

The Perception of Natural Contour

David L. Gilden, Mark A. Schmuckler, and Keith Clayton

The observation that natural curves and surfaces are often fractal suggests that people may be sensitive to their statistical properties. The perceptual protocols that underlie discrimination between fractals and between other types of random contour and fractals are examined. Discrimination algorithms that have precisely the same sensitivities as human observers are constructed. These algorithms do not recognize the integrated scale hierarchy intrinsic to fractal form and operate by imposing a metatheory of structure that is based on a signal-noise distinction. The success of the algorithms implies that (a) self-affinity in random fractals is not perceptually recovered and (b) people have a natural disposition to view contour in terms of signal and noise. The authors propose that this disposition be understood as a principle of perceptual organization.

The environment that we live in has essentially two architectural components: One is carpentered, designed, and built by people; the other is everything else, the material form of nature. If one observes carpentered structures with an unjaded eye, it is difficult not to be struck by the smoothness of the surfaces and the cleanness with which the lines are cut. Even the crudest and least adorned structures have these properties. The things that people make are at least minimally designed, and the primitives of design are lines and planes. This is as true of primitive structures and implements as it is of the things that are built today. An inspection of natural structures reveals an entirely different order. The boundaries that form natural surfaces and contours are often not smooth. Natural form—landscapes, mountain ranges, coastlines, stream paths, clouds, tree lines, vegetation cover—is irregular and rough in appearance. The apparent transparency of this observation belies the subtlety that is required to fully appreciate its import. Geometric descriptions of natural structures required the development of a new set of elements that differ radically from those that comprise Euclidean geometry as well as new modes of analysis that depart from the smoothness assumptions on which differential geometry rests.

Real analysis, as developed by Cauchy, Weierstrass, and Bolzano, treats structures that have specific properties under mag-

nification. The objects that fall within its purview are smooth and regular when looked at with sufficient resolution. The global structure is completely unconstrained, but the local structure is smooth. This conception of structure pervaded the historical development of the calculus. The idea of zooming in on locally smooth neighborhoods and continuing to the limit is what allowed the calculus to be founded on a rigorous method of proof that was independent of the completed infinities that characterized the first formulation by Leibniz. It is an empirical question whether the assumptions of real analysis provide a useful point of departure for describing natural form, and it was in direct response to an empirically motivated question that directed attention to mathematical objects that are not smooth under magnification. Fractal geometry found its first application in the seemingly mundane analysis of the length of the coastline of Britain (Mandelbrot, 1967).

A structural alternative to smoothness under magnification is self-similarity under magnification. Self-similar contours have the property that magnification brings into focus another level of structure that is isomorphic to the global pattern from which it emerged. This form of complexity is characterized by an infinite nesting of structure on all scales, and there is no convergence to a smooth limit on infinite refinement. Geometric objects that have this nesting property have been called *fractals*, because it is possible to define a measure of dimension on them that is nonintegral (Mandelbrot, 1983, and references therein). One consequence of self-similar nesting is that fractals may be continuous but nowhere differentiable. Consequently, there is no differential geometry of fractals; there is no place on a fractal where a derivative can be defined.

The primary application of fractal theory has been in discussions of physical processes that are turbulent. The hydrodynamical equations that describe geophysical change generally have turbulent solutions (Tennekes & Lumley, 1983) because of the high Reynolds numbers that characterize terrestrial gas and fluid flows. Turbulence is not a particularly well understood physical process, but it is generally conceived of in terms of an energy flow that cascades down through a nested hierarchy of fluid structures known as *eddies*. Energy is extracted from the mean flow by eddies that have sizes on the order of the global

David L. Gilden, Department of Psychology, Vanderbilt University, and Department of Psychology, University of Texas at Austin; Mark A. Schmuckler, Department of Psychology, University of Toronto, Toronto, Ontario, Canada; Keith Clayton, Department of Psychology, Vanderbilt University.

This work was supported by a grant from Vanderbilt University to David L. Gilden and by National Institutes of Health Grant P30-EY08126. We wish to thank David Bloom and Michael Bendele for programming assistance. We also wish to acknowledge conversations with David Knill and David Field, whose work inspired our own. We also wish to thank Sandy Pentland and Ennio Mingolla for helpful comments on a draft. Finally, we wish to thank B. Skett for continued support and encouragement.

Correspondence concerning this article should be addressed to David L. Gilden, Department of Psychology, Mezes 330, University of Texas, Austin, Texas 78712.

flow dimension and is finally deposited by eddies on microscopic scales through molecular viscosity. The range of sizes that participate in turbulent flow can be on the order of 10^9 . Natural contours that are created by turbulent processes express their dynamical origin through fractal formations. Such formations reflect the huge range of scales that are entrained by turbulent flows and are natural simulations of an infinite recursive process. It is in the domain of landscape geometry that fractals have had their most celebrated application. The shapes of clouds, coastlines, tree lines, and topography in general are all fractal. This assertion has been verified by ecological surveys of environmental form (Burrough, 1981; Keller, Crowner, & Chen, 1987) as well as by computer simulations of landscape features (Voss, 1985, 1988).

There is a fundamental difference between objects of differential geometry and those envisioned by a fractal theory of form. This difference essentially has to do with the way roughness is regarded; that is, is it an intrinsic part of the object or something extra that has been added on? The classical view of object structure is expressed by the latter proposition, that roughness is a surface phenomenon that does not penetrate into the form that constitutes the object per se. The objects of classical analysis are composed of compact differentiable manifolds, smooth curves or surfaces that include their boundaries. In this view, natural contours consist of a superficial coating of texture or irregularity that is attached to a compact underlying structure.

This understanding is implicit in formal theories of vision. The pervading influence of differential geometry has guided inquiry in virtually every aspect of formal modeling, including that of image segmentation, object parsing, shape from shading, depth, and motion. (Examples of computational theory in this domain are found in Barrow & Tenenbaum, 1986; Biederman, 1987; Carlton & Shepard, 1990; Hoffman & Richards, 1984, 1988; Koenderink & van Doorn, 1988; Richards, Dawson, & Whittington, 1988; Richards, Koenderink, & Hoffman, 1988.) The tacit conception of structure that underlies computational theories of vision is that the global structure can be treated separately from local irregularity. The global structure is presumed to have the necessary feature of smoothness under magnification, whereas the properties of the local structure are of secondary interest and are not explicitly considered in theoretical models. Within this phenomenology, roughness is regarded as a noise that does not add anything coherent to the object; it simply acts as a mask. It is, epistemically, on top of the object, a cloak that surrounds an underlying smooth form.

In this article we contrast these two conceptions of structure in an inquiry that examines how roughness is perceived to be embedded in natural form. An inquiry of this type makes contact with basic issues in epistemology and metaphysics. What distinguishes this work from pure phenomenology is that we contrast mathematical and physical understandings of structure with perceptual understandings that are revealed in a laboratory setting. The difficult part of this investigation is, of course, gaining access to the intensional aspects of perception. The tool that we use to infer subjective understanding is numerical simulation. Discrimination sensitivities in three studies of the perception of rough contour are modeled by explicit algorithms. A particular class of algorithms are developed that have

exactly the same sensitivities to structure as human observers. The instruction set that characterizes these algorithms provide a sufficient account of how natural contour is perceived.

The Phenomenology of Roughness

Analytic work in perceptual theory treats natural objects as having a two-part structure: an intrinsic core that is smooth and differentiable, which is supplemented by added-on roughness. This split is not only a structural distinction, it is a distinction in value as well. The added-on part has reduced value by virtue of its status as being extra, that is, independent of the defining core. This bipartite understanding of object structure must not be regarded as necessitated by physical theory. In fact, theories based on fractals do not make this type of structural distinction. Rather, the partitioning of object structure into an intrinsic component and an added-on component reflects prior conceptual commitments that reveal how people think about structure.

The following equation captures the splitting of structure into two components:

$$\text{phenomenal object} = \text{signal} + \text{noise.}$$

This equation is to be understood as part of a metatheory of object structure and not as an equation of physics nor as an equation that is motivated by a compelling theory of natural structure. This is, however, a familiar equation, and it is particularly important in statistical and psychophysical theory. Appreciating its utility in these contexts will clarify to a large extent the way in which this type of thinking becomes generalized. Distinguishing between these object components is more easily done in practice than in principle. To develop a robust distinction between signal and noise, as well as to develop a language that will allow contact to be made with fractal structures, we articulate the difference in terms of their respective transformational properties. The following examples show how this can be done.

Two contexts in which this equation explicitly appears are data analysis and auditory and visual masking. In theories of data modeling, such as the general linear model, an observed datum is related to some underlying structure through the following equation, $y_{\text{obs}} = y_0 + e$, where y_{obs} is the observed value and y_0 is the underlying structure that would have been observed in the absence of error, e . y_0 is the signal that one attempts to deduce through the logic of hypothesis testing and the imposition of an experimental design. What makes y_0 something that can be distinguished from the observation is that it represents a natural kind, something that exists in nature that structures the outcomes of experiments. It attains its value as a signal by virtue of its invariance in the transformation that consists of forming the ensemble of possible experiments. The noise, however, is not invariant and is conceived to vary across samples.

This notion of invariance also serves to distinguish signal from noise in masking paradigms. Consider a pure tone that is masked by noise. What uniquely identifies the tone as a signal against the noise background is the invariance of the tone in time. Successive time slices of the stimulus yield an ensemble that consists of two parts: discrete samples of a sinusoid super-

imposed on random fluctuations. The tone component is invariant in the sense that the slices sample a single coherent object, that is, a normal mode of a vibrating system. The noise component, however, is not an invariant quantity because slices through it sample the distribution of a stochastic variable.

In spatial masking, the signal is an underlying shape. The notion of shape has been extremely resistant to definition (see Blum, 1973, for one attempt) and an attempt at definition will not be made here. In this article we proceed informally and consider what is entailed by the breaking of camouflage. For a figure to be perceivable in the presence of overlapping contours, it is necessary that some set of contours be perceptually organized into a coherent form. The complement of this set, although not devoid of perceptual organization, is not treated as a figure. Let us consider these two sets in terms of their equivalence classes. The set of contours that are organized into a form belong to an equivalence class that contains just that form; the form uniquely specifies the class. The other contours belong to an equivalence class that consists of all random spatial permutations; because the left-over contours do not make a coherent form, they could in principle be located anywhere in the image. Together these equivalence classes form an ensemble of images. Each member of the ensemble is composed of the form together with a selection from the set of random contours. Within this ensemble, the contours specifying form are invariant, whereas the other contours are not. In this way, spatial masking is placed on the same footing as data structures, with the only difference being that in the latter case the ensemble is generated by equivalent measurements, whereas in the former case the ensemble is generated by equivalent images.

The treatment that we have given here of two-part object structure expresses the tacit understanding that grounds classical theories of object structure. These theories are predicated on the notion of ensemble invariance of meaningful structure. This notion and the entailed conception of noise as coincidental make sense ecologically if we consider what is required in target finding. Often it is the case that noise is not an intrinsic component of the target and must be penetrated for object recognition. In natural environments, the camouflage that hides an animal is composed of whatever the animal happens to be next to, and this changes stochastically as a function of position. The animal is truly the invariant quantity in the ensemble of possible positions. In such circumstances, it makes sense to treat the underlying form as really being independent of the surrounding noise, and a visual system that attempted to do otherwise would be impaired. However, spatial noise need not be understood as comprising the variant component in an ensemble, and, in particular, this conception of noise is not appropriate for describing natural form.

Fractal conceptions of structure are not founded on a metatheory where the rough appearance of objects is secondary to their intrinsic form. In a fractal metatheory, roughness is not a surface property that coats ideal differentiable forms; it is not a layer to be penetrated. Natural contours are to be conceived as nothing but roughness. In a fractal description of contour, there is no platform, no underlying form. Bipartite object structure is replaced by a structural hierarchy. Fractal contours exist only as recursive structures that reiterate roughness throughout all scales.

The self-similarity that is key to fractal structure is a form of invariance. A class of fractals, which we will refer to here as *deterministic fractals*, are explicitly constructed to have the property that the global structure is reiterated on all scales; they have shape invariance under magnification. There are several approaches to constructing deterministic fractals, and all are algorithmic. One class of algorithms uses a base and generator. The fractal is recursively constructed by operating on the base with the generator, then operating on the result, and so on. An example of base-generator recursion is the Koch snowflake, shown in Figure 1 at the third level of recursion. The base of the snowflake is a single line segment, and the generator consists of removing the middle third from each straight line segment and replacing it with two segments that would form the sides of an equilateral triangle. A second class of algorithms uses iterated contractive mappings. These algorithms, known as iterated function systems (IFS; Barnsley, 1988a, 1988b), generate fractals as the attracting sets of dissipative dynamical systems. IFS have the generic property of dynamical systems that, even though they are deterministic, their solutions may be chaotic and therefore unpredictable. This feature gives IFS the ability to render a variety of natural shapes and textures using fractal descriptions. Note that, although it is possible to define the procedure for constructing a deterministic fractal, there is no formula for the contour that is generated in the limit of the completed infinity of iterations. Yet, by the nature of the construction, deterministic fractals are self-similar and invariant under magnification.

