From Splines to Fractals

Richard Szeliski[†] and Demetri Terzopoulos[‡]

[†]Digital Equipment Corp., Cambridge Research Lab, One Kendall Square, Bldg. 700, Cambridge, MA 02139
[‡]Schlumberger Laboratory for Computer Science, P.O. Box 200015, Austin, TX 78720

Abstract

Deterministic splines and stochastic fractals are complementary techniques for generating free-form shapes. Splines are easily constrained and well suited to modeling smooth, man-made objects. Fractals, while difficult to constrain, are suitable for generating various irregular shapes found in nature. This paper develops *constrained fractals*, a hybrid of splines and fractals which intimately combines their complementary features. This novel shape synthesis technique stems from a formal connection between fractals and generalized energy-minimizing splines which may be derived through Fourier analysis. A physical interpretation of constrained fractal generation is to drive a spline subject to constraints with modulated white noise, letting the spline diffuse the noise into the desired fractal spectrum as it settles into equilibrium. We use constrained fractals to synthesize realistic terrain models from sparse elevation data.

Keywords: Fractals, Splines, Constraints, Scattered Data Interpolation, Digital Terrain Models, Physically Based Modeling, Deformable Models
CR Categories: I.3.5—Object Modeling (Curve, surface, solid, and object representations); I.3.7—Three-Dimensional Graphics and Realism; G.1.1—Interpolation (Spline interpolation); G.1.2—Approximation (Spline approximation)

1 Introduction

Over the years, computer graphics researchers have developed numerous mathematical models capable of generating free-form shapes. Such models come in two varieties—deterministic and stochastic. Spline models, which are deterministic, have established themselves as a convenient and powerful technique for modeling smooth, man-made shapes, such as teapots [2]. By contrast, fractal models have become popular for recreating a wide variety of the shapes found in nature [10]. Most fractal models feature a stochastic component, making them well suited to generating nonsmooth, irregular shapes, such as mountainous terrain [4]. In this paper, we develop a model of shape which combines deterministic splines and stochastic fractals to inherit their complementary features.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

1.1 Constraints versus Natural Detail

Splines typically offer precise shape control but they lack natural looking detail. In this paper, we employ a class of variational splines whose positions, slopes, and curvatures are locally controllable though external shape constraints. Furthermore, these variational models can become piecewise smooth, producing local discontinuities such as fractures and creases. However, these splines provide no mechanism for modeling the fine-scale texture of many natural shapes.

By contrast, stochastic fractals provide realistic detail for modeling a wide variety of complex natural phenomena, but they offer little control over shape. This deficiency is most acute in the Fourier based methods [10] which can produce shapes with true fractal distributions but cannot be controlled locally. A common approach for obtaining some control is to first triangulate a given set of points, then add fractal texture by recursively subdividing and randomly perturbing the subtriangles [4]. Unfortunately, this approach produces annoying visual artifacts because the spatial statistics are nonstationary across the original triangle boundaries. Moreover, it makes difficult the imposition of more complicated constraints.

Conventional free-form shape synthesis techniques are therefore inadequate for many graphics applications—splines provide insufficient detail, while fractals provide insufficient shape control. This paper proposes ***constrained fractals***, a new shape modeling technique which simultaneously provides both detail and control.

1.2 Overview

Section 2 gives an intuitive explanation of the constrained fractal technique, and discusses fractal terrain generation which serves as our main application area. Section 3 presents the controlled-continuity splines which form the deterministic component of constrained fractals. The main text develops these variational splines in a single variable, and multivariate extensions are relegated to an appendix. We describe how to discretize the spline energy expressions using finite elements, and present examples of spline surfaces constrained by scattered data. Section 4 explains how to associate a probability distribution with variational splines through the Boltzmann distribution, and shows that samples from this distribution have fractal statistics. Section 5 introduces a multiresolution stochastic relaxation

©1989 ACM-0-89791-312-4/89/007/0051 $00.75

algorithm for sampling the Boltzmann distribution to obtain constrained fractals. Because of length limitations, we present a full description of our algorithm in a companion paper [14]. In Section 6 we demonstrate the application of this algorithm, showing how to include local control over smoothness. Finally, Section 7 closes with some conclusions about our work.

2 The Constrained Fractal Approach

2.1 Deterministic Component

A constrained fractal is a physically-based model which makes use of the energy minimization principles underlying variational splines [1].[1] These principles characterize spline curves and surfaces as the "smoothest" shapes consistent with the given constraints; this gives rise to an interesting relationship connecting variational splines to spatial smoothing filters (see [16,13,12]). What is surprising, however, is that "spline smoothing filters" can provide the spatial frequency characteristics of fractal fields [13].

To exploit the filtering property of variational splines in generating constrained fractals, we employ a general class of multivariate spline models called controlled-continuity splines [16]. These splines allow us to create the desired shape by applying local constraints (say, on position and/or orientation) at arbitrary points. Physically, the constraints are interpretable as external forces which shape the spline [15]. Controlled-continuity splines also afford local control over smoothness, which permits us to introduce jump or crease discontinuities at arbitrary points on the spline to yield a piecewise smooth shape.

2.2 Stochastic Component

To introduce fractal detail into the shape, we go one step further. As it minimizes its deformation energy under the influence of constraint forces, we subject the controlled-continuity spline to white noise. Visualize this physically as continually bombarding the spline shape with point masses impacting at random velocities. The impulses imparted by these projectiles randomly perturb the spline as it "relaxes" into a stochastic equilibrium (characterized by minimal time-averaged energy under the random perturbations). During relaxation the spline diffuses the effects of the perturbations spatially, eventually shaping the flat spectrum of the perturbations into the desired fractal spectrum.[2]

2.3 Bayesian Interpretation

We can also motivate our approach by making use of another interesting relationship, first developed in statistical mechanics, which connects energy functionals to random fields via the Boltzmann probability distribution [7]. The relaxing controlled-continuity spline can therefore be interpreted as a "Gibbs Sampler" which draws sample shapes from a Boltzmann distributed ensemble that has fractal statistics due to the spline's internal energy [13]. The constrained fractal is thus a typical sample from the *posterior* distribution, i.e., a random fractal sample consistent with the observations which are provided by the constraints.

