a target detection where I segmented a jet aircraft from a sea surface background.

FRACTALS AND FRACTAL DIMENSION

A precise physical definition of fractal has not yet appeared, nor is it essential for applications in image processing. More important is the general concept of a

fractal. Falconer¹ defines fractals as objects with some or all of the following properties: fine structure (i.e., detail on arbitrarily small scales) too irregular to be described with Euclidean geometry; self-similar structure, with fractal dimension greater than its topological dimension; and recursively defined. This definition extends fractal into a more physical and intuitive domain than the original Mandelbrot definition whereby a fractal was a set whose "Hausdorff-Besicovitch dimension strictly exceeds its topological dimension."² The fine, irregular, and self-similar structure of fractals can be experienced firsthand by looking at the Mandelbrot set at several locations and magnifications. Better physical examples exist that show the connection between natural processes and fractals in nature. Mandelbrot and Voss² have synthesized realistic landscapes and clouds using fractional Brownian motion. Prusinkiewicz and Lindenmayer³ have syntactically modeled plant growth, producing fractal plant forgeries that are both visually believable and botanically correct. Feder⁴ has shown that diffusion aggregation processes, such as the deposition of films, produce fractals. These physical fractals are not fractals in the mathematical sense, because they have no structure at arbitrarily small length scales. For example, if we look at a cloud with higher and higher magnifications, eventually we begin to see the molecular structure of the atmosphere. This does not present a practical difficulty in target detection, since we are always looking at intermediate length scales in our images.

In our application of target detection, we are most interested in observable differences between the sensor images of natural objects and man-made objects. One such observable quantity is fractal dimension, which is integer for objects with simple Euclidean descriptions (e.g., an aircraft composed of simple curved surfaces) and may be non-integer for complex natural surfaces (e.g., the clutter pattern of the ocean surface).

What do we mean when we assign dimension to some object? When we say that a surface has dimension two, we understand that it takes two coordinates to locate a

point on the surface or that the surface can “live” in spaces with dimension two or greater. Most importantly in this discussion we mean that the “size” of a surface is measured by area, the “size” of a cube is volume, and so on. In short, we can interpret the dimension of an object as the way the capacity (area, volume) changes with respect to a characteristic length. For the normal run of Euclidean objects, dimension is always integer; however, the notion of capacity can extend to non-integral values of dimension.

Consider a bounded surface in space. Let us try covering the surface with line segments, squares, and cubes such that after we have constructed the covering, the set of elements (lines, squares, or cubes) in a covering contain all of the points in the surface. The boxed insert (Measuring the Size of a Surface) illustrates this process. We approximate the area of the bounded surface as the sum of the areas of the squares needed to cover the surface. As we let the edge length of the squares become small, in the limit zero, we have better and better estimates of the surface area. Can we assign a length to the plane in this way? Obviously, we can cover the plane with an infinite number of line segments, and thus the plane has infinite length. Similarly, if we use cubes for the covering elements, we conclude that the surface has zero volume. Thus, we can define the dimension of the surface as the dimension of the covering elements that give us a non-zero but finite capacity. This seems a somewhat awkward definition for dimension. Now let us consider a curve constructed as in the insert (Constructing a Fractal Curve). We start with a line, then replace that line segment with three line segments arranged as shown in the insert, and then replace the smaller line segments each with three smaller line segments, ad infinitum. We can see that the length of the curve is infinite, but its area is

zero. The curve is a fractal, and its dimension is someplace between one and two. Also, the curve is irregular and rough because it is continuous everywhere but differentiable nowhere, is defined recursively, and its self-similar structure is evocative of a snowflake.

What are the geometric objects in our target detection problem? We take an image and make a surface by assigning a height equal to a pixel intensity to each pixel location. When the image is of a natural (fractal) object, the surface may have a fractional dimension. When the image is of an aircraft, for example, the surface will be composed of Euclidean surfaces with integer dimension.

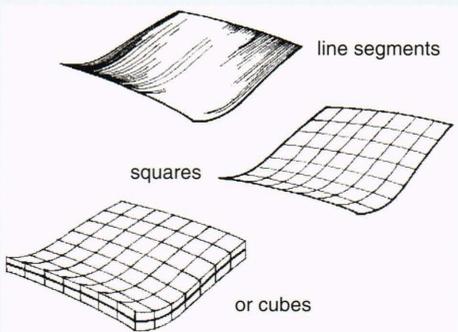
DIMENSION ESTIMATORS

I selected a dimension estimator for images by first testing several dimension estimators on a set of fractals with known dimension. I was interested in estimators that performed well on short data records so that the dimension estimates were localized within small neighborhoods of the image. I was also interested in dimension estimates in the presence of noise, since our images are acquired with real, noisy sensors. For the sake of brevity I discuss only three dimension estimators.

A variety of practical dimension estimators are commonly used. To draw an analogy, when trying to estimate the usually defined Riemann integral of a measured time series, it is not necessarily the best approach to directly use the Riemann sum (i.e., approximation with rectangular areas). Instead, one may be better off using a trapezoidal rule (thus assuming that the signal is best interpolated linearly), or spline fits, or polynomial fits (Lagrangian integration), and so on. All of these methods reduce to the Riemann integral in the limiting case, yet they often have demonstrably better performance than the naive approach. The same is true of dimension estimation: making assumptions about the signal can improve or sometimes degrade the estimation of dimension.

MEASURING THE SIZE OF A SURFACE

We can approximate the size of a surface by covering it with



and adding up the length, area, or volume of the covering elements. As we let the length of the line segments get small, the “length” of the surface becomes infinite. Likewise, the “volume” goes to zero. Only for the covering with squares is the size (area) finite.