In addition to deterministic fractals, there are fractals that are intrinsically random. Random fractals are distinguished from the chaotic attractors of IFS by incorporating stochastic elements in their construction. The algorithms for random fractals are nondeterministic in the sense that random numbers are selected as part of the generation process. Terrestrial structures are typically modeled using random fractals (Mandelbrot, 1983; Pentland, 1986; Voss, 1985, 1988) because they naturally capture the roughness of landscape features. Random fractals generalize the notion of scaling invariance in two ways: (a) The invariants are statistical properties of the contours as opposed to specific pattern designs, and (b) self-similarity is replaced by a less restrictive invariance, namely, self-affinity.

Fractional Brownian Motion

A particular class of random fractals that have been used extensively in modeling natural contours are known as *fractional Brownian motions* (Feder, 1988; Saupe, 1988). Fractional

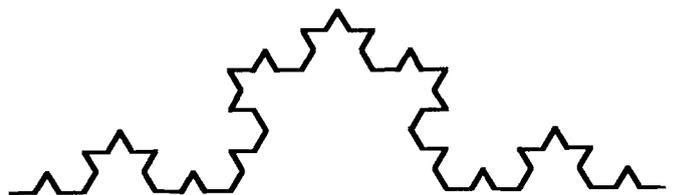


Figure 1. The Koch snowflake is an example of a recursive deterministic fractal. (It is depicted here at a recursive depth of three iterations.)

Brownian motion is a generalization of Brownian motion, and the simplest way to introduce this family is in terms of a generalized random walk. Random walks in statistical mechanics arise from considering the paths of particles as they are buffeted by collisions in a diffusion process (e.g., photons in radiative diffusion and molecules or atoms in diffusion driven by a concentration gradient). The erratic meandering of a diffusing particle is referred to as *Brownian motion*, and the path generated is a random walk with independent increments. To make contact with fractal contour, we consider random walks in one dimension. Consider, then, the time history of a particle's position as it diffuses in one dimension. An example of such a history is shown in Panel A of Figure 2. This graph, within the limits of plotting resolution, is a fractal contour.

In a classical random walk of the type that has been described, the particle position at any given time is highly correlated with where it has been in the past. The positional correlation decreases with increasing look-back time. The presence of an overall trend in Panel A, Figure 2 is the graphical evidence of correlation. Each step in a random walk is, however, statistically independent. A diffusing particle is as likely to be propelled forward as it is backward in any given collision. It is at this level of analysis that Brownian motion is generalized. Mathematically, it is possible to construct motions in which successive increments are negatively correlated (i.e., opposite sign) or positively correlated (i.e., same sign). Examples of these types of contours are shown respectively in Panels B and C. These motions are also fractals; they differ only in their fractal dimension. The amount of correlation that exists between successive increments is a parameter that can be smoothly adjusted and this leads to a continuous family of random fractals.

The type of invariance that fractional Brownian motions satisfy is revealed by considering the power spectra of their time histories. The power spectra of motions in this family is given by $P(f) = f^{-\beta}$, where f is the spatial frequency, and β is the power law exponent that determines the nature of the fractal.

For $\beta > 2$, successive increments are positively correlated, and for $\beta < 2$ they are negatively correlated. $\beta = 2$ generates contours that are classical random walks. The fractal dimension of fractional Brownian motions is given for $1 \leq \beta \leq 3$ by

$$D = E + \frac{3 - \beta}{2},$$

where D is the fractal dimension, and E is the topological dimension; $E = 1$ for curves, and $E = 2$ for surfaces. Computer modeling of landscapes (Voss, 1985, 1988) and ecological inventories (Keller et al., 1987) indicate that natural formations typically have power law exponents near $\beta = 2$.

Power laws have affine symmetry in the following sense. A transformation on a fractal contour of the type $f \rightarrow af$ can be offset by a transformation on the amplitude (i.e., the square root of the power spectrum) by $A(f) \rightarrow a^{-\beta/2}A(f)$. This type of invariance is distinguished from self-similarity in that the statistical structure of the contour is not invariant under magnification until the amplitude has been adjusted. An example will clarify this difference. Suppose that pictures are taken of a tree line with variable settings of a zoom lens as was done in the Keller et al. (1987) study. Zooming in has the effect of multiplying all the spatial frequencies by some constant number $a < 1$; structures that were small (i.e., high spatial frequency) now appear large (i.e., low spatial frequency). The transformation in apparent size can be eliminated, if the tree line is truly fractal, by a vertical stretching or compression. This is the amplitude transformation.

The trade-off between frequency and amplitude will not work if there is some nonfractal feature that appears in one of the photographs. Suppose that individual trees are resolved when the scene is magnified. This fine scale structure will not disappear by an adjustment of vertical scale, and it will be possible to discriminate the two photographs on this basis. Similarly, if the overall extent of the stand is revealed by zooming

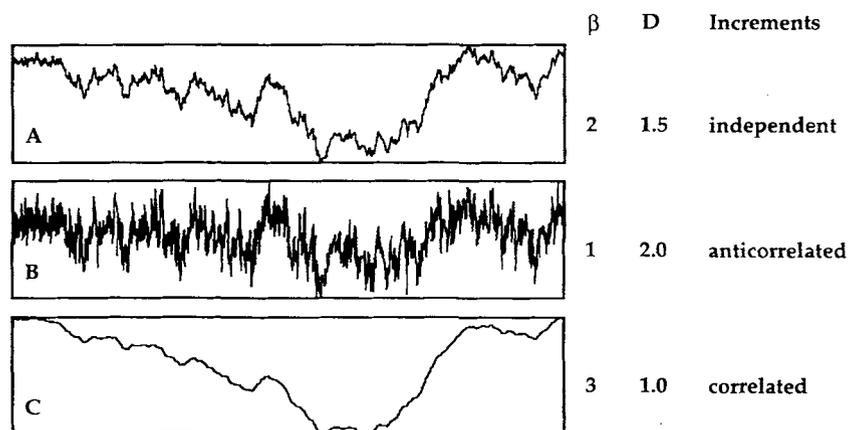


Figure 2. Three fractional Brownian motions. (Panel A shows a classical random walk composed of independent increments. Its power spectrum is a power law with exponent $\beta = 2$. It has a fractal dimension of 1.5. Panel B depicts a random walk with anticorrelated successive increments. Its power law exponent is unity, and contours of this type are referred to as $1/f$ noise. It has the maximum fractal dimension of 2. Panel C shows a random walk composed of correlated successive increments. It is quite smooth and has a steep power law with an exponent of three. It has a fractal dimension equal to its topological dimension.)

out, no transformation on the amplitude spectrum can eliminate this feature. In this sense, tree lines and all landscape forms are fractal only within certain scales of size.

Nonfractal features generally impose a signature in the power spectrum that is not self-affine. Figure 3 illustrates how this occurs. Panel A shows a power spectrum that is a pure power law in the $\log(\text{frequency})$, $\log(\text{power})$ plane. A global change of scale induces a multiplicative transformation on the frequency with the result that the entire spectrum shifts horizontally. This horizontal shift can be offset by a vertical shift—a multiplicative transformation on the amplitude or spectral power. The combined action of shifts in both amplitude and scale lays the original spectrum precisely down onto the transformed spectrum. This is self-affinity. Now suppose that there is some feature in the power spectrum that deviates from the power law. This is shown in Panel B as increased power at some frequency. Horizontal and vertical shifts move the feature to a new frequency, providing a basis for discrimination.

Figure 3 also clarifies what kind of objects natural fractals are. Theories of object recognition (see Biederman, 1987, for an example) generally treat the case in Panel B in which there is some feature that stands out. Spectral analysis of the shapes of everyday objects would reveal a multitude of nonfractal signatures that would be associated with the existence of the various parts that together compose the object. Self-affine contours, however, do not have parts in this sense. There is no sense in which a piece of the contour can be isolated and analyzed separately. All pieces of the contour are essentially equivalent in containing the same hierarchy of structure. What makes a contour a fractal is the way all the scales are linked together. What makes a contour fractal is that all scales embed and are themselves embedded in a hierarchy.

The idea that objects have parts is a metatheoretical notion that is allied to the notion of bipartite structure. The signal-noise distinction is, after all, fundamentally a separation into parts. Now, natural contours do not mathematically fit into this metatheory of structure by virtue of not having parts. They are not decomposable into anything more primitive. The underlying form of a fractional Brownian motion is, perhaps, best thought of negatively; it is something that does not decompose.

Experiments in the Discrimination of Rough Contour

Overview to the Studies

That the environment is teeming with fractal structure does not mean that people are sensitive to the information implicit

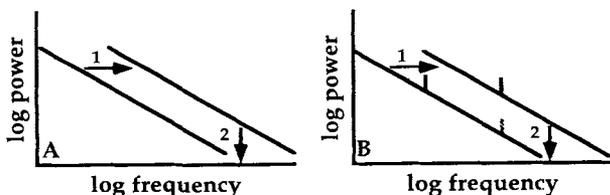


Figure 3. Double shift transformations illustrate why power spectra are self-affine. (Any deviation from a power law, as in Panel B, creates a feature in the power spectrum that is not invariant under a horizontal and vertical shift.)

in a hierarchy of structure. A visual system that operates in terms of a bipartite metatheory of structure may still be able to distinguish between fractal structures in certain regimes of correlation. From an ecological point of view, it makes little difference how information is processed so long as the animal is informed. Ramachadran (1990) has articulated this observation in terms of perception as being a “bag of tricks.” In the case of fractal structure, however, the nature of processing is of interest in itself. Fractals have a definite and delineated structure that permits a focused inquiry into the concordance between distal hierarchies and perceptual logic.

We have attempted to distill how fractional Brownian motions are perceived as exemplars of natural forms in three experiments. In Experiment 1, fractals were discriminated from other fractals. Discrimination sensitivity for fractional Brownian motions was assessed as a function of the power law exponent, β . In Experiments 2 and 3, subjects discriminated fractals from other types of rough contour that were not self-affine. The non-self-affine contours were parametrically related to fractional Brownian motions, allowing β to serve again as an independent variable.

The experimental work that we present here must be placed into context with the theoretical ideas that motivate this inquiry. First, we use stimuli that are abstract representations of natural form. Inferences from our discrimination studies about the perception of natural contour are allowable only to the extent that fractional Brownian motions cover the set of contours that occur in nature. Ecological studies (Burrough, 1981; Keller et al., 1987; Voss, 1985, 1988) provide evidence that this inference can be made. Second, the theoretical points that we wish to make are quite general, whereas the experimental work is quite specific. This incongruity is inherent to the task we have set for ourselves; it is necessary to conduct highly constrained experiments to collect data that are interpretable. What permits generalization from the empirical work to statements about process are the numerical simulations. The simulation technique that we present is extremely powerful, and it is by virtue of the clarity that simulation imposes that we are able to establish general claims concerning the processing of contour information.

Previous psychophysical studies in the perception of fractals have focused on establishing a correlation between some distal aspect of fractal structure and the perception of that structure. Early studies demonstrated that perceived roughness or complexity is correlated with the fractal dimension. Pentland (1986) showed that estimates of the roughness of 2D drawings of fractional Brownian surfaces were highly correlated ($r = .98$) with the distal roughness (i.e., fractal dimension) that the drawings represented. In related work, Cutting and Garvin (1987) found that complexity ratings of iterated geometric fractals were explained by both fractal dimension ($r = .68$) and generation characteristics such as recursive depth ($r = .86$) and the number of segments displayed ($r = .47$).

Knill, Field, and Kersten (1990) performed a correlational study that is relevant to the work described here. They presented subjects with 2D, rasterized images in a discrimination paradigm. The graininess of the images determined the roughness or fractal dimension. It was found that people’s ability to discriminate between fractal images was a function of how

rough the images were, and the sensitivity curves correlated quite well with the frequency with which surfaces of a given fractal dimension appear in nature. Most terrain surfaces have fractal dimension on the order of 2.5, and it was found that discrimination sensitivity for images with this dimension was maximal. Knill et al. also attempted to suggest how discrimination might be accomplished by human observers. Their model was based on band-pass filtration, which is similar in many respects to the models we develop later. However, it is clear from inspection of the relevant figures that Knill et al. were largely unsuccessful in characterizing the processes mediating discrimination; the theoretical curves do not bear even a faint resemblance to the empirical discrimination functions.