2.4 Algorithmic Features

The stochastic equilibrium of the spline is computable by a stochastic relaxation algorithm implemented as a local, iterative, numerical process. We obtain such an algorithm by locally discretizing the continuous spline energy functionals using regular finite difference or finite element grids.[3]

Each relaxation step replaces the values of nodal variables on the grid by a noise-perturbed weighted combination of neighboring variable values, thereby propagating information from node to node across the grid. Because of the local nature of the communication, a large number of iterations may be necessary for the stochastic relaxation process to converge on large, fine grids. If the variational spline problem is discretized on a sequence of successively coarser grids, however, the convergence rates improve dramatically [15]. In this paper we propose a multiresolution stochastic relaxation process which computes the interpolated fractal shape efficiently. The multiresolution process is readily parallelizable, and is therefore suitable for massively parallel computers.

2.5 Fractal Terrain Generation

As an application of our constrained fractal technique, we concentrate on the problem of terrain generation for the synthesis of realistic outdoor scenes.

Fractal terrain generation [10] has traditionally been approached using a variety of stochastic subdivision techniques. Fournier *et al.* [4] employ a technique called random midpoint displacement which creates a tessellated surface by recursively subdividing triangles. The height of each newly created interior point is randomly perturbed away from its original interpolated value, and the magnitude of each perturbation is related to the level of the subdivision. Varying this relationship results in fractals of arbitrary degree. Voss [19] uses successive random additions, which differs in that all of the points are randomly perturbed at each subdivision step (not just the newly cre-

[1] Variational splines are simple instances of "deformable models" (see [18,17]); hence, the constrained fractal technique extends to this general class of physically-based models as well.

[2] As Lewis [9] has argued, the spectrum of the random process need not be truly fractal in order to provide realistic detail. We use the term "fractal" to describe a random function whose spectrum we can control.

[3] Regular ("fine-grained") grids are independent of the spatial organization of the data and lead to algorithms that are readily implementable on massively parallel computational structures. These properties account for the prevalent use of regular grids in the computer vision community, as does the massive parallelism of the early human visual system [8]. Our use of regular grids stands in contrast to conventional techniques for interpolating sparse data in computer graphics applications. The latter usually involve irregularly triangulating the domain in a data dependent fashion [5] (techniques have also been developed for the simpler case of interpolating "gridded" data [3]).

ated ones). Lewis [9] proposes generalized stochastic subdivision, a refinement of random midpoint displacement. Instead of displacing each midpoint independently, he adds correlated Gaussian noise, which alleviates the artifacts due to spatially nonstationary statistics across triangles that are sometimes evident with the preceeding methods.

Unlike the above methods, our multiresolution stochastic relaxation algorithm can easily accommodate arbitrary constraints as additional terms in its energy minimization principle. The solution is guaranteed to have spatially stationary statistics, so long as we perform a sufficient number of iterations on the fine level (in practice, this number proves to be quite low).

To apply our constrained fractal technique to terrain modeling, we may begin either with synthetic elevation landmarks selected by a user or with true elevation data acquired in surveys or through aerial photogrammetry. The first step is to create the terrain model from these typically sparse elevation data (e.g., isoelavation contours) by interpolating a dense digital elevation map in the form of a gridded, single-valued surface $f(u, v)$. To do so, we supply these data to our algorithm as elevation constraints d_i, letting the controlled-continuity spline serve as an interpolant whose parameters afford local control over the continuity of the surface and the tightness of fit to the data. At the same time, the stochastic mechanism invents fractal detail that enhances the realism of the interpolated terrain model. Thus our technique solves the surface fitting and detail generation problems simultaneously.

3 Variational Splines

Variational splines are characterized by the minima of energy functionals. To construct a variational spline, we first define a deformation energy functional $\mathcal{E}$ over a suitable class of functions, and then compute a function which minimizes $\mathcal{E}$. In the univariate case, the minimizing functions approximate the steady-state deformations of elastic strings and beams subjected to applied forces.

Suppose that we stretch an elastic string horizontally along the x axis. Let $x = u$ denote the coordinate along the string and let $y = f(u)$ give the shape of the string as a function of u. The (linearized) deformation energy associated with f is given by

$$\mathcal{E}(f) = \frac{1}{2}\int f_u^2\, du,$$

where f_u indicates differentiation with respect to u.

Next, suppose that $f(u)$ gives the shape of an elastic beam. The (linearized) deformation energy of the beam is

$$\mathcal{E}(f) = \frac{1}{2}\int f_{uu}^2\, du,$$

where f_{uu} denotes the second derivative. Note that the minimum of this expression characterizes the common cubic spline [1]. Because the string has tension, it defines a continuous curve f. Since the beam resists bending, it defines a smoother curve f that is not merely continuous in position but also in tangent.

3.1 Controlled-Continuity Splines

We can blend the above energies to create a model that combines the properties of the string and beam. Introducing nonnegative rigidity $\rho(u)$ and tension $[1-\tau(u)]$ parameter functions that take values between 0 and 1 inclusively, the hybrid energy functional is

$$\mathcal{E}_{\mathrm{p}}(f) = \frac{1}{2}\int \rho(u)\{[1-\tau(u)]f_u^2 + \tau(u)f_{uu}^2\}\, du. \quad (1)$$

Note, that if we restrict $\rho(u) = \rho$ and $\tau(u) = \tau$ to be global constants independent of u, the above spline reduces to the globally continuous "spline under tension" proposed by Schweikert [11].

As functions of u, however, $\rho(u)$ and $\tau(u)$ provide local control over the continuity of the curve. Consequently, the function f which minimizes (1) is a *controlled-continuity spline* [16]. In regions of u where $\rho(u) > 0$ and $\tau(u) > 0$, the spline is continuous in both position and tangent, tending towards the string as $\tau(u) \to 0$ and towards the beam as $\tau(u) \to 1$. In regions where $\tau(u) = 0$, the spline is free to have discontinuous tangents (creases), and in regions where $\rho(u) = 0$, the spline is free to have discontinuous positions (fractures).