Box-Counting Dimension

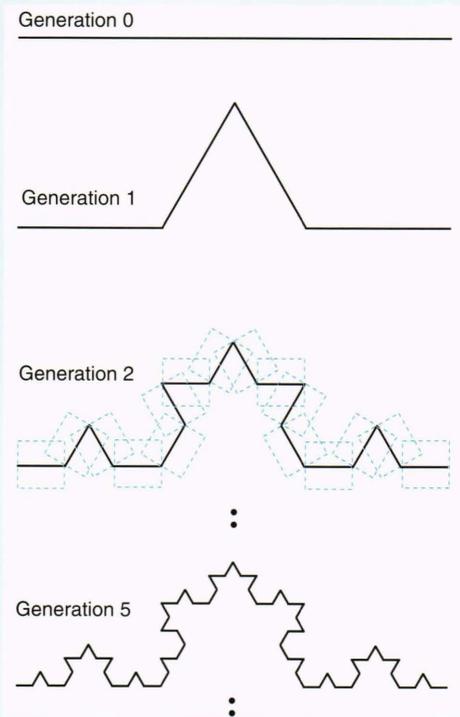
The first step in finding the box-counting dimension estimate is to compute a function $M_b[A, \delta]$ of the set A . This measure counts the number of boxes, fixed on a lattice of edge length δ , that are needed to cover A . Then the dimension of A , D_b , is the number d that satisfies

$$\delta^2 M_b[A, \delta] = \text{capacity} = C\delta^d, \tag{1}$$

where C is a constant. Note that the slope of the graph of the logarithm of capacity versus the logarithm of δ is simply related ($D - 2 = d \ln M/d \ln \delta$) to the dimension. In fact, when we are interested in dimension over limited scales, we will estimate the slope over those scales. Note also that box counting is ill-defined under affine transforms of the graph itself. In the Kiesswetter⁵ curves shown in Figure 1, if the vertical scale is magnified or minified, the measure of the graph will change. Box-counting dimension is problematic for self-affine graphs since one can manipulate scales to get any dimension. A practical problem associated with all dimension estimat-

CONSTRUCTING A FRACTAL CURVE

Take a line segment (generation 0) and replace its center third with two line segments, each having the same length as the center third (generation 1). Take each new line segment and replace its center third with two line segments as before (generation 2). Do this ad infinitum.



If we approximate the area of this curve by covering it with square tiles that have edges as long as the line segments, we have

$$\begin{aligned} \tilde{A} &= \text{tile area} \times \text{number of tiles} \\ &= \left(\frac{1}{3}\right)^{dn} \times 4^n = \left(\frac{2}{3}\right)^{2n}, \end{aligned}$$

where n is the generation number, and d is the tile dimension, two. Since

$$\lim_{n \rightarrow \infty} \tilde{A} = 0,$$

the area of the curve is zero. If we approximate the curve length by tiling with line segments,

$$\begin{aligned} \tilde{L} &= \text{segment length} \times \text{number of segments} \\ &= \left(\frac{1}{3}\right)^{dn} \times 4^n = \left(\frac{4}{3}\right)^n \end{aligned}$$

and

$$\lim_{n \rightarrow \infty} \tilde{L} \rightarrow \infty,$$

the length of the curve is infinite. To find the dimension, let us define capacity as

$$M(d) = \left(\frac{1}{3}\right)^{dn} \times 4^n.$$

Then the capacity is finite when

$$1 = \left(\frac{1}{3}\right)^{dn} \times 4^n \Rightarrow d = \frac{\ln 4}{\ln 3} \cong 1.26,$$

so the fractal dimension is 1.26.

ing methods that use a measure and then estimate d is that a real fractal may only scale over a limited range of δ 's. How do we choose the range to estimate d ? In Figure 2, a saturation region seen on the left should be excluded, for example. As the size of the boxes approaches the size of the data sample lattice, the boxes will contain at most one sample point from the graph, so that $M_b[A, d]$ approaches N , where N is the number of samples.

Hurst Dimension

The Hurst dimension⁴ is estimated by computing the cumulative range of the graph over intervals $\{\delta_i\}$. Thus, the Hurst method avoids the complications that box counting has with self-affine graphs, since the vertical scale is self-defined. First the measure $M_h[A, \delta]$ is computed by

$$\begin{aligned} M_h[A, \delta] &= \Sigma (\text{Max}[A_i] - \text{Min}[A_i]) \\ A_i &= \{y: i\delta \leq x < (i+1)\delta\}, \end{aligned} \quad (2)$$

where the sum ranges over $\{i\}$ that cover the domain of A (Fig. 3). Then the term d that satisfies the hypothesis that

$$\delta M_h[A, \delta] = \text{capacity} = C\delta^d \quad (3)$$

is the Hurst dimension D_h (Fig. 4).

Spectral Dimension

Spectral dimension estimates are based on the assumption of fractional Brownian motion⁶ (fBm). If the function is fBm, then the correlations scale as

$$\langle |x[t+T] - x[t]| \rangle = T^{-H} \langle |x[t+1] - x[t]| \rangle, \quad (4)$$

where $\langle \rangle$ is the expectation operator, $D_s = 2 - H$, and it can be shown⁶ that