Stimulus Construction and Methodology

The three experiments described in this article have a common methodology and logic in the construction of the stimuli. In all of the experiments, subjects made a same-different judgment on the presentation of two contours. *Same* in this context meant that the two contours came from the same family. What constituted a family depended on the experiment. The families are shown heuristically in Figure 4 in terms of their spectra in the log(power), log(frequency) plane. In Experiment 1 the two families were both self-affine fractional Brownian motions (*fBms*) and were distinguished by the power law exponent (β). In Experiments 2 and 3, *fBms* were compared with hybrid contours that had a break in the power spectrum. Hybrids were formed by smoothly joining two power laws at the geometric mean of the frequency range. In Experiment 2, *fBms* were compared with hybrids that had a larger power law exponent for low spatial frequencies and the same exponent for high spatial frequencies. In Experiment 3, the construction of the hybrids was reversed: larger exponent for high spatial frequencies, the same exponent for low spatial frequencies. Examples of the contours and the sense in which they were paired in the respective experiments are shown in Figure 5.

Individual contours were created by assigning a random phase to each frequency component and computing the inverse Fourier transform. For the total range of excursion to be eliminated as a cue for discrimination, all contours were normalized to the same range. This normalization has the consequence of effectively randomizing the absolute power at any given frequency across contours within a given family. The only consis-

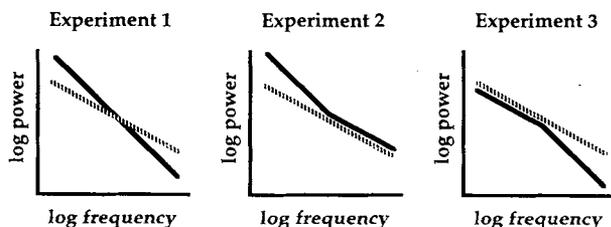


Figure 4. The stimuli used in Experiments 1, 2, and 3 are most clearly distinguished by their power spectra. (Fractals were discriminated in Experiment 1. In Experiments 2 and 3, fractals were discriminated from hybrids that were composed of two power laws.)

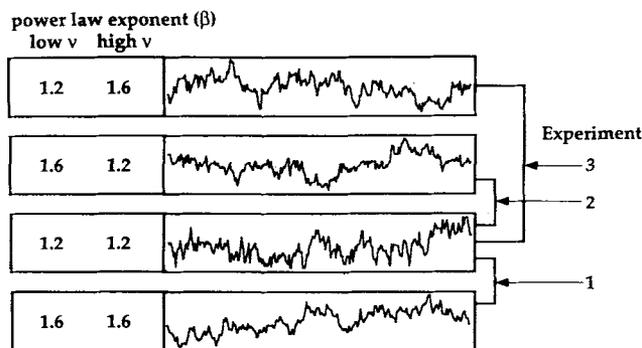


Figure 5. Exemplars of the stimuli used in discrimination Experiments 1, 2, and 3. (The lines connecting the contours illustrate the comparisons that distinguished the experiments.)

tent features within a family of normalized contours are the power law exponents.

In Experiment 1, 14 *fBm* families were created by spectral synthesis in the range $0 \leq \beta \leq 3.9$, where the k th family had an exponent $\beta_k = 0.3(k - 1)$, $k = 1, 2, 3, \dots, 14$. Each family consisted of 200 exemplars of *fBms* with power law exponent β , and each contour consisted of 256 random numbers. All noises were scaled to the same range. The noises subtended vertically 10° and horizontally 14.0° at a viewing distance of 49 cm. A block of trials consisted of showing all 200 exemplars from two adjacent β families. On a given trial a pair of contours would be displayed. Pairs were presented simultaneously and were spatially adjacent. Four contour pairings were used: (β_k, β_k) , (β_{k-1}, β_k) , (β_k, β_{k-1}) , and $(\beta_{k-1}, \beta_{k-1})$. The subject's task was to respond *same* if they thought that the pair came from the same β family and to respond *different* otherwise. There were 100 same pairs and 100 different pairs in each comparison block. Presentation of the pairs was random. An example of the three different types of pairings (two same types, one different type) is shown in Figure 6 for $\beta = 2.1$ and $\beta = 2.4$. The proportion of correct and incorrect different responses was used to compute the area under the receiver operator characteristic (ROC) curve. This

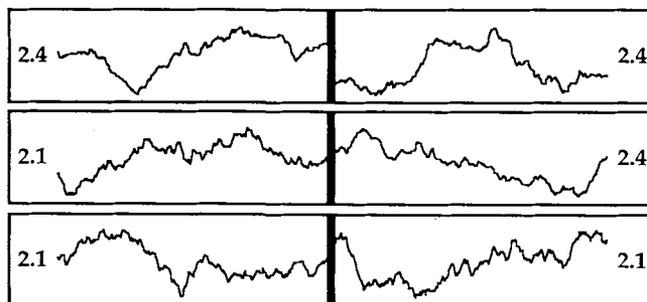


Figure 6. Examples of the three types of trials that occur in Experiment 1. (Depicted are fractional Brownian motions with $\beta = 2.1$ and $\beta = 2.4$. The $\beta = 2.1$ motions are slightly rougher and have greater point-to-point fluctuation than do the $\beta = 2.4$ motions. The correct response was *same* for trials illustrated by the top and bottom panels, and *different* for trials illustrated by the middle panel.)

area was computed by calculating d' from the number of hits and false alarms (treating a correct response of different as a hit) and then inverting the cumulative Gaussian (Falmagne, 1986). The areas under the ROC curve in each of the 13 family comparisons constitute the data for this experiment.

Pilot experiments in discriminating $fBms$ from contours with breaks in the power spectrum revealed that this was a more difficult task than discriminating two $fBms$. Recall that the power law exponents for these stimuli were equated on one side of the break point, f_{break} . To make the task in Experiments 2 and 3 more tractable, we increased the exponent separation from 0.3 to 0.4. Thus, in Experiment 2, there were 10 fBm families with power law exponents $\beta_k = 0.4(k - 1)$, $k = 1, 2, 3, \dots, 10$. The nonself-affine contours in this experiment had these exponents for frequencies $f > f_{break}$ and exponents $\beta_k + 0.4$ for $f < f_{break}$. The four contour pairings used in each block were of the form (β_k, β_k) , $(\beta_k, \beta_k + 0.4)$, $(\beta_k + 0.4, \beta_k)$, and $(\beta_k + 0.4, \beta_k + 0.4)$, where the pairings here refer only to the power law exponent for $f < f_{break}$; both contours had the same power law exponent, β_k , for $f > f_{break}$. In Experiment 3, the formal construction was the same except that the pairings now refer to frequencies $f > f_{break}$, and where the contours were equated on β_k for $f < f_{break}$. In all other respects, the design of the experiments and the presentation of the stimuli were the same.

The procedure for the three experiments consisted of a training period in which the subjects became familiar with the stimuli and the sense in which random contours could belong to the same family without appearing identical on a point-to-point basis. This task is intuitive, and subjects took only a few minutes of practice before they understood what was required of them. Feedback was given on each trial. The number of subjects was 8, 5, and 5 in Experiments 1, 2, and 3, respectively.

Discrimination Sensitivity for Rough Contour

In Figure 7 the results of the three experiments are shown. On the abscissa is the mean power law exponent for a family comparison. So, for example, if a block of trials in Experiment 1 compared $fBms$ with $\beta = 2.4$ and $\beta = 2.7$, the data point is plotted at $\beta = 2.55$. Sensitivity is measured on the ordinate as the area under the ROC curve. The discrimination profiles for all three experiments are inverted U-shaped curves with maxima in the region of power law exponent $1.5 < \beta < 2.5$. Other studies of fractal discrimination are consistent with the results from Experiment 1 (Knill et al., 1990). In particular, the finding that maximum discrimination sensitivity occurs near $\beta = 2$ appears to generalize across different methodologies and stimulus appearance.

We were concerned that the discrimination profiles may not have reflected observer's sensitivity to β , the power law exponent used in the construction of the $fBms$ and non-self-affine contours, but rather to some artifact in the way the contours were presented. For example, the power spectrum of each contour is determined not only by its exponent but also by a multiplicative constant, which sets the absolute amplitudes. To assess the role of this multiplicative constant, we repeated Experiment 1 with three new subjects using a larger stimulus display (roughly twice as large) that had vertical height subtending 19.6° at a viewing distance of 49 cm while maintaining the same

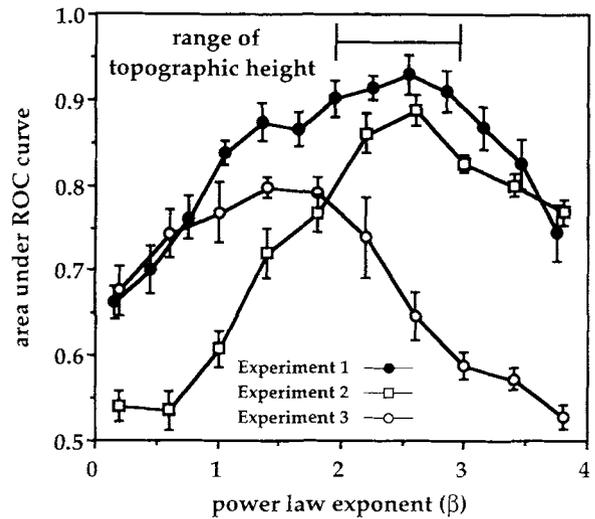


Figure 7. The discrimination data for Experiments 1, 2, and 3 as a function of the mean power law exponent (β) of the fractional Brownian motions in each family comparison. (The measure of discrimination used here is area under the receiver operator characteristic [ROC] curve. An area equal to 0.5 corresponds to chance guessing, whereas an area equal to 1.0 implies perfect discrimination. Also shown is the region in power law exponent that corresponds to fractals encountered in nature. Error bars depict the standard error.)

horizontal extent as in the original experiment. In this experiment all 13 comparison blocks were repeated twice for both large and small displays. The results of this study, averaged over subjects, are plotted in Figure 8. The two curves illustrate the

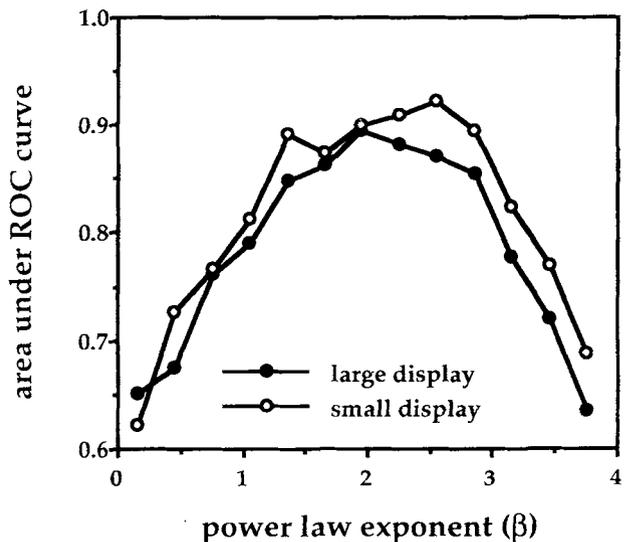


Figure 8. The mean discrimination curves for three subjects viewing fractals subtending vertically 10.0° and horizontally 14.0° (small) and subtending vertically 19.6° and horizontally 14.0° (large). (Stimulus size does not influence the shape of the sensitivity profile and has only a marginal influence on absolute discrimination. ROC = receiver operator characteristic.)

discrimination accuracy for large and small displays of fractals and are shown in the same format as Figure 7. The main effect of stimulus size was quite small in magnitude; the average difference was .038 but was significant, $F(1, 2) = 19.77, p < .047$. The size by exponent interaction was not significant, $F(12, 24) = 0.93, p < .54$. It is evident from these data that the shape of the sensitivity function is invariant under transformations of stimulus size and that size has only a marginal effect on absolute discriminability.