The univariate controlled-continuity splines generalize to any number of variables, to any embedding space dimensionality, and to arbitrary order of continuity (see Appendix A), thereby providing a unified treatment of curves, surfaces, solids, or higher-dimensional models (in spacetime) [16]. Since in this paper we concentrate on terrain modeling applications, we make use of the bivariate case—surfaces. Appendix A defines controlled-continuity surface splines based on the membrane and plate, the natural bivariate analogues of the string and beam.

3.2 Fitting Data

To formulate the spline fitting problem, we combine the energy functional of the controlled-continuity spline with a data compatibility constraint. In the univariate case, the data is a collection of values $\{p_i\} = \{(u_i, d_i)\}$. A simple data compatibility constraint measures the squared (weighted Euclidean) distance between the data and the spline $f(u)$ using the constraint functional

$$\mathcal{E}_{\mathrm{d}}(f; \{p_i\}) = \frac{1}{2}\sum_i c_i\,(f(u_i) - d_i)^2 \quad (2)$$

where the weights c_i are inversely related to the variance of the uncertainty (noise) in the data ($c_i = \sigma_i^{-2}$).[4]

To find the approximating spline, the functionals (1) and (2) are combined to form the energy functional

$$\mathcal{E}(f) = \lambda\mathcal{E}_{\mathrm{p}}(f) + \mathcal{E}_{\mathrm{d}}(f; \{p_i\}), \quad (3)$$

where λ is known as a regularization parameter. This continuous formulation of the variational spline fitting prob-

[4] We can handle orientation (slope) data o_i by adding the term $\frac{1}{2}\sum_i c_i\,(f_u(u_i) - o_i)^2$ to (2) [16].

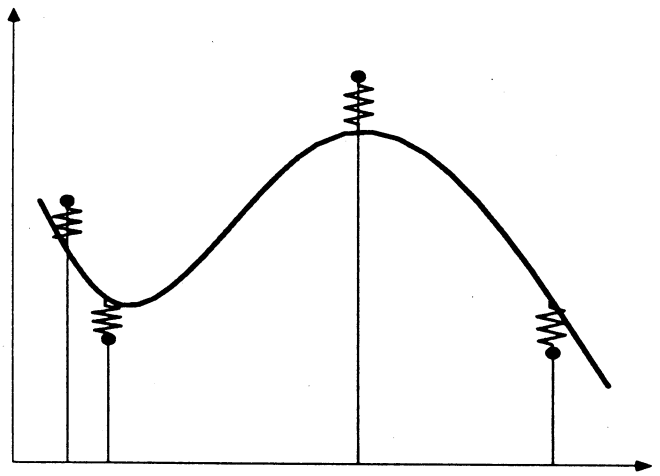

Figure 1: Spline approximation

lem in terms of a functional applies to arbitrarily structured constraint data, whether gridded, contoured, or scattered.

Figure 1 illustrates a physical interpretation of (3) in the univariate case. The first term defines the energy of the spline curve. The second term corresponds to the energy of a collection of zero-length ideal springs (with spring constants c_i) connecting the spline to the data points. The springs apply forces which deflect the spline such that it approximates the data (infinitely stiff springs result in strict interpolation). When (3) is at a minimum, the physical system is at equilibrium, such that the forces exerted by the springs balance the reluctance of the spline to deform.

3.3 Finite Element Discretization

To minimize (3), we apply the finite element method, which provides a systematic approach to the discretization and solution of variational spline problems [15]. We discretize $f(u)$ on a regular "fine grained" mesh of *nodal variables* (see Footnote 1).

When finite element analysis is applied to quadratic functionals such as (1),[5] we obtain a discrete energy function expressible as a quadratic form:

$$E_{\mathrm{p}}(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T \mathbf{A}_{\mathrm{p}} \mathbf{x} \tag{4}$$

where $\mathbf{x}$ is the vector of nodal variables.[6] The *prior model* matrix $\mathbf{A}_{\mathrm{p}}$ is sparse, having at most 13 entries per row (see [15,12] for details). Similarly, the discrete data constraint in (2) can be written as

$$E_{\mathrm{d}}(\mathbf{x}, \mathbf{d}) = \frac{1}{2}(\mathbf{x} - \tilde{\mathbf{d}})^T \mathbf{A}_{\mathrm{d}} (\mathbf{x} - \tilde{\mathbf{d}}), \tag{5}$$

[5] Actually, (1) is quadratic only if $\rho(u)$ and $\tau(u)$ are fixed functions.

[6] For example, in the univariate case $\mathbf{x} = [f(0), \ldots, f(jh), \ldots, f((M-1)h)]'$, where h is the spacing of a linear M-node mesh.

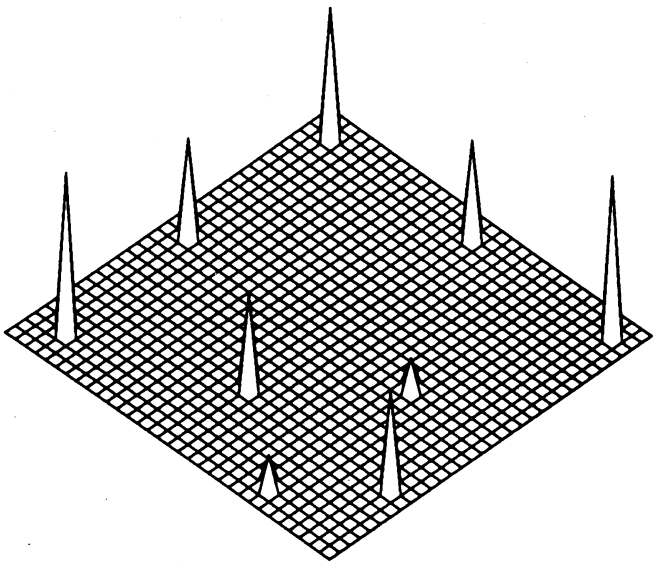

Figure 2: Sparse data points

where $\tilde{\mathbf{d}}$ is a zero-padded vector of data values, and the diagonal matrix $\mathbf{A}_{\mathrm{d}}$ has entries w_i where data points coincide with nodal variables and zeros elsewhere.