$$|X[k]|^2 = Ck^{-2H-1}, \quad (5)$$

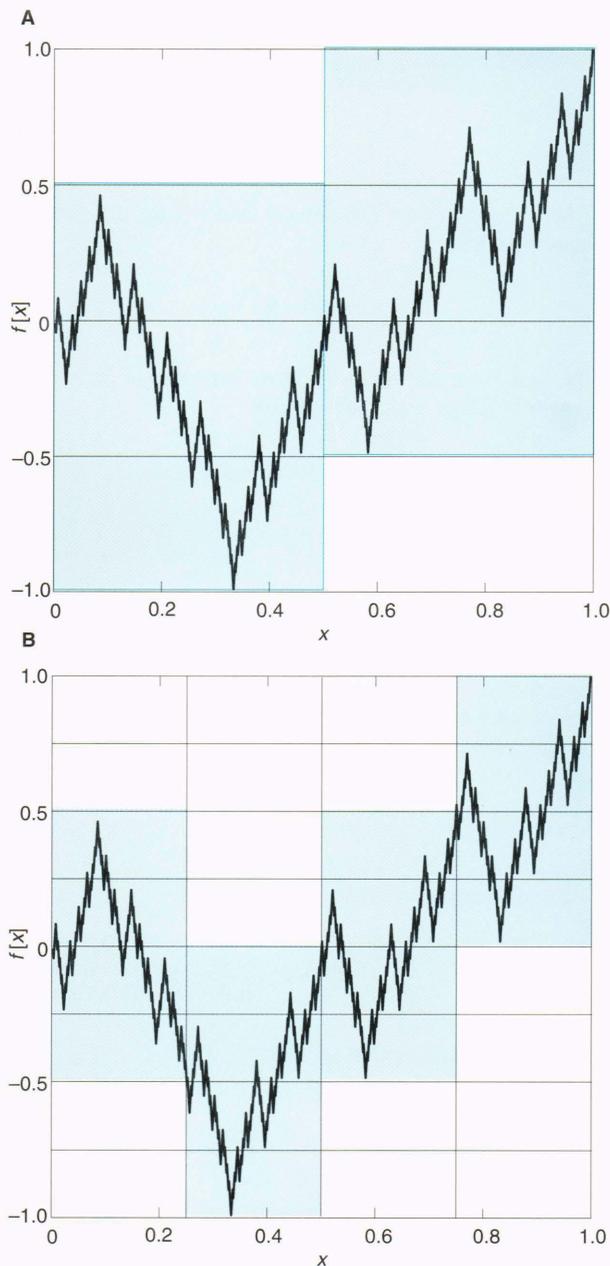


Figure 1. Box-counting measure is the number of boxes the graph intersects as a function of the box size. **A.** Boxes of $1/2 \times 1/2$. **B.** Boxes of $1/4 \times 1/4$. ($f =$ Kiesswetter curve.)

where $X[n]$ is the discrete Fourier transform of a discrete signal $x[n]$. Thus, the estimation of spectral dimension involves computing the sample power spectral density $|X[k]|^2$ and then finding the slope of the log-log presentation.

Figure 5 is an example where the function is not fBm but still has correlations of the form (Eq. 4), and thus the estimate is relatively good. A counterexample is a sawtooth wave that has $D_b = D_h = 1.0$, whereas $D_s = 1.5$.

Evaluating Dimension Estimators

The application of several dimension estimators to the functions in Table 1 are summarized in Table 2. The box

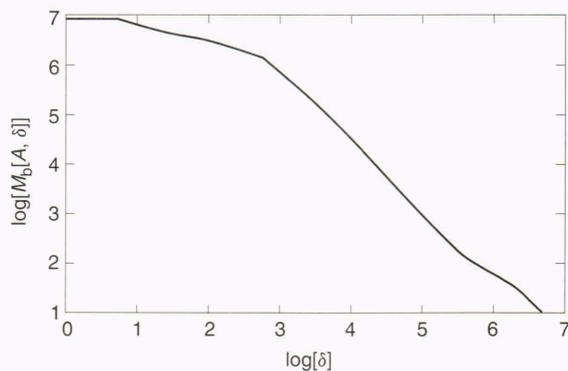


Figure 2. A box-counting measure in log-log presentation. The slope of the line equals $D_b - 2$.

estimate is consistently low. The Hurst estimate is compressed about $D = 1.5$ but accurately estimates the dimension of the ramp and the noise process. The spectral estimate performs better than the Hurst estimate for the three fractals but is grossly incorrect for the ramp and noise process. The parenthetic entries in the spectral estimate column are the expected results: It is possible to show that the spectral estimation algorithm should yield $D_s = 1.5$ for a ramp and $D_s = 2.5$ for an independent noise process. Figure 6 presents the results juxtaposed against the ideal, $D = D_{b,h,s}$. This clearly demonstrates the superiority of the Hurst estimate in this context.

ESTIMATING DIMENSION FROM NOISY RECORDS

The estimation of fractal dimension begins with the calculation of an approximate capacity measure, the rate at which covering area changes with covering tile size. The subsequent step involves finding the region of this curve that satisfies a hypothesis of an exponential relationship between (using a two-dimensional example) the covering area and tile size. Unfortunately, we know of no first-principles method for determining what constitutes a good scaling region that satisfies the hypothesis well.

Although one can extract scaling regions by eye, a subjective approach compromises quantitative results. More important, when computing a dimension image from an infrared image where the scaling region determination must be made for each pixel in the image, there must be an automatic and thus objective criterion function for doing so. In addition, the image probably contains high-dimensional noise and will thus have two scaling regions if the signal-to-noise ratio is moderate and the underlying image is pseudofractal rather than not fractal at all. We must consider how to segment the capacity measure into multiple scaling regions, since

1. The set may be a *pseudofractal* that only scales as a fractal over a limited range of lengths; or
2. The set may be a mixture of several fractals, pseudofractals, or nonfractals and thus may have different scaling characteristics in different length ranges; or
3. The set may not be a fractal at all.