An issue that has arisen in the discrimination of fractal texture (Knill et al., 1990) is whether sensitivity is correlated with environmental frequency. As mentioned earlier, landscapes have fractal forms that are typically characterized by $\beta = 2$, random walks with independent increments (Burrough, 1981; Keller et al., 1987; Sayles & Thomas, 1978a, 1978b; Voss, 1985, 1988). It turns out that discrimination sensitivity is also maximal at power law exponents near $\beta = 2$. In Figure 7 we have also indicated the range in β characteristic of topography. This range should be compared only with the results of Experiment 1 in which only fractals were discriminated. Knill et al. interpreted the agreement between maxima in frequency of environmental occurrence and sensitivity as evidence that perception is efficient; resources are allocated so as to be able to distinguish those structures that are likely to be met in nature. As a statement limited to a summary of the data, it is true that perception is efficient. However, the implication is that perception is efficient because perception has been tutored by natural form. The data suggest such an interpretation and we have also been tempted by it (Gilden & Schmuckler, 1989). However, this agreement is at best correlational, and there may be other constraints that lead to U-shaped fractal discrimination functions that peak near $\beta = 2$. An alternative interpretation of this agreement is that the shape of the discrimination curve is not produced by environmental training but rather is a reflection of the particular logic that is used in discrimination. This logic may be quite general and subsume fractal discrimination as a specific case. The idea here is that perception may have its own protocols; these protocols exist for reasons unrelated to the occurrence of fractal structure, and it is primarily coincidence that leads to the agreement between environmental frequency and discrimination sensitivity. Note that a discrimination process that fails for extremely rough contours (i.e., small β) and extremely smooth contours (i.e., large β) will naturally generate an inverted U-shaped discrimination profile as a function of β . The correlation that exists is limited to the observation that the top of the inverted U is near the center of the β range where environmental fractals are encountered. Both the maximum of the discrimination profile and the range of fractal occurrence are ill defined and the stated correlation is a weak result. In what follows, we attempt to determine if there is, in fact, a logic of discrimination.

Theory of Rough Contour Discrimination

The primary goal of this article is to produce a definite theory of how rough contours are discriminated. There are several distinct problems that must be considered: How is rough contour organized in perception, what aspects of contour serve as a basis for discrimination, and what kind of metric associates

difference in fractal dimension with perceived difference. These sorts of questions require the construction of explicit models of discrimination. The models contemplated here consist of algorithmic instructions that serve as mathematical solutions to the problems of organization, discrimination criteria, and metric. In establishing criteria for agreement with the data to be simulated, we adopt a conservative approach and shall require models to actually reproduce the exact shapes and amplitudes of the discrimination profiles. The task of replicating data through simulation is much more demanding than establishing correlation with data (e.g., as in Knill et al., 1990). As we show later, replication places severe constraints on admissible algorithms.

A central issue in these studies is whether people perceive the fractal properties of fractals. There are a number of ways of framing this issue. Is the roughness of fractal contour perceived in terms of a scaling hierarchy? Is self-affinity a perceptible transformational invariant? Is discrimination based on an awareness of fractal dimension? We know from earlier studies (Cutting & Garvin, 1987; Pentland, 1986) that people can rank order random fractals in terms of their fractal dimension, and this suggests that fractal structure may be manifest. This proposition may be tested by the construction of explicit models of discrimination. Now there is no sampling distribution of models; models come from insight, and it is necessary to develop some ideas about how fractals are perceived and what aspects of fractal contour are used in discrimination. The idea that fractals are perceived in terms of a hierarchy of structure and that discrimination is based on fractal dimension is only one model.

Ways of Looking at Rough Contour

We begin by simply looking at some fractals. Consider the fractals illustrated in Figure 2. These contours were constructed to be self-affine up to the limits of resolution permitted by the graphical interface and the finiteness of the random number sequence. Even within these limitations, the contours are self-affine over several orders of magnitude in resolvable scale. There is nothing in principle that prevents some portion of the hierarchy from being perceived. Yet the dominant impression that one derives from inspecting these contours, an impression congruent with the reports of subjects, is that they are composed of two parts: a smooth trend that supports rough oscillations. This impression is tacitly founded on a signal-noise bipartite decomposition, and, though informally derived, it is not trivial. It suggests that visual understanding of rough contour is based on a perceptual splitting. Within this way of looking at contour, fractals are perceptually differentiated by the variations in the smooth and rough components.

Rough contours may also be analyzed in terms of a complete decomposition into individual increments. There are several measures defined on the increments that distinguish fractals from one another. One measure that has been mentioned previously as a formal index of contour structure is the interincrement correlation. Contours with positive increment correlation are smoother than contours with negative correlation. A related property of the increments is the width of their distribution. The distributions may be inferred from Figure 2; one of the features that distinguishes $\beta = 1$ from $\beta = 3$ *fBms* is that the

former has increments of all sizes, whereas the latter has increments only within a restricted range of size.

These ideas can serve as a basis for formulating models of discrimination. This task is best served by considering the discrimination problem in the respective experiments using contours widely separated in power law exponent, β . In Figure 9 we show how self-affine and non-self-affine contours constructed from $\beta = 1$ and $\beta = 2$ power laws appear when decomposed into smooth and rough components. The contours are distinguished in this figure by the slope of the power law at low and high spatial frequencies, ν . *fBms* (Columns 1 and 2) have the same power law at all frequencies, whereas the hybrids (Columns 3 and 4) are formed by joining two power laws at the geometric mean frequency. The lines connecting the boxes illustrate the comparisons that were made in the respective experiments. The top contours are stimuli representative of the types used in the experiments, the middle contours are smooth approximations, and the bottom contours are the rough components computed as contour minus smooth.

The smooth-rough decomposition displayed in Figure 9 was computed using the smoothing algorithm suggested by Press, Flannery, Teukolsky, and Vetterling (1986), which applies a Gaussian window in the Fourier transform domain. Other smoothing algorithms could have been used, but they are all essentially identical to band-pass filtering in the Fourier domain. For example, smoothing may be achieved through convolution in which a set of adjacent points are averaged with variable weights. Convolution in the spatial domain is equivalent to multiplying by a window in the Fourier domain. If the

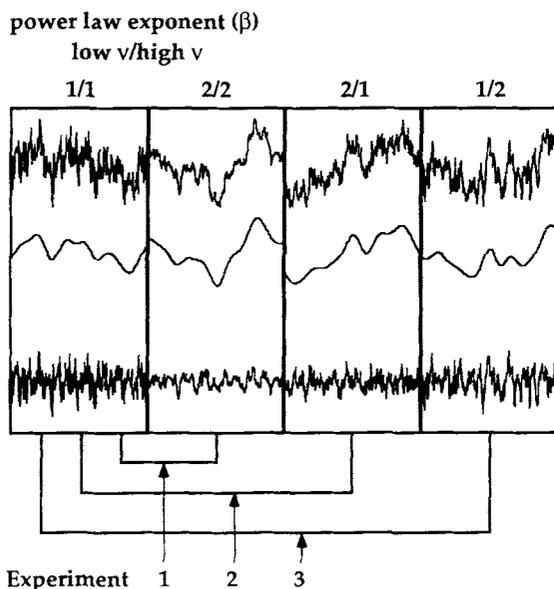


Figure 9. Examples of self-affine contours and hybrid contours composed of two power laws. (Also shown are smooth and rough extractions. Numbers across the top refer to the power law exponent at low and high spatial frequencies. The lines connecting the panels illustrate the sorts of comparisons that were made in Experiments 1, 2, and 3. We depict contours with a wide spacing in power law exponent to highlight the information that could be used in discrimination.)

weights decrease with increasing distance from the center of the smoothing window, then the window in the Fourier domain can be approximated by a Gaussian.

The Press et al. (1986) algorithm accepts one parameter that determines how faithfully the original contour is reproduced by the smooth approximation. This parameter is roughly the number of adjacent points that are averaged together and is formally the full width of the convolving Gaussian. The full width enters as a degree of freedom in all theoretical models that are based on smooth-rough decomposition; models based on increment analysis do not have any free parameters. Full widths of approximately 25 points produced the best fits to data, and it is this value that is illustrated in Figure 9. In practice we found that simulated discrimination sensitivities were slowly varying functions of the Gaussian full width. Perturbation of the full width by as much as 25% produced little effect in simulated discrimination. In the simulations presented later the exact value of the full width parameter plays a minor role. Models that are rejected did not provide adequate fits to data for any value of the full width.

We consider the problem of feature extraction separately in the three experiments. The task of identifying information that is both available and accessible for discriminating *fBms*, the task in Experiment 1, is informed by comparing the fractal $\beta = 1$ and $\beta = 2$ contours depicted in Columns 1 and 2 of Figure 9. These contours are easily distinguished on the basis of several different but related characteristics. Consider first the increment distributions. The $\beta = 2$ distribution contains mostly small increments, whereas the $\beta = 1$ distribution is more uniform. This relationship holds in general. The standard deviation of the increment distribution decreases monotonically with power law exponent as illustrated in Panel A of Figure 10. For contours that are normalized to the same total vertical range, small increments are associated with smooth profiles and positive increment correlation, whereas large increments are associated with spiky profiles and negative correlation. The linkage between distribution width and increment correlation is illustrated in Panel B by the increasing monotonic trend of increment correlation with β . Monotonicity in both functions suggests that either the distribution width or the two-point increment correlation may provide an adequate measure for the purposes of discrimination.

There is further information in the rough and smooth extractions from a bipartite decomposition. Most salient is the differences in range. Treating range as an operator on the contour, we have two manifest relations:

$$\text{range}(\text{smooth}, \beta = 1) < \text{range}(\text{smooth}, \beta = 2),$$

$$\text{range}(\text{rough}, \beta = 1) > \text{range}(\text{rough}, \beta = 2).$$

These inequalities express an important and perceptually transparent property of fractal contour; overall trend increasingly dominates point-to-point fluctuation with increasing β . Either or both of the inequalities may be capitalized on in discrimination. The general relationship between the respective ranges and β is illustrated in Panel A of Figure 11.

There is a second operator that may be defined on the smooth and rough extractions that is motivated by a signal processing approach to the measurement of fractal parameters

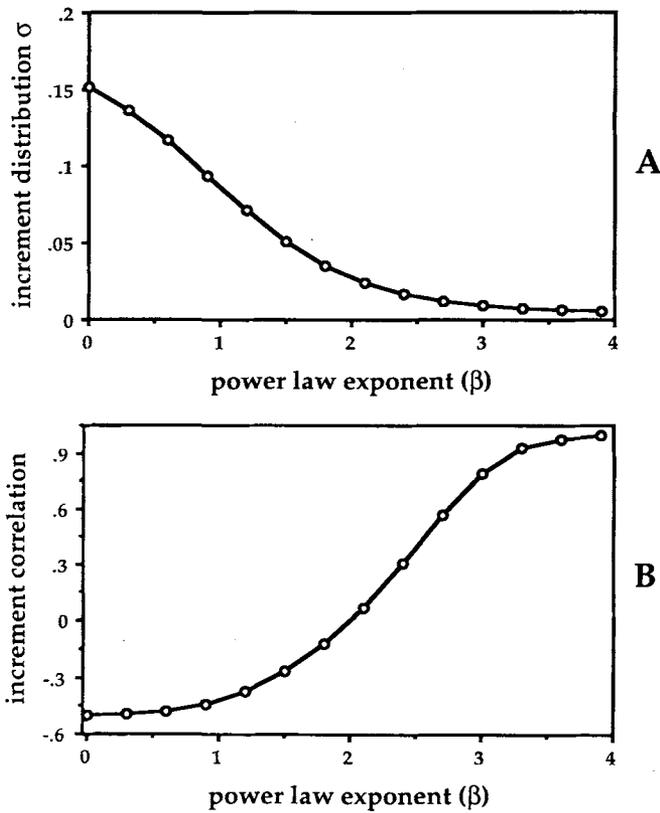


Figure 10. The relation between increment measures of fractal contour variation and power law exponent, β . (Panel A shows the relationship between β and the standard deviation of the distribution of increments. Panel B shows the relationship between β and the correlation between successive increments.)

(Pentland, 1986). This approach uses several formal constructions based on Fourier analysis to separate out the power spectrum. The contour is first band-pass filtered at all scales. The outputs of these filters are squared and spatially averaged to yield estimates of the local band-pass spectral power (i.e., Parseval's Theorem). A regression of log power versus log spatial frequency (i.e., spatial scale) yields a slope that is a direct estimate of the power law exponent, β , of the fBm . The fractal dimension is also estimated by this construction because it is related to β by a linear transformation for $1 \leq \beta \leq 3$.

This procedure requires only slight modification to define an operator that is less oriented to machine vision and is more psychologically plausible. Band-pass filtering will be presumed to be coarse and limited to smooth-rough decomposition. Coarse filtering removes the need for a regression analysis. The modified procedure also consists of three steps. First, the fractal contour is decomposed into smooth and rough components by low- and high-pass filtration, respectively. The Fourier power in each of the two components is then estimated from Parseval's Theorem. Panel B of Figure 11 shows how the rough and smooth amplitude spectra vary with β . The power law exponent, β , is itself estimated from the difference between log power(smooth) and log power(rough). Over a wide range of β ,

this difference, denoted as $\Delta \log$ band-pass power, is linearly related to β . This is shown in Panel C of Figure 11.