Using (4) and (5), we write the combined energy (3) in discrete form as

$$E(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T \mathbf{A} \mathbf{x} - \mathbf{x}^T \mathbf{b} + k \tag{6}$$

where k is a constant, while

$$\mathbf{A} = \lambda \mathbf{A}_{\mathrm{p}} + \mathbf{A}_{\mathrm{d}} \quad \text{and} \quad \mathbf{b} = \mathbf{A}_{\mathrm{d}} \tilde{\mathbf{d}}.$$

This energy function has a minimum at $\mathbf{x} = \mathbf{x}^*$, the solution to the linear system of algebraic equations

$$\mathbf{A}\mathbf{x} = \mathbf{b}. \tag{7}$$

In principle, the solution to (7) can be found using either direct or iterative numerical methods. Direct methods are impractical for solving large systems associated with fine meshes because of excessive storage requirements. Alternatively, simple iterative schemes such as Gauss-Seidel relaxation require relatively little storage but can be very slow to converge. We resolve the difficulty using multigrid relaxation, which accelerates convergence by solving the problem at multiple resolution levels [15].

3.4 A Surface Fitting Example

To better visualize the fitting of variational spline models to scattered data, let us consider a small example involving the bivariate controlled-continuity spline model given in Appendix A. Figure 2 shows nine data points. The interpolated thin plate ($\tau = 1$) solution on the finite element mesh is shown in Figure 3. Note that a jump discontinuity ($\rho = 0$) has been introduced along the left edge and an orientation discontinuity ($\tau = 0$) along the right edge.

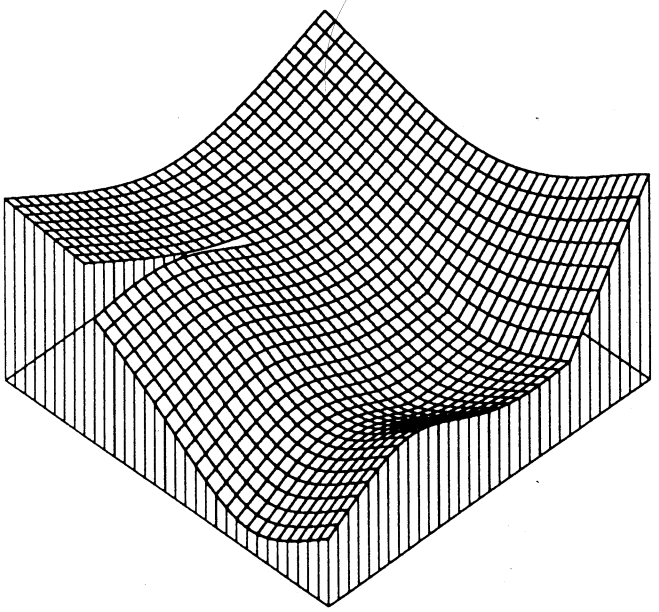

Figure 3: Interpolated piecewise continuous solution

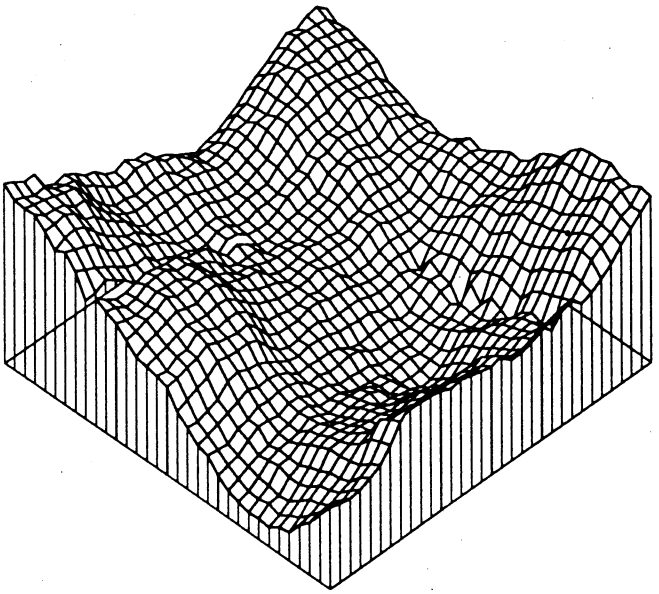

Figure 4: Sample from posterior distribution

4 Converting Energies into Probabilities

The controlled-continuity splines presented in the previous section give us a powerful approach to interpolating (or approximating) data with differing amounts of smoothness. For many graphics applications, however, the resulting shapes have insufficient detail or random texture to look natural. We will now show how to convert our spline energy functions into probability distributions from which we can draw random samples having fractal textures.

The idea of converting an energy function into a probability distribution comes from statistical mechanics. In many stochastic physical systems, the probability of a particular configuration is inversely related to its energy. Many different equations for transforming energies to probabilities are possible. For our application, we will use the Boltzmann distribution, where the energy of a state $E(\mathbf{x})$ is related to its probability $p(\mathbf{x})$ through a negative exponential:

$$p(\mathbf{x}) = \frac{1}{Z} \exp\left(-E_{\mathrm{p}}(\mathbf{x})/T\right) \tag{8}$$

(the partition function Z is used to normalize the distribution). The temperature parameter T controls how "peaked" the distribution is with respect to its low-energy states. The Boltzmann distribution has properties which make it useful for modeling random fields [7]. For our purposes, the most important of these is that multiplicative interactions between probability distributions can be converted into additive interactions between energies. This becomes particularly useful when we look at fractal generation as sampling from a constrained (conditional) distribution using the Gibbs Sampler algorithm to be presented in the next section.

Interestingly, we can use the energy of the controlled-continuity spline model to define a Boltzmann distribution with a fractal spectrum. A Fourier analysis of this prior model reveals that the resulting distribution is correlated Gaussian noise with a fractal spectrum (i.e., self-affine over scale) for the membrane or string ($\tau = 0$) and for the thin plate or beam ($\tau = 1$) [13] (see Appendix B).

When data $\mathbf{d}$ constrain the fractal shape, we may obtain it by sampling from the conditional distribution $p(\mathbf{x}|\mathbf{d})$. This posterior distribution can be calculated using Bayes' rule

$$p(\mathbf{x}|\mathbf{d}) = \frac{p(\mathbf{d}|\mathbf{x})\, p(\mathbf{x})}{p(\mathbf{d})} \tag{9}$$

where the distribution $p(\mathbf{d})$ is a normalization factor. The conditional distribution $p(\mathbf{d}|\mathbf{x})$ can be derived from a measurement model which describes how the data $\mathbf{d}$ was acquired from a sample shape $\mathbf{x}$ [12].