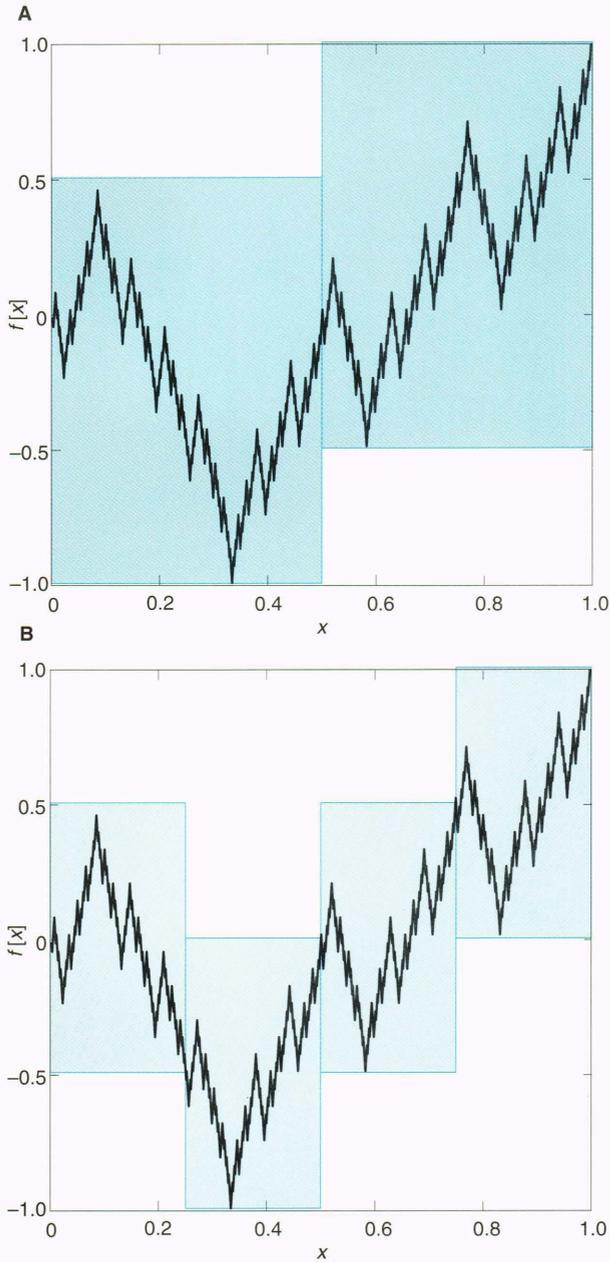


Figure 3. The Hurst measure is the area covered by the ranges of the graph in intervals. **A.** An interval of $1/2$. **B.** An interval of $1/4$. ($f =$ Kiesswetter curve.)

The way we handle these different possibilities can cause significant disagreement in dimension estimates that purportedly relate to a specific set. For example, if one were to estimate the dimension of cloud infrared radiance using a very noisy infrared imaging sensor, and also depended on short length scales on the order of the imager pixel size, then the estimate would likely be around $D = 2$. Conversely, a low-noise imager and longer length scales provide an estimate of $D = 1.16$. Because of these problems, dimension estimates should always be accompanied by the range of scales used to make the estimate. We assume that the log-log capacity measures

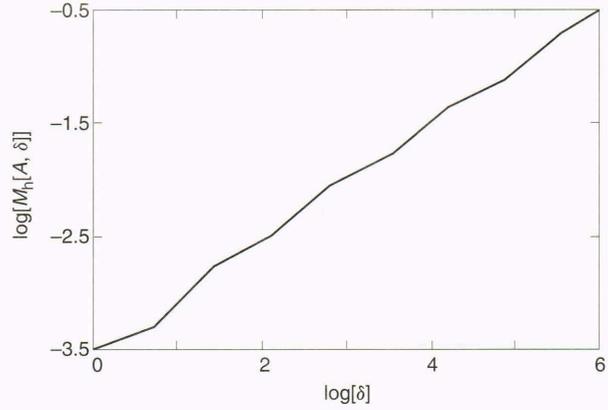


Figure 4. The Hurst measure of Figure 3 in log-log presentation. The slope of the line equals $D_h - 1$.

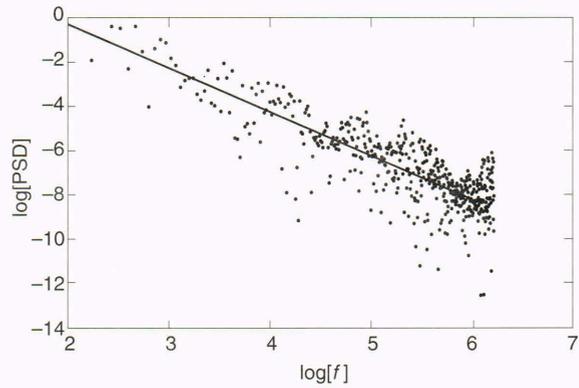


Figure 5. The power spectra density (PSD) of the Kiesswetter curve (f) and the least-square-error line fit to the data ($D = 1.53$).

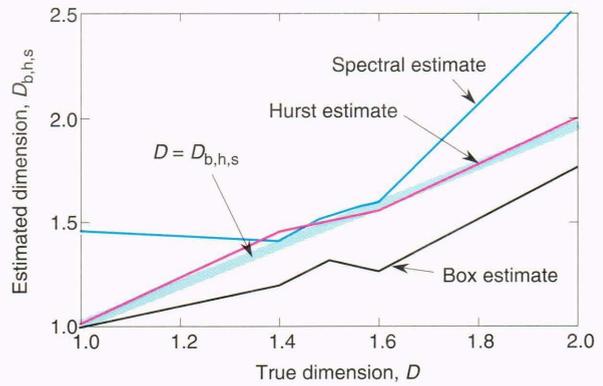


Figure 6. Comparison of dimension estimator performance.

are invariant, except for translation of the function origin, under affine transforms of the fractal graph. This is not true, as previously noted, for the box-counting method.