Figures 10 and 11 make it clear that there are a number of functions that may be defined on fractals that are monotonic with the power law exponent and the fractal dimension. Thus, the result that people can rank order fractals in terms of their fractal dimension does not have a unique interpretation. Any of the functions that we have discussed would be adequate to this task. In this article we argue that Monte Carlo simulation of discrimination data may help decide the empirical matter of how contours are decomposed for analysis and what contour features are perceptually extracted.

In Experiments 2 and 3 fractals were discriminated from hybrid contours that were constructed from joining two distinct power laws at the geometric mean frequency. The hybrid

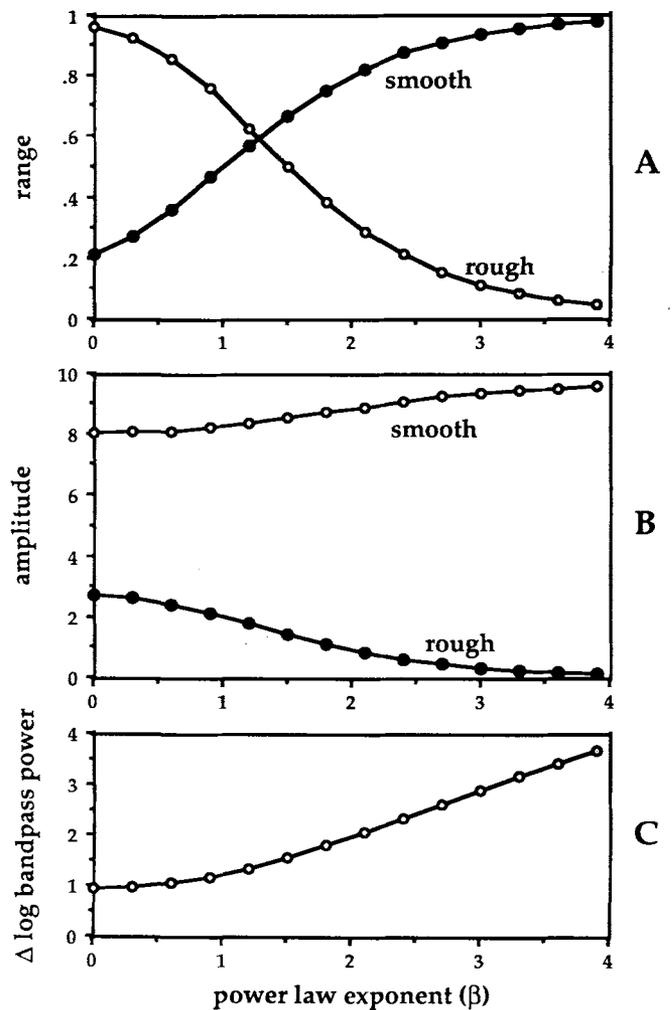


Figure 11. Measures of range and spectral power for the smooth and rough extractions from fractal contour as a function of power law exponent, β . (Panel A shows that the range in the smooth and rough extractions are monotonically, but not linearly, related to β . Panel B shows a similar monotonic relation between the band-pass amplitude and β in the respective extractions. Panel C shows that an estimate of β , $\Delta \log$ band-pass power, is linearly related to β over most of the domain.)

contours have different increment and range properties than do the fractals, and the issue of discrimination in these experiments must be dealt with separately.

The contour discriminations that were required in Experiment 2 are illustrated by comparing Columns 1 and 3 of contours in Figure 9. In this experiment, the hybrids and the fractals had the same power law at high spatial frequencies. The equation of β at high spatial frequency has discernible implications for comparisons on the range of the smooth component. A ($\beta = 2$ low ν , $\beta = 1$ high ν) hybrid will appear to be more rough than a $\beta = 2$ fBm . Consequently, the smooth component of the hybrid has a reduced range compared with the smooth component of the $\beta = 2$ fBm ; that is, it approximates the smooth range of the $\beta = 1$ fBm . This makes the range of the smooth component less useful as a point of comparison for discriminating ($\beta = 2$ low ν , $\beta = 1$ high ν) hybrids from $\beta = 1$ $fBms$. Decomposition-based discrimination between fractal and hybrid in Experiment 2 may be reduced to analysis of the rough extraction alone.

The equation of range in the two components is completed in Experiment 3 in which the hybrids had the same power law at low spatial frequencies. Compare the smooth and rough extractions in Columns 1 and 4 of contour in Figure 9. There is virtually no difference in range within the respective components, and these contours are much more widely separated in power law exponent than was the case in the actual experiment. Distinguishing fractal from hybrid in this experiment may require an explicit analysis of increments. The sufficiency of increment information is demonstrated in Figure 12 in which the increment distributions of the $\beta = 1$ fBm and the ($\beta = 1$ low ν , $\beta = 2$ high ν) hybrid are compared. The hybrid contour has fewer large increments than does the fBm contour, and so its distribution of increments has a smaller standard deviation. Note that the range of the distribution does not provide useful information as both contours have the full range of increment size represented.

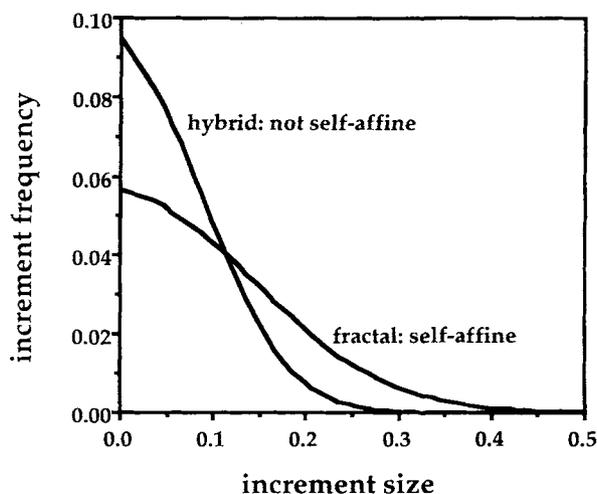


Figure 12. The distribution of unsigned magnitudes of successive increments for a $\beta = 1$ fractal and a hybrid with $\beta = 1$ at low spatial frequencies and $\beta = 2$ at high spatial frequencies. (The width of these distributions may be used for discrimination.)

The difference is only in the probability with which various increments are encountered.

Implications of Bipartite Decomposition

The decomposition of fractals into smooth and rough components does not in itself imply that fractal structure is not perceived or used in discrimination. Of critical importance is how the two components are used in articulating structure. Decomposition may be used in a coarse representation of the hierarchy of structure, or it may be a first step in a complete dismantling of the fractal where the two components are split for the purposes of separate analyses. In the first case, decomposition is interpreted as initiating an analysis of contour structure that preserves the relations among scales that distinguishes fractals. In the second case, decomposition is viewed as a visual heuristic that is insensitive to these relations and expresses a logic that has more to do with how people are prepared to see the world and less to do with the structure of that which is available to be seen.

The key aspect of the signal processing approach to fractal dimension estimation is that decomposition is used only to resolve scales and is not preparatory to an analysis of the components as separate kinds of structure. This distinction is critical because it bears on how information in the two components is used. Regardless of what features in the two components are extracted, once the fractal has been split there is a decision problem that must be addressed: Are the two components to be treated as being related in some way, or are they to be treated as unrelated and separate entities? From the point of view of fractal structure, the two components are related; they form a nested hierarchy. However, there is another point of view, which states that decomposition of the contour makes two things: a smooth thing and a rough thing. This is the signal-noise point of view. Signals and noise always appear in conjunction, but they are causally unrelated and are to be regarded as independent channels of information. In this second case, the decomposition results in a decision theory in which each component is treated in isolation.

The difference between these two views of structure can be clarified by formalizing the way in which parts of an object may be used to establish difference. Let A and B be the result of bipartite decomposition of a fractal, and let a fractal be denoted by $O_i = (A_i, B_i)$. The two decision theories are distinguished by how the A_i and B_i are grouped in discrimination. If the two components are regarded as having some relation to each other (e.g., nesting), and this relation is important in discrimination, then the perceived difference between two fractals will have the following form:

$$\text{perceived difference}(O_i, O_j) = \text{Difference}[F(A_i, B_i), F(A_j, B_j)],$$

where F is some procedure for characterizing the fractal on the basis of features in both components, and Difference is some measure of psychological distance along a continuum induced by this characterization. In the specific case of fractal parameter estimation, F computes the difference of the logarithm of the power in the two components. The computation of this difference effectively recreates the hierarchy of nested struc-

ture. The important point is that the two components are kept together to serve as arguments within a single function.

Consider now the case in which the two components are not regarded as having any particular relation to each other. Here, the two components are regarded as forming two separate types of information. In this case, discrimination will have a different form:

$$\text{perceived difference } (O_i, O_j) \\ = \text{Difference } [F(A_i, A_j), G(B_i, B_j)],$$

where now F and G are feature analyses that are applied to contours within a given type, and Difference is some measure of psychological difference in a two-dimensional space of types. Explicit models in which F and G are the same functions and are interpreted either as spectral power or range are constructed in the following section.

In this second way of looking at structure, decomposition produces separate categories; it is not just an operation for resolution of scale. This distinction is critical for how we interpret what people see when they look at fractal contour. An observer that computes fractal parameters sees a relation between the components and discriminates on the basis of that relation. Such an observer could be said to perceive fractal structure. An observer that discriminates on the basis of the separate components has produced a category distinction that violates the structural integrity of the fractal. This observer provides the sense in which fractals are not perceived in terms of the structural properties that define them. One goal of the simulations is to clarify whether people perceive the characteristics that define fractals or whether they make implicit category distinctions that are incompatible with fractal structure.

Implications of Increment Decomposition

An analysis of rough fractal contour in terms of its increments leads to several measures that are monotonically related to the power law exponent, β , and so can be used as features for the purpose of discrimination. The one-to-one relationships depicted in Figure 10 establish that either correlation or distribution width may be used to estimate fractal parameters. The observation that these measures may be used to estimate fractal parameters might suggest that were people to discriminate on the basis of the measures, they would be perceiving fractal structure. There are several arguments against this inference. First, these measures require that the contour be decimated. The self-affine property of fractals, that they contain a nested hierarchy of structure, is lost in this procedure of deconstruction. Second, there are many functions that can be defined on fractal contour that are monotonically related to β . For example, both range and spectral power defined on either the smooth or rough components can be used as estimates of β . Now the entire class of functions that permit monotonic mappings into β for fractal contours can be also used to analyze nonfractal contours. These operators are general tools for taking apart contour and provide measures of structure regardless of the domain of application. It is questionable to suppose that a procedure that may be applied universally has a special interpretation when applied to fractals.

A final point has been raised by Westheimer (1991) in a study of fractal border discrimination. Westheimer showed that people can easily discriminate between contours that have the same fractal dimension when the increments are chosen from different statistical distributions. This result implies that fractal dimension does not uniquely specify the perceived attributes of random fractal contour. The features that people use to discriminate rough contour are not completely captured by properties of the power spectrum, such as fractal dimension. The phase spectrum plays an important role here as it does generally in object recognition (Piotrowski & Campbell, 1982).

A second issue is whether these measures are perceptually accessible. In one sense, inquiring into the perceptual status of the increment correlation is idle because the increment correlation is completely confounded with the manifest appearance of the contour. However, it is not idle to inquire whether the increments of a $\beta = 2$ contour appear to be independent with zero correlation or whether the increment correlation of $\beta = 1.5$ and $\beta = 2.5$ contours can be distinguished by sign. In this sense, it is not clear that correlation is itself a perceptual property of contour. On the other hand, the width of the increment distribution emerges as a contour property that is perceptually penetrable. Breaking up contour in terms of the distribution of increments does not require a perceptual analysis that goes beyond seeing (a) mostly little jumps or (b) both big and little jumps. The width of a distribution does not share the subtlety of correlation in terms of reliance on sign; it is an unsigned magnitude. The issue of accessibility will ultimately be decided by comparing simulated increment decomposition observers with the discrimination data. The critical test bed for this comparison is Experiment 3 in which a bipartite decomposition does not yield distinguishable ranges and in which an increment analysis appears to be forced.

Models of Rough Contour Discrimination

In preceding sections we have developed a number of ways of thinking about the structure of rough contour and have identified several mathematical functions that could be used in explicit algorithms for modeling discrimination. In this section, we formalize this procedure and spell out in detail how discrimination is represented in computer algorithms.