For a linear measurement model, the negative logarithm of the posterior distribution $p(\mathbf{x}|\mathbf{d})$ can be written as the sum of two energy functions

$$-\log p(\mathbf{x}|\mathbf{d}) = E(\mathbf{x};\mathbf{d}) = E_{\mathrm{p}}(\mathbf{x}) + E_{\mathrm{d}}(\mathbf{x},\mathbf{d}); \tag{10}$$

i.e., the posterior distribution is itself a Boltzmann distribution. We thus have a correspondence between our Bayesian models and the constraints used in the previous section to fit the generalized spline. The measurement model $p(\mathbf{d}|\mathbf{x})$ corresponds to the constraint data. The prior model $p(\mathbf{x})$ corresponds to the controlled-continuity spline. As shown in Appendix B, we can control the spectrum of this prior model through our choice of the order of the spline.

The posterior distribution $p(\mathbf{x}|\mathbf{d})$ thus defines a class of random shapes which are drawn from a family of fractals

and which are also consistent with given shape constraints (observations). The most likely sample from this distribution, the *Maximum a Posteriori* (MAP) estimate, corresponds to the minimum energy solution of (7), as shown in Figure 3. A *typical* sample from this distribution is shown in Figure 4. Devising an efficient method for generating such a random sample is the subject of the next section.

5 Multiresolution Stochastic Relaxation

The explicit evaluation of the Boltzmann distribution given in (8) is very difficult because the computation of the partition function Z requires a summation over all possible states. Fortunately, we can generate random samples from this distribution using a simple algorithm known as the Gibbs Sampler [7]. At each step of this iterative algorithm, a new random state is chosen from the Boltzmann distribution corresponding to the local energy function of the variable being updated. This updating rule is guaranteed to convergence to a random sample from the overall distribution (the ensemble is then said to be at thermal equilibrium).

For our quadratic energy function (6), the local energy function for the node x_i (with all other nodes fixed) is

$$E(x_i) = \frac{1}{2}a_{ii}x_i^2 + \left(\sum_{j\in N_i} a_{ij}x_j - b_i\right)x_i + k, \qquad (11)$$

where a_{ij} are the entries of $\mathbf{A}$, and N_i expresses the fact that the a_{ij} are nonzero only for certain neighbors of node i (see [14] for details). When using Gauss-Seidel relaxation, we choose the new node value which minimizes this local energy

$$x_i^+ = \frac{b_i - \sum_{j\in N_i} a_{ij}x_j}{a_{ii}} \qquad (12)$$

(for pure interpolation, we set $x_i^+ = d_i$ at points coincident with data). We can thus rewrite the local energy as

$$E(x_i) = \frac{1}{2}a_{ii}(x_i - x_i^+)^2 + k'. \qquad (13)$$

For the Gibbs Sampler algorithm, we choose the new value for x_i from the local Boltzmann distribution

$$p(x_i) \propto \exp\left(-E(x_i)/T\right) \propto \exp\left(-\frac{1}{2}\frac{(x_i - x_i^+)^2}{T/a_{ii}}\right) \qquad (14)$$

which is a Gaussian with mean equal to the deterministic update value x_i^+ and a variance equal to T/a_{ii}. Thus, the Gibbs Sampler is equivalent to the usual relaxation algorithm with the addition of some locally controlled Gaussian noise at each step. For instance, the sample surface in Figure 4 exhibits the rough (wrinkled) look of fractals. The amount of roughness can be controlled with the "temperature" parameter T.

5.1 Multiresolution Acceleration

Although the above iterative algorithm will eventually achieve stochastic equilibrium, the convergence may be unacceptably slow in practice. To accelerate convergence, we use coarse-to-fine relaxation on a multiresolution pyramid. The problem is first solved on a coarser mesh, and the solution is used as an initial condition for the next finer level.

The coarse-to-fine technique is thus similar to recursive subdivision and successive addition techniques [4,19,9]. Unlike these approaches, however, we do not just add random detail as the resolution is increased. Instead, we use the coarse (low resolution) solution as an initial condition for the stochastic relaxation algorithm. The iterative nature of this algorithm removes the nonstationarities that are present in the interpolated initial condition, and it also allows us to impose constraints at finer resolutions. The disadvantage of this approach is that the coarse level sample may no longer look like a subsampled version of the fine level sample. This lack of "internal consistency" becomes a problem when we wish to generate renderings at a variety of scales [4].

We can avoid this discrepancy either by limiting the number of iterations at the finer levels, or by "freezing" the points derived from the coarse level and only iterating on the new points (which resembles random mid-point displacement). The second approach is usually preferable, since it allows the fine level shape to relax sufficiently so that "creases" or other artifacts are not visible. However, if this freezing procedure occurs at a resolution much coarser than that of the data, the resulting shape will not fit the data as well as it could. A strategy which allows full relaxation until the data resolution is reached and then freezes coarse level points as more detail is added works well in practice. In [14] we present examples of this approach applied to the generation of "zoom sequences" over fractal terrain.

Using a multiresolution pyramid also gives us additional control over the spectrum of the fractal shape. When we use a membrane (string) or thin plate (beam) model in our stochastic relaxation algorithm, the resulting spectrum is fractal, but it is limited to the form

$$S_f(\boldsymbol{\omega}) \propto |\boldsymbol{\omega}|^{-2m},$$

which corresponds to fractal βs of 2 and 4 respectively (Appendix B). The simplest way to approximate an intermediate fractal degree is to use the controlled-continuity spline. If we choose $w_1 = \omega_0^2 w_2$ in (18), we obtain a power spectrum which behaves as $S_u \propto |\boldsymbol{\omega}|^{-3}$ in the vicinity of ω_0. At lower frequencies, the model behaves as a membrane, while at higher frequencies it behaves like a thin plate.

We can extend the range of intermediate fractal behavior by modifying the coarse-to-fine Gibbs Sampler. Instead of implementing the same energy equation at each level, we modify the algorithm so that a different blend of

membrane and thin plate is used at each level [14]. Since the relaxation at a given level mostly affects the short wavelength Fourier components, we can use the coarse levels to shape the low frequencies and the fine levels to shape the high frequencies. The effective power spectrum of the resulting interpolator (or random sample) now depends on the number of iterations performed at each level.