We create a graph with two scaling regions and different dimensions by adding a normally distributed, independent, random noise process to a Kiesswetter curve (Fig. 7). The Kiesswetter curve has $D = 1.5$, and the noise process has dimension $D = 2.0$. The signal-to-noise ratio

Table 1. Summary of function properties.

Function	Type	Scaling	Dimension
Weierstrass–Mandelbrot	Deterministic	Self-affine	1.6
Kiesswetter	Deterministic	Self-similar	1.5
Fractional Brownian motion	Random	Self-affine	1.4
Ramp	Deterministic	Self-similar	1
Uniform pseudorandom process	Random	Self-affine	2

Table 2. Dimension estimation results.

	Box estimate	Hurst estimate	Spectral estimate	Hausdorff–Besicovitch dimension
Ramp	0.93	1.00	1.45 (1.50) ^a	1.00
Fractional Brownian motion	1.21	1.45	1.40	1.40
Kiesswetter	1.32	1.51	1.53	1.50
Weierstrass–Mandelbrot	1.26	1.54	1.59	1.60
Uniform independent noise	1.76	1.99	2.52 (2.50) ^a	2.00

^aExpected results.

of the sum is about 4. In Figure 8, we display the expected range of the Kiesswetter curve, the best exponential regression, and note that the slope of the line implies a dimension $D_h = 1.488$ such that

$$M_h[f, T] = T^{(2-D_h)} M_h[f, 1]. \quad (6)$$

The measure $M_h[T]$ for the noisy Kiesswetter curve $f_n[x]$ will obviously satisfy

$$M_h[f_n, T] = M_h[f, 1] \quad (7)$$

for small T , since the (independent) noise process will dominate the measure and will satisfy Equation 6 for $D_h \cong 1.5$ and large T .

Figure 9 shows how both noise and Kiesswetter curve dimensions are extracted from different sections of the measure. The sections were selected by computing the curvature of a polynomial fit to the measure and then selecting scaling regions where the curvature was less than some threshold. Here, the portions of a 9th-degree polynomial fit with curvature less than 0.15 were selected (thresholded) as scaling regions.

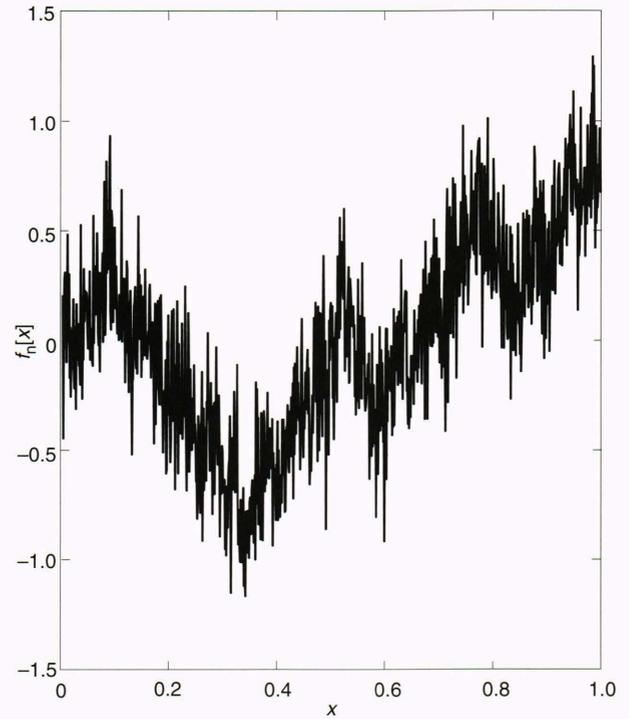


Figure 7. A noisy Kiesswetter curve (f_n) with a signal-to-noise ratio of 4.

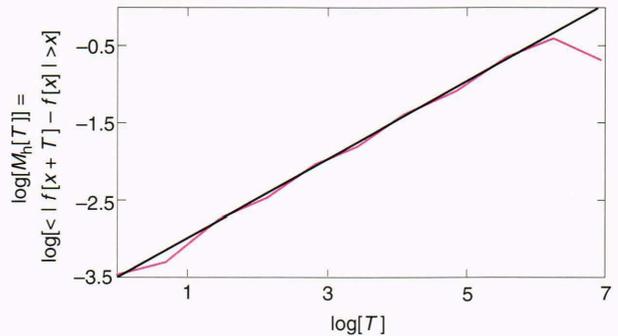


Figure 8. A log-log graph of the expected range of the Kiesswetter curve without noise (blue curve) and the best linear fit (black curve). (All logarithms are taken to base e .)

IMAGE SEGMENTATION USING DIMENSION ESTIMATES

We now look at an application of fractal dimension estimation to a real signal processing problem: We segment a T-38 jet aircraft from a benign clutter background in an infrared image. The data are taken from experiments in support of the algorithm design for an infrared seeker demonstration program. We find that by first computing the local Hurst dimension of the surface of infrared irradiance as a function of position (which results in a dimension image) and then thresholding the dimension image, we obtain a more robust estimate of target size and location than that provided by immediate thresholding of the infrared image.

Although the experiment is not conclusive, it suggests a possible scheme for improved detection of extended targets in clutter. The algorithm for computing the local Hurst dimension of images is straightforward, requiring only the maximum and minimum of the image over a small set of neighborhood sizes and the computation of the least-squared-error fit to the function of image range over neighborhood. The experiment has also suggested that ways be found to incorporate *a priori* knowledge of target shape into the dimension segmentation scheme. In this example we use a square neighborhood for the estimate, but using a target-shaped neighborhood would be better.