Perceptual organization in terms of increment decomposition leads to two discrimination models: one based on correlation and the other based on the width of the distribution of increments. An algorithm that discriminates in terms of increment correlation is formed by computing the correlation of increments for the respective contours and then forming a correlation contrast:

$$\text{corr contrast}(\text{contour}_1, \text{contour}_2) = \frac{|\text{corr}(\text{inc}_1) - \text{corr}(\text{inc}_2)|}{|\text{corr}(\text{inc}_1) + \text{corr}(\text{inc}_2)|}$$

An algorithm that discriminates on the basis of increment distributions is formulated by computing the width of the increment distributions, $\sigma(\text{dist})$, for each contour and then forming a standard deviation contrast:

$$\sigma \text{ contrast}(\text{contour}_1, \text{contour}_2) = \frac{|\sigma(\text{dist}_1) - \sigma(\text{dist}_2)|}{\sigma(\text{dist}_1) + \sigma(\text{dist}_2)}$$

Discrimination is based on the magnitude of the real number associated with the σ contrast or corr contrast functions. For a

given threshold of discrimination, λ , the two contours are perceived to be in the same family if contrast $< \lambda$, and are perceived to be in different families otherwise. Discrimination based on either of these contrasts has no free parameters. Simulations based on these contrasts are applied to data from all three experiments.

Models of fractal discrimination (i.e., Experiment 1) based on bipartite decomposition fill the cells of a 2×2 design, feature extracted (i.e., range or spectral power) crossed with contrast procedure (i.e., within-component contrast or between-fractal contrast). A model that discriminates on the basis of estimates of power law exponent first forms the two quantities

$$\Delta \log \text{power}_1 = \log \text{power}(\text{smooth}, \beta_1) - \log \text{power}(\text{rough}, \beta_1),$$

$$\Delta \log \text{power}_2 = \log \text{power}(\text{smooth}, \beta_2) - \log \text{power}(\text{rough}, \beta_2),$$

where the subscripts refer to the two contours to be discriminated. $\Delta \log \text{power}_1$ and $\Delta \log \text{power}_2$ are estimates of β_1 and β_2 , respectively. A measure of the difference between the contours is formed by computing a between-fractal contrast of the two slope estimates:

$$\text{slope contrast} = \frac{|\Delta_1 - \Delta_2|}{|\Delta_1 + \Delta_2|},$$

or by computing a slope difference $= |\Delta_1 - \Delta_2|$. In practice, these two different estimates of difference lead to virtually identical discrimination sensitivities.

Spectral power can be used for other purposes than computation of fractal parameters. Power may be used to form contrasts defined on the separate components; that is, where the two components are treated as forming two distinct and independent sources of information. In this case, spectral power is not used to estimate β but serves only as a feature that might be useful in characterizing contour. This leads to a model based on within-component contrasts:

power contrast(smooth)

$$= \frac{|\log \text{power}(\text{smooth}, \beta_1) - \log \text{power}(\text{smooth}, \beta_2)|}{\log \text{power}(\text{smooth}, \beta_1) + \log \text{power}(\text{smooth}, \beta_2)},$$

power contrast(rough)

$$= \frac{|\log \text{power}(\text{rough}, \beta_1) - \log \text{power}(\text{rough}, \beta_2)|}{\log \text{power}(\text{rough}, \beta_1) + \log \text{power}(\text{rough}, \beta_2)}.$$

When the smooth and rough extractions are treated as separate dimensions of comparison, the respective contrasts must be combined in some way to provide a unitary judgment of same or different. A procedure for combining separate measures of difference into a single score entails the use of the Minkowski metric:

$$f(x, y) = [x^r + y^r]^{1/r}.$$

In practice, we have used the city block metric, $r = 1$, but the simulations gave essentially the same results for the Euclidean metric, $r = 2$. The function, f , effectively maps the two contours into a single real number, the total difference score. The difference score serves analytically as a basis for discrimination. For a given threshold of discrimination, λ , the two contours are per-

ceived to be in the same family if $f < \lambda$, and are perceived to be in different families otherwise.

Discrimination models for range can be constructed analogously to models based on spectral power. Again, we distinguish between a decision procedure that treats the fractal as the fundamental unit of comparison and one in which the decomposed parts form the units of comparison. In the former case, we define an operator similar to $\Delta \log$ power:

$$\Delta \text{range}_1 = \text{range}(\text{smooth}, \beta_1) - \text{range}(\text{rough}, \beta_1),$$

$$\Delta \text{range}_2 = \text{range}(\text{smooth}, \beta_2) - \text{range}(\text{rough}, \beta_2).$$

The difference metric is formed by computing a range contrast between fractals as was done for spectral power above. Alternatively, we can compute contrasts within the separate components as

range contrast(smooth)

$$= \frac{|\text{range}(\text{smooth}, \beta_1) - \text{range}(\text{smooth}, \beta_2)|}{\text{range}(\text{smooth}, \beta_1) + \text{range}(\text{smooth}, \beta_2)},$$

range contrast(rough)

$$= \frac{|\text{range}(\text{rough}, \beta_1) - \text{range}(\text{rough}, \beta_2)|}{\text{range}(\text{rough}, \beta_1) + \text{range}(\text{rough}, \beta_2)}.$$

Again, the two contrasts are combined within a city block metric to provide a total difference score for the contours under comparison. This last model is also applied to the fractal-hybrid discrimination data.

One additional model that arises in consideration of the comparisons made in Experiment 2 uses only the range of the rough component. This leads to the following contrast:

contrast(rough)

$$= \frac{|\text{range}(\text{rough}, fBm) - \text{range}(\text{rough}, \text{hybrid})|}{\text{range}(\text{rough}, fBm) + \text{range}(\text{rough}, \text{hybrid})}.$$

This model is also applied to the data from Experiments 1 and 3.

As fBm contours are characterized by a power-law power spectrum, discrimination data from Experiment 1 are well suited to an analysis that is based on band-pass power estimation. The discriminations in Experiments 2 and 3, however, required comparisons between $fBms$ and non-self-affine contours. The hybrids in these experiments were formed by smoothly joining two power-law power spectra at the geometric mean frequency and so are examples of multifractals; they are self-affine over restricted intervals in frequency. Now the signal processing approach for estimating β that has been outlined for fractals can be generalized for multifractals. The complication that arises for hybrids is that the separate frequency intervals within which the contour is self-affine must be identified. This identification requires band-pass filtering with an array of filters that cover a wide range of scales and an ability to resolve breaks in the power law in the Fourier domain. The types of filters that are associated with visual receptive fields are not relevant here; these contours are not formed by fluctuations in brightness. The type of Fourier analysis that is contemplated for decomposition of line contour is purely formal and

has not been given a physiological basis. Recognizing that this analysis is possible in principle, the problem of interval identification seems patently psychologically intractable. For these reasons we do not attempt to construct models based on spectral power for fractal-hybrid discrimination data. As we demonstrate in the following section on simulated discrimination, even for the simplest case in which fractals were compared, spectral power models do not generate discrimination functions that bear any resemblance to human performance.

Construction of Algorithmic Observers

A discrimination strategy, when expressed as a formal algorithm, is a mapping defined on two contours into the real numbers. For simplicity, we refer to the strategies collectively as some function F that operates as $F(\text{contour}_1, \text{contour}_2) = X$. An algorithmic observer is defined by the following procedure.

1. Construct the stimuli that would be used in a block of trials. These stimuli are made exactly as described in the procedure sections of the respective experiments. Thus, these blocks consist of 200 pairs of contours, 100 from the same family and 100 from different families. There were 13 blocks in simulations of Experiment 1 and 10 blocks in simulations of Experiments 2 and 3.

2. For each block of trials compute $F(\text{contour}_1, \text{contour}_2) = X_i, i = 1, 2, \dots, 200$.

3. Compute the minimum and maximum values of $\{X_i\}$. These values serve to delimit the range of possible thresholds, λ .

4. Construct a uniform sequence λ_j on the interval $[\min\{X_i\}, \max\{X_i\}]$. This sequence discretely resolves the range of thresholds.

5. Count the hits and false alarms: If $X_i > \lambda_j$, and the two contours were from different families, then the discrimination is a hit. Alternatively, if $X_i > \lambda_j$, and the two contours were from the same family, then the discrimination is a false alarm. Compute the fraction of hits and false alarms on the entire λ grid. These fractions generate the entire ROC curve.

6. Integrate under the generated curve in the (hit, false alarm) plane to compute the area under the ROC curve. This is the algorithm's discrimination performance for that block of trials.

This procedure was performed for all blocks in all three experiments using subsets of the various strategies based on spectral power, range, increment correlation, and width of the increment distributions. In each simulation, Steps 1–6 were repeated 10 times to reduce the noise associated with sampling error.

Simulated Discrimination of Rough Contour

The simulations of fractal contour discrimination are of primary importance because they allow us to distinguish relevant contour features and to place constraints on the type of decision process that occurs in human judgment of rough contour. There are two issues that we address at the outset. First, is spectral power a better predictor of human discrimination performance than range? Second, are the smooth and rough components of a given contour related to each other in some way, or are they compared within types across contours? These two

ways of treating the components have been formalized earlier and are now evaluated on the basis of the simulations. Figure 13 shows a 2×2 comparison of simulations of data in Experiment 1. The two variables are comparison feature and comparison type. The two features that are isolated in these simulations are range and the logarithm of spectral power. These features are used in two ways, depending on the structure of the decision theory. Panels A and C show results for simulations in which (a) the smooth and rough components are contrasted separately and (b) the separate contrasts are combined in a city block metric to give a total difference score. Panels B and D show results for simulations in which the two components are combined within a single function to provide a unitary measure of fractal contour. In these latter simulations, only one contrast is needed to yield an index of difference. We wish to emphasize that in all simulations that are presented in this article, no attempt has been made to scale the algorithmic sensitivities. No slope or intercept has been fit in a regression. The algorithms are simulating the absolute, unadjusted values produced by human subjects.

Panel A shows the results for a model that is based on informal observations of fractal contour: that range is a salient feature of contour and that comparing fractal contours involves comparing the ranges of the separate smooth and rough components. It is evident that this model provides an excellent fit to human data across the entire range of power law exponent. Panel D of Figure 13 illustrates the discrimination sensitivity of a model that is motivated by current signal processing approaches to fractal recognition. This model estimates the power law exponent, β , directly as $\Delta \log$ band-pass power and should perform optimally in regimes in which fractals are compared with each other, that is, when each contour is characterized by a single power-law power spectrum, as in Experiment 1. This observer is not as sensitive as the human subjects except when discriminating very smooth contours at large values of β . This result is seemingly paradoxical because an algorithm that discriminates on the basis of estimated power law exponent should be quite adept at distinguishing contours that actually do differ in the single respect of power law exponent. To resolve this paradox, it is helpful to consider the distributions of the quantities that serve as the basis of discrimination in the respective models.

A semianalytic theory of the algorithmic observers can be developed analogous to Thurstone's (1927) derivation of the law of comparative judgment. In this theory, discrimination sensitivity is proportional to the distance between the feature distributions of neighboring β families. These distributions are naturally generated by variation among the exemplars within each family. In practice, the distribution of a feature is computed by taking a large sample of exemplars from each family and then computing the numerical value of that feature for each exemplar. The distance between adjacent distributions is expressed as

$$d'(\beta_i, \beta_j) = \frac{\mu_{\beta_i} - \mu_{\beta_j}}{\sqrt{\sigma_{\beta_i}^2 + \sigma_{\beta_j}^2}},$$

where β_i and β_j are exponents for adjacent families and where μ and σ are the mean and standard deviation of the feature

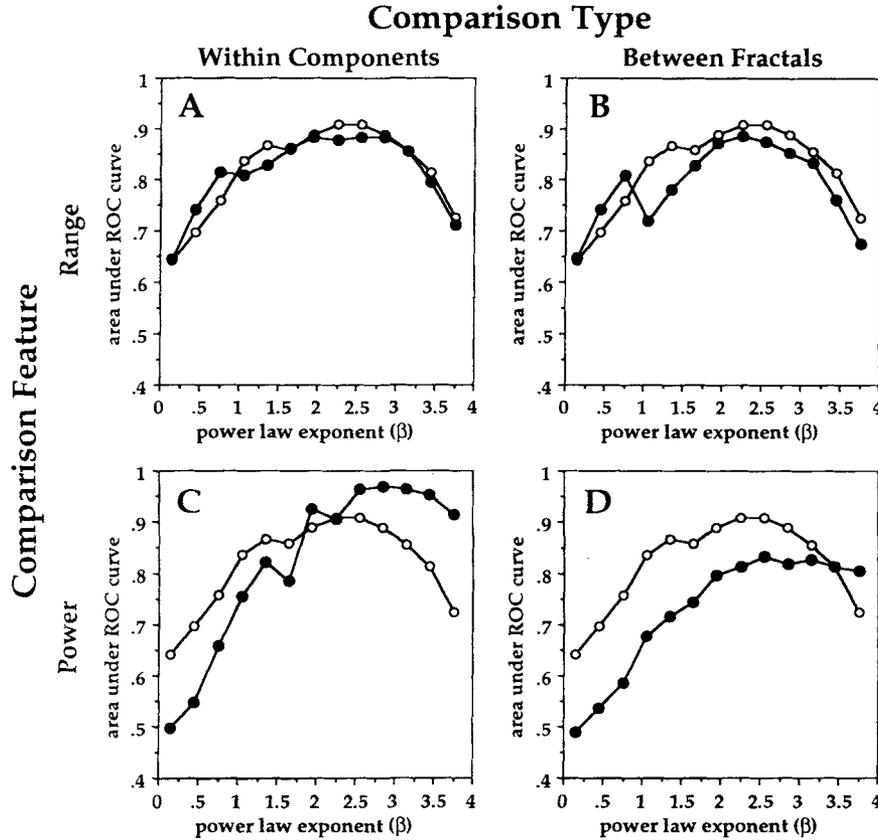


Figure 13. The 2×2 factorial block of Experiment 1 models. (The models are distinguished by the extracted contour feature: In Panels A and B it is range; in Panels C and D it is log spectral power. They are also distinguished by the logic of comparison: In Panels A and C distinctions are made within the smooth and rough components across fractals; in Panels B and D the two components are integrated into a unitary measure and distinctions are made between fractals on the basis of this measure. ROC = receiver operator characteristic.)

distribution. $d'(\beta_i, \beta_j)$ is a measure of distance. In signal detection theory and theories of comparative judgment, the distributions are of immanent neural activity and d' is a measure of discriminability.