An alternative to the Gibbs Sampler algorithm for generating random samples is diffusion [6]. This approach is closer in spirit to the usual forward difference equations used in physically-based modeling [18]. Instead of adding controlled (spatially varying) Gaussian noise to the result of the relaxation step, we add uniform white noise to a procedure which descends along the energy gradient $\nabla E(\mathbf{x})$:

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \Delta t \nabla E(\mathbf{x}_k) + \sqrt{2T\Delta t}\,\mathbf{g}_k, \qquad (15)$$

where $\mathbf{g}_k$ is a random Gaussian vector with unit variance. We can thus think of the fractal generation process as a dynamic system (with internal and external forces ∇E) driven by small-scale uniform noise. The system eventually diffuses into a stochastic equilibrium which reflects the likelihood of each state as a function of its energy. For this diffusion equation to converge to the correct equilibrium distribution, the time step must be very small. In practice, the Gibbs Sampler is preferred, since it attains equilibrium much faster. Nevertheless, the diffusion equation provides an elegant intuitive model for the generation of constrained fractals.

6 Constrained Fractal Examples

The multiresolution algorithm proposed in the previous section can be used to generate a wide variety of constrained fractals. The controlled-continuity spline models used to generate these fractals include parameters which control the appearance of the final shape. In this section, we will employ cartographic data to show how these and the other parameters of the stochastic relaxation algorithm lend flexibility to our method.

We present an example of fitting constrained fractal surfaces to isoelevation contours [5] extracted from a 256×256 digital terrain map. Figure 5 shows the digital terrain map rendered as a Phong-shaded surface. Aside from the addition of the horizontal blue plane to introduce a "sea level," we forego the usual texture mapping and environmental embellishments in this and subsequent images, using instead a uniform surface reflectance to better reveal the features of the synthesized terrain. We subsampled the digital terrain map using 200m elevation contour lines and we input to our algorithm the resulting sparse contour data rendered in Figure 6.

First we illustrate the reconstruction of smooth controlled-continuity spline surfaces to the contours in Figure 6. The result of using membrane reconstruction (equation (16) with $\tau(u,v) = 0$) on these data is shown in Figure 7. Evidently, the membrane is insufficiently smooth because it has C^0 continuity only. It looks like a "tent"

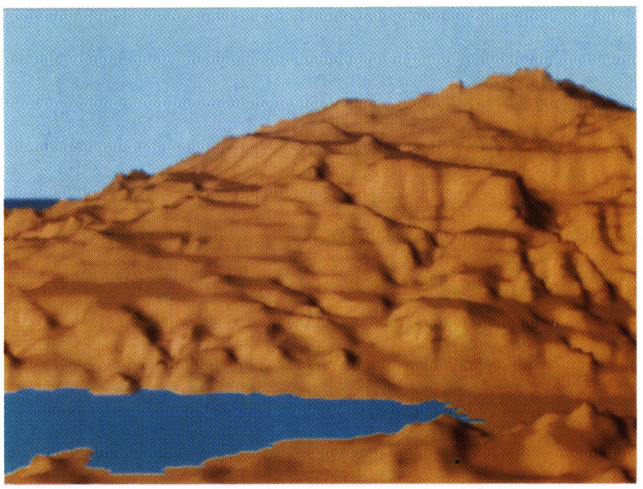

Figure 5: Original digital terrain map

or "soap film" stretched over the contours, and it flattens peaks and valleys. Figure 8 shows the thin plate reconstruction ($\tau(u,v) = 1$). In this case, the surface is perhaps too smooth, producing a large depression that forms the "lake" at the lower right. The thin plate under tension model (intermediate $\rho(u,v)$ and $\tau(u,v)$) appears to yield a good compromise, as shown in Figure 9.

Next, we invoke the noise process in our multiresolution stochastic relaxation algorithm to synthesize fractal-textured terrain which fits the contour data of Figure 6. Figure 10 shows the result of using the thin plate under tension model with a temperature $T = 50$. If we desire a "rougher" looking surface, we can use a higher temperature; for example, $T = 500$ yields the reconstructed surface shown in Figure 11. The roughness of the surface can also be controlled by changing the fractal degree.

The degree of smoothness and the amount of noise can, of course, be controlled locally to generate spatially varying textures and different kinds of terrain. Figure 12 shows an example of the complicated fractal surfaces that can be generated using our multiresolution Gibbs Sampler algorithm. This fractal scene is constrained by some data points defining the peaks of the mountains and the bottom of the "ravine." Note that it also has a crease coinciding with the "ravine" on the upper right side, and a depth discontinuity on the lower right. Four different parameter value combinations are specified for the controlled-continuity spline in each quadrant of the (u,v) domain—a thin plate under tension in the upper quadrant, a membrane in the right quadrant, a thin plate in the left quadrant, and a very stiff thin plate in the lower quadrant.

7 Conclusion

We have proposed constrained fractals, a new free-form shape modeling technique which combines the complementary features of deterministic spline and stochastic fractal models. This combination enables us to synthesize realis-