The segmentation of candidate threats from background is an early step in the signal processing associated with infrared seekers. Typically, the image is segmented by computing some threshold irradiance value R_T based on the sample statistics of the entire image (e.g., variance), and then the image is thresholded such that pixels exceeding R_T are labeled as threat and all other pixels are

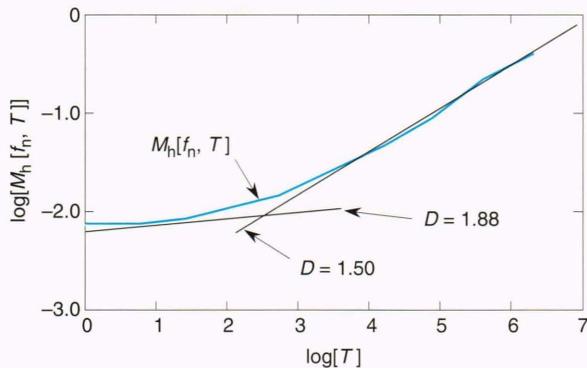


Figure 9. By selecting portions of M_h with low curvature, both the noise dimension (for small T) and the Kiesswetter curve dimension (for large T) are estimated. (f_n = noisy Kiesswetter curve).

labeled not-threat. This constitutes a needed data compression step: We reduce a large number of possible threat locations to a relatively smaller set. Normal practice is to assign a nonzero constant value (e.g., 1 to values that exceed the threshold and 0 otherwise) to minimize memory requirements. The segmented image is then processed with a variety of false-alarm suppressing algorithms such as *m-of-n* detection, track logic, and so on.

The Infrared Data

I selected three images of a T-38 jet against a variety of sea surface conditions for this experiment. The data are from a 3- to 5- μm -waveband mercury-cadmium-telluride 128×128 staring-array infrared detector. The images were digitized on the Mercury image processing workstation in APL's Electro-Optics Systems Group Image and Signal Processing Laboratory and then transferred to a DEC MicroVax GPX II image and signal processing workstation for further processing using Image Description Language as a programming environment.

The images have clutter backgrounds that are benign (low-level sea irradiance), structured solar sea glint, and saturated solar sea glint (Fig. 10). In Figure 10A, the jet is faintly visible near the center of the infrared image as a lighter spot. In Figure 10B, the jet is invisible, although it can be discerned when the motion of the jet against the background is observed in the video tape. Humans appear to discriminate jet from sea glint in Figure 10B (with motion) by discriminating on the basis of subtle textural cues: Fractal dimension correlates well with human perception of texture,¹ and thus a time series from the scenario in Figure 10B is an ideal candidate for fractional dimension/motion detection study. In Figure 10C, the spatial density and size of glint patches cause each image pixel to include at least one solar image, and the jet is discerned by its obscuration of the glint.

In what follows we will only use Figure 10A for analytical and experimental purposes. Simple application of dimension estimation on a single image was insufficient

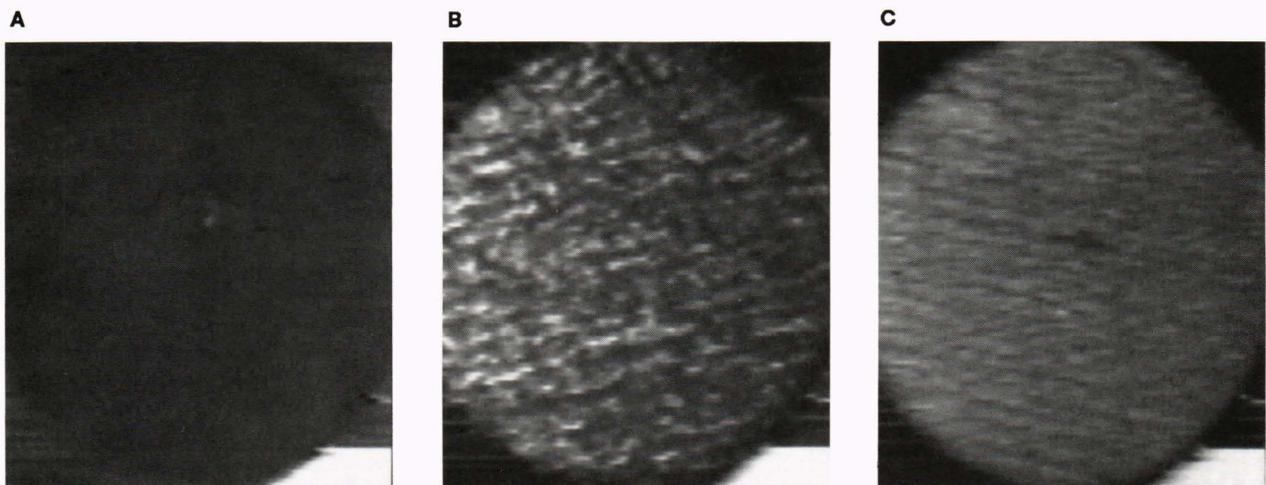


Figure 10. Examples of targets in background clutter. **A.** Jet in benign clutter. **B.** Jet in structured glint. **C.** Jet in saturated glint.

to extract the target in Figure 10B, and Figures 10A and 10C are similar enough that we suppose dimension estimation should work on both.

Computing the Dimension Image from the Infrared Image

We compute the dimension image using the Hurst dimension estimation algorithm, previously described, which is easily extended to two dimensions by finding the range of the image intensity value over a square area rather than an interval.