Figure 14 shows the results of calculations of $d'(\beta_i, \beta_j)$ for the features range(smooth), range(rough), and Δ log band-pass power. The values of $d'(\beta_i, \beta_j)$ were computed from distributions consisting of 2,000 exemplars from each fractal family. These calculations serve two purposes. First, the discrimination sensitivities of the algorithmic observers can be accounted for by establishing that the simulated areas under the ROC curves are highly correlated with the values of $d'(\beta_i, \beta_j)$ associated with relevant feature distributions. We find that the summed variable [smooth] $d'(\beta_i, \beta_j)$ + [rough] $d'(\beta_i, \beta_j)$ is almost perfectly correlated with the city block range algorithm ($r = .98$) and so too is [Δ log band-pass power] $d'(\beta_i, \beta_j)$ with the spectral power between-fractal algorithm ($r = .99$). Second, it is clear from the graphs of $d'(\beta_i, \beta_j)$ that range(smooth) and range(rough) taken together generate distributions that are more widely separated than Δ log band-pass power. Although Δ log band-pass power estimates the power law exponent, β , it is

not optimal for accuracy. The range distributions are more widely separated and are, therefore, intrinsically capable of supporting more refined discrimination. Note that it is possible to construct models based on spectral power that discriminate much better than do humans, perfectly in fact. Band-pass filtration of the contours that use a number of narrow filters, followed by a regression analysis of the slope in the log (power), log (frequency) plane yields an almost error-free estimate of β . However, models based on coarse band-pass filtering generate estimates of β that have error distributions, and these distributions have considerable overlap.

The case against spectral power being a discrimination feature is strengthened in Panel C of Figure 13. The algorithm displayed here discriminates on the basis of separate power contrasts within the rough and smooth components. This algorithm differs from that depicted in Panel A only in terms of the feature extracted, power versus range. The logical forms of the contrasts are identical. The discrimination sensitivity for the power observer bears little resemblance to the data. That the range observer provides an excellent model of the data, and that neither of the power observers is successful, is good evidence—

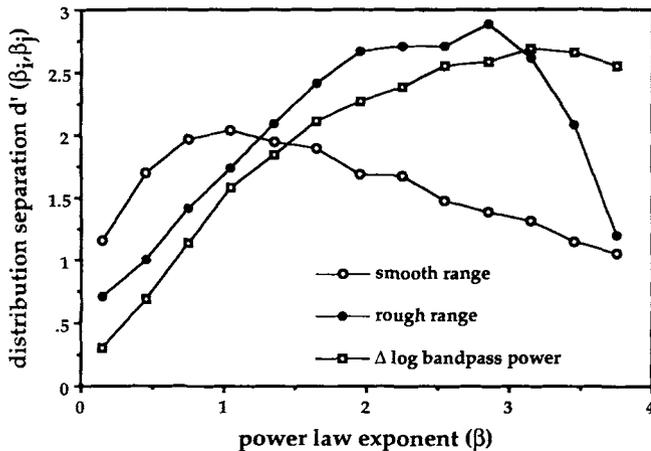


Figure 14. $d'(\beta_i, \beta_j)$ plotted for three features: smooth range, rough range, and $\Delta \log$ band-pass power. ($d'(\beta_i, \beta_j)$ is a measure of the distance between the distributions of these features defined on the fractal families compared in Experiment 1.)

as good as can be had within a simulation methodology—that spectral power is not an extracted feature. We conclude that there is no support for the notion that human observers use spectral power in discriminating between fractals. The evidence is clearly that they use the much simpler and more visually salient feature of range.

The model illustrated in Panel B of Figure 13 is critical for the interpretation of how people discriminate fractals. This model uses range as an index of contour variation to characterize individual fractals, not as a point of comparison between the rough and smooth extractions from two fractals. The quantity that this simulated observer uses is the range difference, $\text{range}(\text{smooth}) - \text{range}(\text{rough})$. The analogue to the range-difference observer is the power law exponent estimator depicted in Panel B. Range, when used in this way, generates an unavoidable pathology; there is a large dip in sensitivity near $\beta = 1$. This dip is caused by the intersection of the respective ranges that is illustrated in Panel A of Figure 11. Near the intersection, the range difference becomes quite small and this leads to a severe drop in discrimination sensitivity. The conclusion to be drawn by the failure of this model to describe human performance is that people do not use $\text{range}(\text{smooth})$ and $\text{range}(\text{rough})$ together as a single index of fractal contour structure. The two components are not viewed as forming an integral structure for the purposes of discrimination. Rather, the smooth component is taken to be one category of structure, and the rough component is taken to be another. Range comparisons are made within the respective categories.

The multifractal, non-self-affine contours used in Experiments 2 and 3 motivated additional models. In Figure 15 we show results of simulations for all three experiments for three discrimination algorithms. The first algorithm is the one that produced the best replication of Experiment 1 data; range contrasts within the smooth and rough components. The second algorithm takes into account only the range contrast of the rough component. The third algorithm computes the contrast of the standard deviation of the increment distributions. The

most important panels in this figure are those along the diagonal. They represent models based on the informal observations of contours in Figures 9 as to what information might be used for discrimination in the respective experiments. It is evident that the fits along the diagonal are uniformly excellent. (The rough-smooth simulation of Experiment 1 is the same as in Panel A of Figure 13.) In each experiment, the algorithm's discrimination sensitivity over the entire range of power law exponent, β , is virtually identical to the sensitivity displayed by subjects.

Figure 16 depicts the results from simulations based on increment correlation. The correlation observer suffers from two problems. In Experiments 1 and 3 it does not make any mistakes. These results suggest that either people do not use correlation in a discrimination task, or they do so in some incomplete way that induces error. However, in Experiment 2, the correlation observer makes too many mistakes for small β and not enough for large β . Although it might be possible to develop models for how people derive partial calculations of increment correlation, it is not obvious how to proceed. These problems do not arise for the observer that computes the width of the distribution of increments. The increment σ observer successfully models data from Experiment 3 without any tinkering with the algorithm.

These simulations allow us to draw several conclusions about how rough contours are discriminated. The first is that when people look at rough contour, they essentially see two things, noise and trend. These two forms of structure are not regarded as having a coherent relationship but rather are treated as two separate categories of information. Discrimination is based on range comparisons within components. It is important to underscore that people have access to a discrimination strategy that is generally more sensitive; in Experiments 1 and 2 the observer that discriminates in terms of the width of the increment distribution is generally more sensitive over the domain of fBm families. Yet, our simulations suggest that if there are range differences in at least one component, discrimination is based on bipartite decomposition. In Experiment 1, in which both smooth and rough contrasts are salient, ranges from both components are used. In Experiment 2 in which the smooth contrast is suppressed, the rough range is still used.

Access to a second strategy is revealed in Experiment 3 in which we designed hybrid contours that could not be discriminated from fractals on the basis of range within a smooth-rough decomposition. It was evident from the contours displayed in Figure 9 that the natural decompositions were not informative. In fact, when bipartite decompositions are simulated for Experiment 3 contours, the algorithms perform worse than did subjects. When neither the smooth nor the rough extraction provides an adequate basis for discrimination, subjects can extract the more sensitive measure of the width of the increment distributions.

A strategy of discrimination that we have not discussed, but which should be mentioned for the sake of completeness, is one where only the smooth extraction is used for discrimination. Such a strategy could have been attempted in Experiment 1. We have performed simulations of this strategy in which a contrast function was computed on the range of the smooth component only. Smooth contrast observers perform above chance, but the

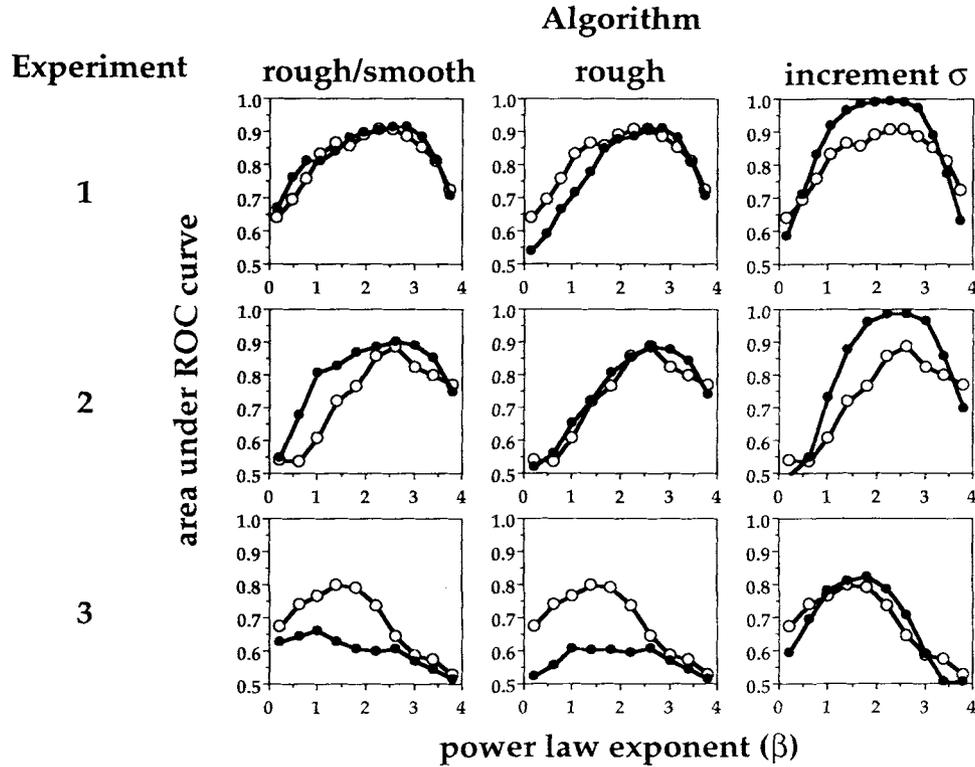


Figure 15. Simulated observers based on decomposition and the standard deviation of the distribution of increments compared with data in Experiments 1, 2, and 3. (Open circles are data; filled circles are simulation. On the diagonal are simulations suggested by analysis of the contours illustrated in Figure 9. ROC = receiver operator characteristic.)

area under the ROC curve never exceeded .75, whereas human observers in Experiment 1 had areas near .9 over a wide range of power law exponent. In fact, the discrimination data exceeded that of the simulated smooth contrast observer for each family comparison. Human observers evidently use both the rough and smooth components for the purpose of discrimination when the smooth component is sufficiently salient.

The agreement between simulation and data demonstrates that observers perceptually organize rough contour, extract features within these organizations, and apply them in a metric of comparison that is consistent in detail with the algorithms that we have formulated. The simulations and data converge on a point-by-point basis, and this is a much stronger result than the correlation between sensitivity and environmental fractal frequency. The notion that the visual system is tuned to the statistical distribution of environmental form (Gilden & Schmuckler, 1989; Knill et al., 1990) is evidently founded on a fortuitous coincidence. The sensitivity curves in Experiment 1 for the discrimination of fractal contours have the shape they do because of the logic of discrimination. This logic not only has nothing to do with fractal form, it violates fractal form by imposing a smooth-rough decomposition on hierarchically integrated structures and by using the extracted components as separate categories of information. Such are the dangers of drawing conclusions based on correlation.