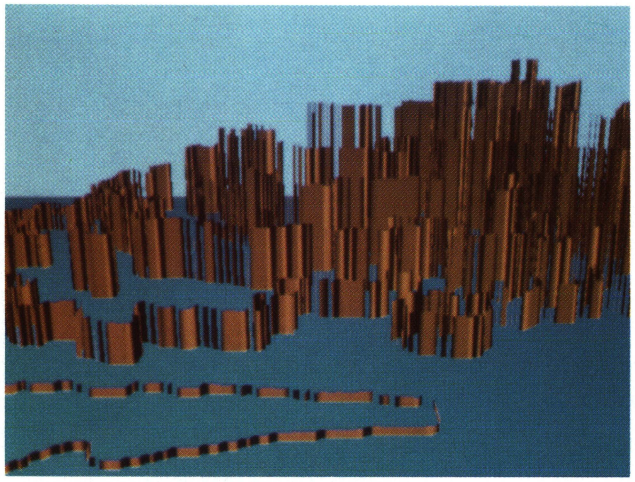

Figure 6: Rendered contour data

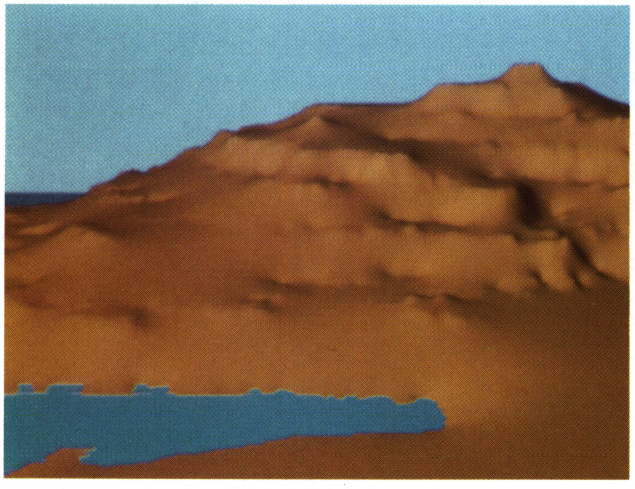

Figure 7: Interpolated membrane surface

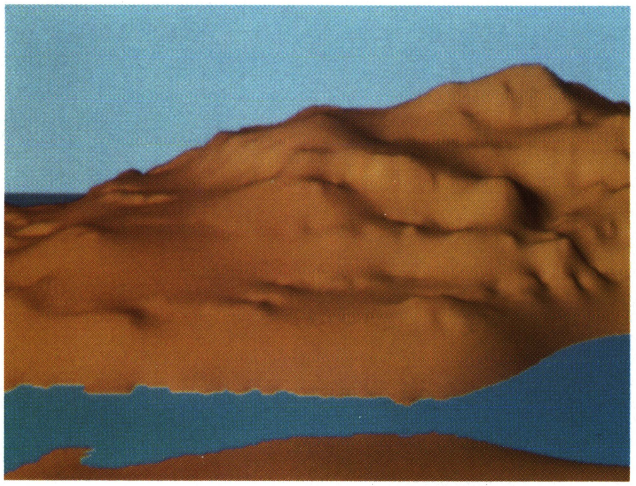

Figure 8: Interpolated thin plate surface

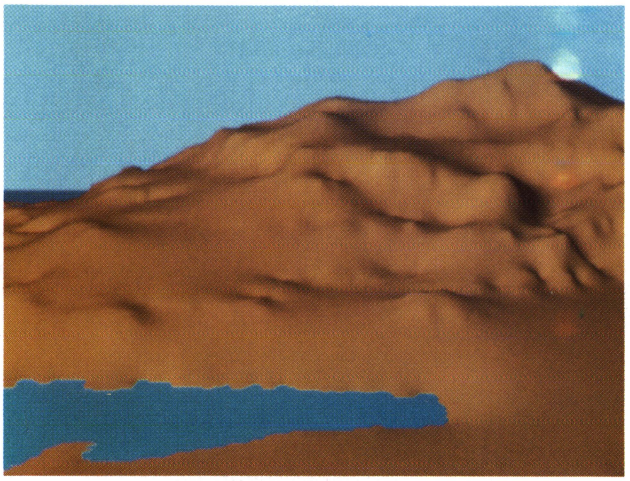

Figure 9: Interpolated thin plate under tension surface

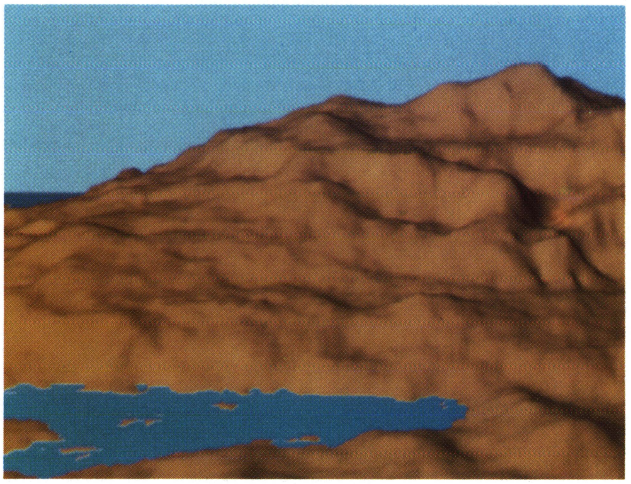

Figure 10: Fractal surface (thin plate under tension model)

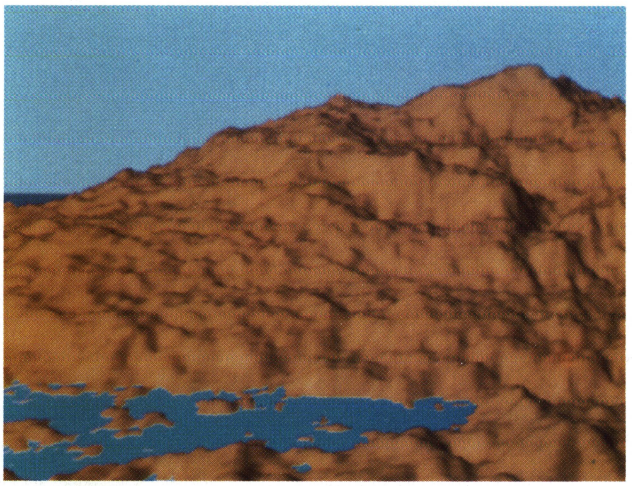

Figure 11: Fractal surface with more noise

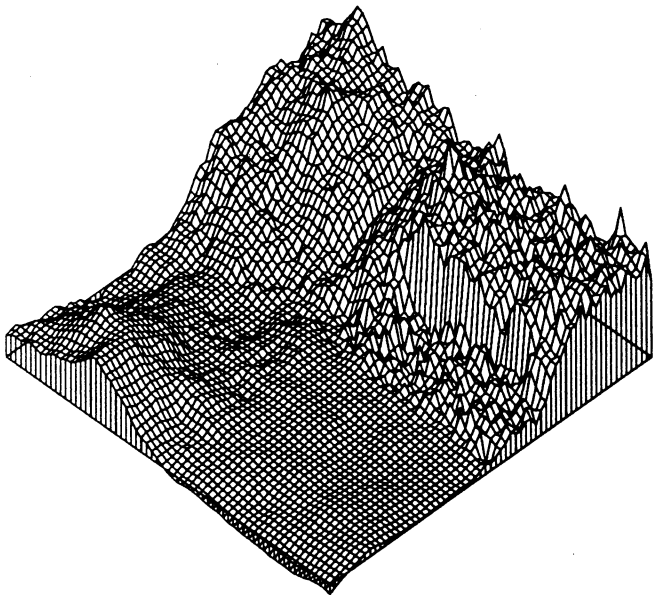

Figure 12: Constrained fractal with spatially varying fractal degree and variance