The image in Figure 10A can be thresholded immediately to segment target from background. Figures 11A and 11B show Figure 10A thresholded at gray levels of 170 and 160, respectively. In Figure 11B the reader can discern the T-38 location and outline, but it is unlikely

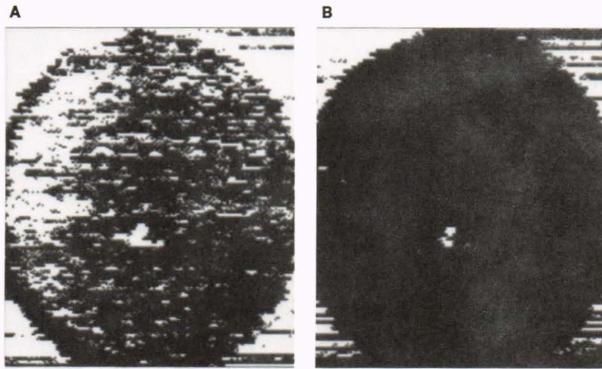


Figure 11. Examples of simple threshold sensitivity. **A.** The image of Figure 10A thresholded at 170. **B.** The image of Figure 10A thresholded at 160.

that an image processing algorithm could do so generally. When we compare Figure 11B with Figure 10B, we see that the clutter structure is similar and thus we suppose that the clutter in Figure 10A is infrared skylight reflection from the sea surface or some similar phenomenon. Figure 11A shows Figure 10A thresholded at a higher level, and we see that whereas the number of false detections is greatly reduced, we have also lost the outline of the T-38.

Figure 12A is the dimension image computed from the image in Figure 10A by the Hurst method, with ranges measured over intervals of {1, 2, 4, 8, 16, 32}. The dimension at a pixel is computed from the log-log slope of the average ranges over these intervals. Figure 12B is the result of thresholding the dimension image at $D \approx 2.0$. The bounding polygon in Figure 12A represents a border about 16 pixels from the jet outline proper. Note that the interior of the bounding polygon has $D \approx 2.0$, whereas most of the remaining image has $2.5 \leq D < 3$, as we would expect.

Let us compare the two methods of segmentation—thresholding the irradiance image versus thresholding the dimension image:

1. The threshold for the dimension image is selected on first principles, since the dimension of the jet should be $D \approx 2.0$, whereas an irradiance threshold is selected from image statistics.
2. The thresholded dimension image provides an outline of the target. The thresholded irradiance image gives only a few pixels on the target when the threshold is low enough to give a reasonable false-alarm rate.
3. The method using estimation of dimension is relatively simple, although not as simple as thresholding the irradiance image.

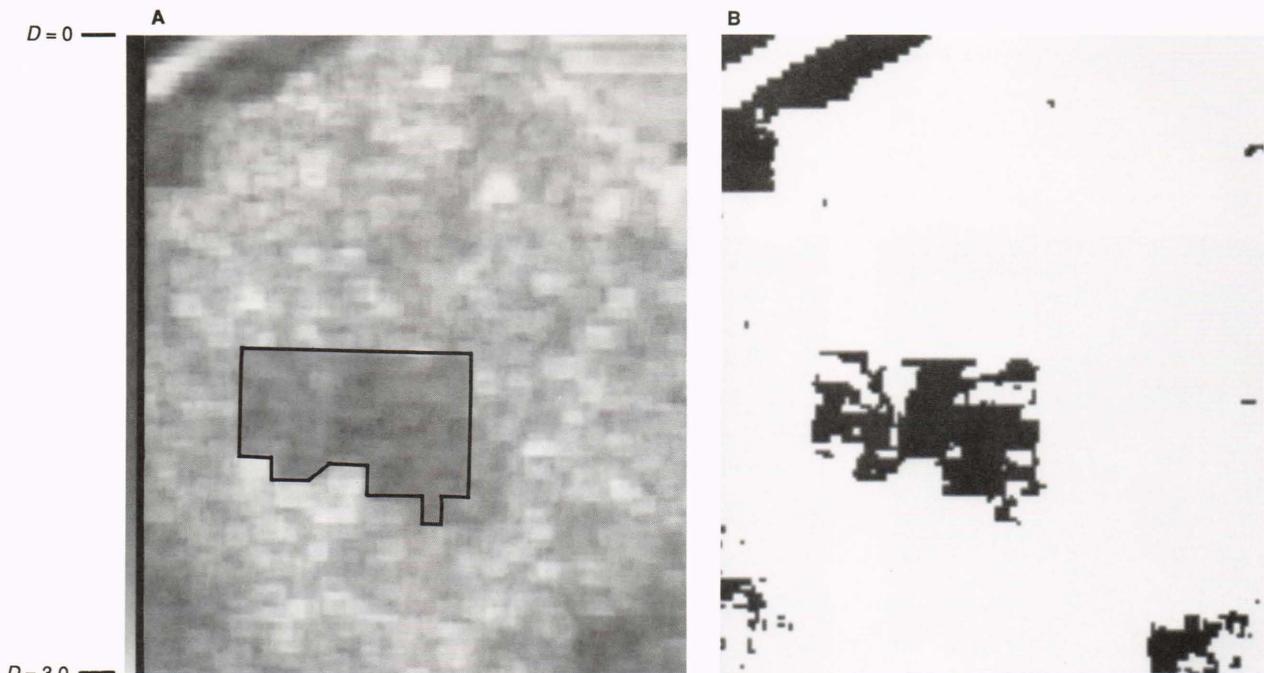


Figure 12. Locating the jet with fractal dimension. **A.** The dimension image corresponding to Figure 10A. The outline bounds a region of low dimension. The strip to the left is the dimension-to-gray scale correspondence. **B.** The dimension image thresholded at $D \approx 2.0$.