Bipartite Decomposition as a Principle of Organization

It is evident that fractal contours have no special status or priority in contour discrimination. They are perceived in terms of a bipartite metatheory of structure, a theory founded on the notion that surface roughness is supported by a coherent underlying form. This is a theory that does violence to the nested hierarchic structure of self-affine contour. The hierarchy that unites the different scales of roughness in a fractal is either not perceptually penetrable or not used in discrimination. The way people appear to think about contour is in terms of signals and noise. They look for trend and treat the structure that is carried by the trend as the noise, even if the distal contour does not support this structural distinction. This way of looking at the contour is not necessitated; it is a form of perception. This form persists even with immersion in a world of fractal structures.

The perception of roughness in terms of a bipartite decomposition should be viewed as a principle of organization. As a principle it is distinguished from the protocols that establish figure-ground relationships. In figure-ground organization, both components mutually support each other, and both are essential for either to be perceived. In a bipartite decomposition that treats part of a contour as signal and part as noise, this mutuality is lacking. The signal is not ground because the noise is not figure. Neither is the noise ground; it does not provide a

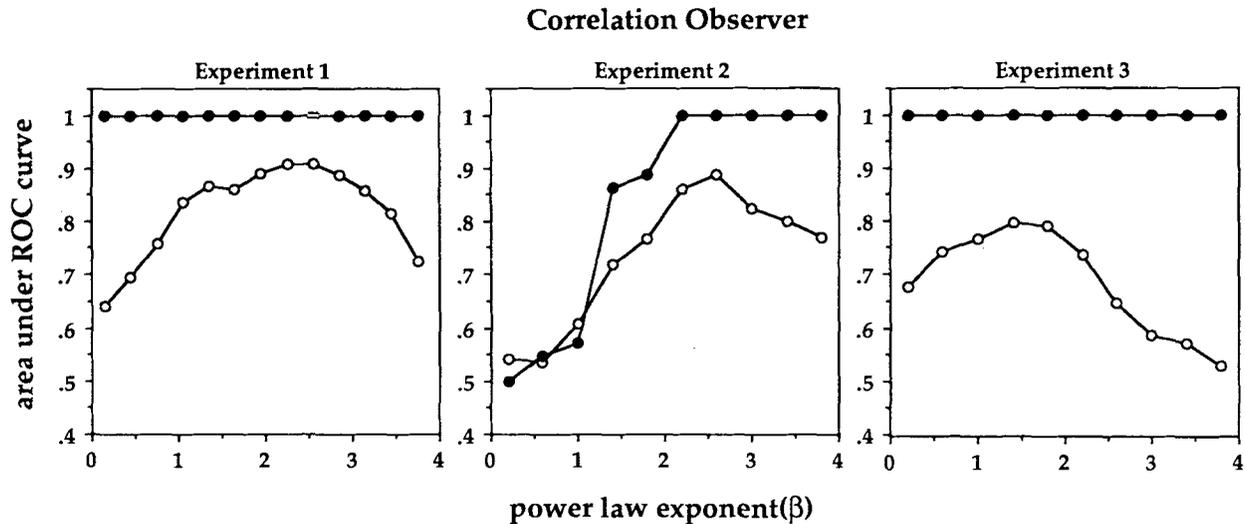


Figure 16. A simulated observer based on the correlation between successive increments is compared with data. (Open circles are data; filled circles are simulation. ROC = receiver operator characteristic)

context for the signal. It is treated in perception simply as something extra, added on.

Bipartite decomposition is the way that perception routs out structure. It is a fundamental principle of organization in that it is implicit in every perceptual act. Yet, this principle does not fall within the purview of traditional theories of perception. It is not a minimum principle (Köhler, 1947) because it does not recognize hierarchic structures, and hierarchies are the minimal encodings of natural form. The entire nested structure of a fractional Brownian motion is specified by a single parameter, β . Neither is it a principle based on likelihood (Helmholtz, 1910/1962) or intelligent inference (Rock, 1983, 1984). The principle of decomposition contains the irony that it resolves the ambiguity of proximal stimulation in terms of an agenda that is insensitive to the information available in natural structure. This principle can only be understood as a metatheory of structure, as a metaphysical statement about the way the world is constructed.

The metaphysics that is expressed by bipartite decomposition is not recondite or arcane and should be placed into the context of the philosophical literature on what constitutes structure. It is a mode of thought that attempts to distinguish what is essential from what is not, and it is based on the idea that there is indeed something essential, that there are essences. It is founded in Plato's distinction between ideal form and phenomenal appearance. It is reiterated in British empiricism in the distinction between primary and secondary qualities. It reappears in 19th-century German idealism as the thing-in-itself distinguished from the thing perceived. All of these distinctions are at root an attempt to distinguish invariant structure. For Plato, invariance is achieved through an ideal realm where change does not occur. For Locke, it is the Newtonian properties of objects that are invariant. The thing-in-itself is essentially invariance per se.

It is arguable that most perceptual tasks are linked to distal

events that do have a signal-noise structure. This way of perceiving has clear ecological utility: Prey and predators are, in fact, smooth invariant structures that are often masked by camouflaging noise. Measurement of any kind is inevitably the distillation of invariance. Indeed, if it were not for turbulent solutions to the equations of hydrodynamical flow, fractals might not be alternative structural possibilities in nature.

What we wish to point out is not that this mode of perceptual analysis exists, but that it has become completely generalized. As a principle of organization, it reduces all proximal data in the same way. This tyranny has the odd consequence that it was a mathematical discovery that natural form is hierarchic. The fractal structure of natural form is not a perceptual truism. That it is not is the best evidence that there is a general principle of perceptual organization that preempts a hierarchical analysis of form.

Summary

In this article we have inquired into the way that rough contour is perceived. We were particularly interested in fractal contours as a special class because of the seminal mathematical work by Mandelbrot, which has redefined the geometry of natural form. The observation that fractals are part of everyday natural experience led to the question of how people view hierarchic structure. Three discrimination experiments were conducted to establish a database for theoretical modeling of discrimination sensitivity. The results of Experiment 1 were of interest in themselves because they suggested that people's sensitivities might be tuned to the frequency with which fractals occur in nature.

Theoretical models of discrimination were constructed in terms of procedures that did not recognize the structural properties of random fractals. Instead, they incorporated principles common to the notion that rough contours can be decomposed

into smooth differentiable trends and rough additions. This way of decomposing contour was discussed in terms of a general theory of bipartite structure that is based on a logic of dividing objects into signal and noise. We found that the theoretical models produced discrimination sensitivities that were virtually indistinguishable from those produced by our subjects. This result validated the logic of the models and also forced us to conclude that fractal structure is either not perceptually penetrable in random contours or not sufficiently salient to be used as a basis for discrimination.

This inquiry motivated a fresh look at the way structure is perceived in general. We suggested that the signal-noise way of thinking about structure be viewed as a principle of perceptual organization. It is a principle of perception distinct from a principle of natural order. The work of Mandelbrot and others has made it clear that the natural order consists of integrated hierarchies. Perception, however, does not appear to operate in terms of hierarchies. Rather, its agenda is one of establishing a two-part division in which coherency and invariance are distinguished from noise and variability. This agenda is enforced in spite of the prevalence of structural alternatives. The perceptual analysis that is brought to fractal structure is apparently guided by a design more tuned for the breaking of camouflage than for the appreciation of statistical self-affinity.

References

- Barnsley, M. F. (1988a). *Fractals everywhere*. San Diego, CA: Academic Press.
- Barnsley, M. F. (1988b). Fractal modeling of real world images. In H.-O. Peitgen & D. Saupe (Eds.), *The science of fractal images* (pp. 219–242). New York: Springer-Verlag.
- Barrow, H. G., & Tenenbaum, J. M. (1986). Computational approaches to vision. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of perception and human performance* (Vol. 2, pp. 38-1–38-70). New York: Wiley.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115–147.
- Blum, H. (1973). Biological shape and visual science (Part 1). *Journal of Theoretical Biology*, *38*, 205–287.
- Burrough, P. A. (1981). Fractal dimensions of landscapes and other environmental data. *Nature*, *294*, 240–242.
- Carlton, E. H., & Shepard, R. N. (1990). Psychologically simple motions as geodesic paths: I. Asymmetric objects. *Journal of Mathematical Psychology*, *34*, 127–188.
- Cutting, J. E., & Garvin, J. J. (1987). Fractal curves and complexity. *Perception & Psychophysics*, *42*, 365–370.
- Falmagne, J. C. (1986). Psychophysical measurement and theory. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of perception and human performance* (Vol. 1, pp. 1-1–1-66). New York: Wiley.
- Feder, J. (1988). *Fractals*. New York: Plenum Press.
- Gilden, D. L., & Schmuckler, M. A. (1989, November). *The perception of fractal contour*. Paper presented at the annual meeting of the Psychonomic Society, Chicago, IL.
- Helmholtz, H. von. (1962). J. P. Southall (Ed. and trans.), *Treatise on physiological optics* (Vol. 3). New York: Dover. (Originally published 1910)
- Hoffman, D. D., & Richards, W. A. (1984). Parts of recognition. *Cognition*, *18*, 65–96.
- Hoffman, D. D., & Richards, W. A. (1988). Representing smooth plane curves for recognition: Implications for figure-ground reversal. In W. Richards (Ed.), *Natural computation* (pp. 76–82). Cambridge, MA: MIT Press.
- Keller, J. M., Crownover, R. M., & Chen, R. U. (1987). Characteristics of natural scenes related to the fractal dimension. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *9*, 621–627.
- Knill, D. C., Field, D., & Kersten, D. (1990). Human discrimination of fractal images. *Journal of the Optical Society of America*, *7*, 1113–1123.
- Koenderink, J. J., & van Doorn, A. J. (1988). The shape of smooth objects and the way contours end. In W. Richards (Ed.), *Natural computation* (pp. 115–124). Cambridge, MA: MIT Press.
- Köhler, W. (1947). *Gestalt psychology: An introduction to new concepts in modern psychology* (Rev. ed.). New York: Liveright.
- Mandelbrot, B. B. (1967). How long is the coast of Britain? Statistical self-similarity and fractional dimension. *Science*, *155*, 636–638.
- Mandelbrot, B. B. (1983). *The fractal geometry of nature*. San Francisco: Freeman.
- Pentland, A. (1986). Fractal-based description of natural scenes. In A. Pentland (Ed.), *From pixels to predicates* (pp. 227–252). Norwood, NJ: Ablex.
- Piotrowski, L. N., & Campbell, F. W. (1982). A demonstration of the visual importance and flexibility of spatial-frequency amplitude and phase. *Perception*, *11*, 337–346.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. (1986). *Numerical recipes*. Cambridge, MA: MIT Press.
- Ramachadran, V. S. (1990). Visual perception in people and machines. In A. Blake & T. Troscianko (Eds.), *AI and the eye* (pp. 21–77). New York: Wiley.
- Richards, W., Dawson, B., & Whittington, D. (1988). Encoding contour shape by curvature extrema. In W. Richards (Ed.), *Natural computation* (pp. 83–98). Cambridge, MA: MIT Press.
- Richards, W., Koenderink, J. J., & Hoffman, D. D. (1988). Inferring 3D shapes from 2D silhouettes. In W. Richards (Ed.), *Natural computation* (pp. 125–137). Cambridge, MA: MIT Press.
- Rock, I. (1983). *The logic of perception*. Cambridge, MA: MIT Press.
- Rock, I. (1984). *Perception*. San Francisco: Freeman.
- Saupe, D. (1988). Algorithms for random fractals. In H.-O. Peitgen & D. Saupe (Eds.), *The science of fractal images* (pp. 71–136). New York: Springer-Verlag.
- Sayles, R. S., & Thomas, T. R. (1978a). Surface topography as a nonstationary random process. *Nature*, *271*, 431–434.
- Sayles, R. S., & Thomas, T. R. (1978b). Topography of random surfaces. *Nature*, *273*, 573.
- Tennekes, H., & Lumley, J. L. (1983). *A first course in turbulence*. Cambridge, MA: MIT Press.
- Thurstone, L. L. (1927). A law of comparative judgment. *Psychophysical Review*, *34*, 273–286.
- Voss, R. F. (1985). Random fractal forgeries. In R. A. Earnshaw (Ed.), *Fundamental algorithms for computer graphics* (NATO ASI Series, Vol. F17, pp. 805–835). Berlin: Springer-Verlag.
- Voss, R. F. (1988). Fractals in nature: From characterization to simulation. In H.-O. Peitgen & D. Saupe (Eds.), *The science of fractal images* (pp. 21–70). New York: Springer-Verlag.
- Westheimer, G. (1991). Visual discrimination of fractal borders. *Proceedings of the Royal Society of London: Series B. Biological sciences*, *243*, 215–219.

Received October 1, 1991

Revision received October 19, 1992

Accepted January 4, 1993 ■