The selection of threshold on *a priori* dimension difference between target and clutter fosters robustness in the segmentation. Target outline may be better than target hot spot in endgame guidance, since the best aim point is not always the hottest spot on the target. Finally, dimension images probably can be computed in flight hardware within time, size, power, and weight constraints, since the algorithm is so simple.

FUTURE DIRECTIONS

Although fractal dimension estimation has been used for segmentation of images, unresolved practical issues remain. Dimension estimates should give good results when applied to small data sets, determine over what range of scales the estimate holds, work when the observations include modest high-dimensional noise, and have low computational complexity. I have performed a numerical experiment to investigate these problems and have found an estimator, called the Hurst estimate, that gives good results on one-dimensional fractal approximations. I demonstrated a technique for finding scaling regions for the estimate. When the estimator was extended to the two-dimensional case, I was able to segment an airplane from a cluttered image.

My work so far has addressed the issues of noise, short data records, best estimators, and object shape in image segmentation. Many issues remain to be addressed, such as characterizing the fluctuation of dimension estimates around the true dimension of an image, and implementation of matched filters on dimension images. Work on practical implementations (e.g., algorithm design and potential for application-specific integrated circuit implementations) of these algorithms is yet to be undertaken, but the approach discussed in this article holds promise for new missile signal processing applications.

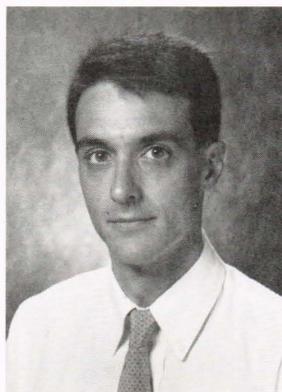
REFERENCES

- ¹Falconer, K., *Fractal Geometry: Mathematical Foundations and Applications*, John Wiley & Sons, New York (1990).
- ²Mandelbrot, B. B., *The Fractal Geometry of Nature*, W. H. Freeman, San Francisco (1982).
- ³Prusinkiewicz, P., and Lindenmayer, A., *The Algorithmic Beauty of Plants*, Springer-Verlag, New York (1990).
- ⁴Feder, J., *Fractals*, Plenum Press, New York (1988).
- ⁵Dubuc, B., Quiniou, J. F., Roques-Carnes, C., Tricot, C., and Zucker, S. W., "Evaluating the Fractal Dimension of Profiles," *Phys. Rev. A* 39(3), 1500-1512 (1989).
- ⁶Peitgen, H.-O., and Saupe, D. (eds.), *The Science of Fractal Images*, Springer-Verlag, New York (1988).

SELECTED BIBLIOGRAPHY ON FRACTALS

- Barnsley, M. F., *Fractals Everywhere*, Academic Press, New York (1988). An almost self-contained presentation of Iterated Function Systems. Includes tutorial introduction of the needed mathematics, chaotic dynamics on fractals, fractal dimension, fractal interpolation, and applications of Iterated Function Systems in data compression.
- Devaney, R. L., and Keen, L. (eds.), *Chaos and Fractals: The Mathematics Behind the Computer Graphics; Proceedings of Symposia in Applied Mathematics* 39, American Mathematical Society, Providence, R. I. (1989). A series of lectures on fractals and chaos and the interrelations between the two. Includes an introduction to Iterated Function Systems.
- Falconer, K., *Fractal Geometry: Mathematical Foundations and Applications*, John Wiley & Sons, New York (1990). A rigorous introduction to fractals, including a variety of dimension definitions and techniques for calculating dimension and operations on fractals. Includes physical and mathematical examples of fractals and their applications.
- Feder, J., *Fractals*, Plenum Press, New York (1988). A physically conceptual introduction to fractals and discussion of fractals in viscous fingering and diffusion-limited aggregation, percolation, fractal time series, Brownian motion, and fractal surfaces. Discussions of empirical evidence for fractal clouds, waves, and landscapes.
- Mandelbrot, B. B., *The Fractal Geometry of Nature*, W. H. Freeman, San Francisco (1982). The classic work on fractals in nature. Can be difficult going for the uninitiated and requires some background in physics and mathematics.
- Pietgen, H. O., and Saupe, D., *The Science of Fractal Images*, Springer-Verlag, New York (1988). A survey of topics in fractals, including an introduction by Mandelbrot, fractals in nature, algorithms for generating fractals, chaotic system and fractals, and modeling real world images with fractals. Many excellent color plates.
- Prusinkiewicz, P., and Lindenmayer, A., *The Algorithmic Beauty of Plants*, Springer-Verlag, New York (1990). A presentation of plant models using syntactic systems and virtual laboratories for studying plant growth and form. Includes a discussion of the fractal properties of plants and many stunning graphical examples of plant models.

THE AUTHOR



KIM T. CONSTANTIKES is a senior staff engineer in APL's Fleet Systems Department. He received his B.S. degree in applied physics in 1983 from Ohio University. He then attended Carnegie Mellon University with a Benjamin Franklin Fellowship and received an M.S. degree in electrical and computer engineering in 1985. At Carnegie Mellon, he researched topics in pattern recognition and sensors in the Smart Sensors Laboratory of the Robotics Institute. Mr. Constantikes joined APL's Electro-Optic Systems Group in 1985, where he has worked on various

computational and algorithmic aspects of imaging systems used for target acquisition, tracking, and recognition, as well as image-matching navigation. His current interests include image and scene modeling, custom VLSI implementation of image processing algorithms, image-matching navigation, target detection in clutter, computer graphics, and computational theory of vision.