tically detailed shapes that interpolate data or match prescribed shape constraints. The deterministic component of our technique employs variational splines which can generate piecewise continuous shapes that satisfy sparse, irregular position and orientation constraints. The stochastic component of constrained fractals injects noise into the spline energy minimization procedure in order to imbue the shapes with fractal characteristics. Using the constrained fractal model, we can locally control the continuity of the underlying shape and the amount of random detail added. Such controllability makes our approach extremely flexible.

The constrained fractal computation is performed on a regular, fine-grained mesh using stochastic relaxation; hence, it is amenable to massively parallel implementation. We use a multiresolution pyramid to accelerate the convergence of the relaxation process. This coarse-to-fine procedure bears some similarity to existing recursive subdivision schemes. Unlike these other schemes, however, constrained fractals avoid the artifacts that are introduced during the subdivision process, and they can assimilate constraints at any resolution (instead of merely refining a coarse map).

The ideas developed in this paper are applicable to other deformation energy based models, such as those developed in [18,17]. By augmenting these dynamic models with a closely related method for fractal detailing, we can synthesize a greater variety of realistic, three-dimensional shapes and motions which, in principle, can be readily controlled and constrained. The generality of our approach makes it suitable for a wide range of modeling and computer graphics applications.

Acknowledgements

This research was carried out in part while the authors were at Schlumberger Palo Alto Research, Palo Alto, CA, and while Richard Szeliski was at SRI International, Menlo Park, CA. The digital elevation data was provided courtesy of Jeff Rodriguez of the University of Texas at Austin. The color images were rendered using a modeling testbed system implemented by Kurt Fleischer on Symbolics Lisp Machines.

A Multivariate Spline Interpolation

There exist natural multivariate extensions to the deformation energy functionals given in Section 3 [16]. The bivariate extension of (1) is the "thin plate under tension"

$$\mathcal{E}_{\mathrm{p}}(f) = \frac{1}{2}\iint \rho(u,v)\{[1-\tau(u,v)](f_u^2+f_v^2) + \tau(u,v)(f_{uu}^2+2f_{uv}^2+f_{vv}^2)\}\,du\,dv, \quad (16)$$

where, the $(f_u^2+f_v^2)$ gives the stretching energy density of a membrane, while the $(f_{uu}^2+2f_{uv}^2+f_{vv}^2)$ term gives the bending energy density of a thin plate. As in the univariate case, the rigidity and tension functions can be used to introduce discontinuities in position and orientation. The bivariate extension of the data compatibility constraint (2) is

$$\mathcal{E}_{\mathrm{d}}(f;\{p_i\}) = \frac{1}{2}\sum_i c_i\,(f(u_i,v_i)-d_i)^2\,. \quad (17)$$

The functional (16) is a second-order instance of the general d-variate, p-order controlled-continuity spline model

$$\mathcal{E}_{\mathrm{p}}(f) = \frac{1}{2}\sum_{m=0}^{p}\int w_m(\mathbf{u})\sum_{j_1+\cdots+j_d=m}\frac{m!}{j_1!\cdots j_d!}\left|\frac{\partial^m f(\mathbf{u})}{\partial u_1^{j_1}\cdots\partial u_d^{j_d}}\right|^2 d\mathbf{u} \quad (18)$$

where $\mathbf{u}$ is the d-dimensional domain of f. This more general formulation allows us to specify interpolators of arbitrary smoothness. A generalized version of the data compatibility term is

$$\mathcal{E}_d = \frac{1}{2}\int c(\mathbf{u})|f(\mathbf{u})-d(\mathbf{u})|^2\,d\mathbf{u} \quad (19)$$

where $d(\mathbf{u})$ and $c(\mathbf{u})$ are now continuous functions.

B Fractal Nature of Prior Model

By taking the Fourier transform of the function $f(\mathbf{u})$ and expressing the energy equations in the frequency domain, we can analyze the filtering behavior and the spectral characteristics of the spline model.

Using Rayleigh's energy theorem we can rewrite $\mathcal{E}_p(f)$ in terms of the Fourier transform $F(\boldsymbol{\omega}) = \mathcal{F}\{f(\mathbf{u})\}$ to obtain the new energy function $\mathcal{E}'_p(F)$. Using the general form in (18) with the simplifying assumption that the

weighting functions are constant, $w_m(\mathbf{u}) = w_m$,[7] we obtain

$$\mathcal{E}'_p(F) = \frac{1}{2}\int |G(\boldsymbol{\omega})|^2 |F(\boldsymbol{\omega})|^2 d\boldsymbol{\omega} \tag{20}$$

where

$$|G(\boldsymbol{\omega})|^2 = \sum_{m=0}^{p} w_m |\boldsymbol{\omega}|^{2m}. \tag{21}$$

For the membrane interpolator, $|G(\boldsymbol{\omega})|^2 \propto |\boldsymbol{\omega}|^2$ and for the thin plate model, $|G(\boldsymbol{\omega})|^2 \propto |\boldsymbol{\omega}|^4$.

We note that since the Fourier transform is a linear operation, if $f(\mathbf{u})$ is a random variable with a Boltzmann distribution with energy $\mathcal{E}_p(f)$, then $F(\boldsymbol{\omega})$ is a random variable with a Boltzmann distribution with energy $\mathcal{E}'_p(F)$. Thus, $p(F)$ is proportional to $\exp\left(-\frac{1}{2}\int |G(\boldsymbol{\omega})|^2 |F(\boldsymbol{\omega})|^2 d\boldsymbol{\omega}\right)$ from which we see that the probability distribution at any frequency $\boldsymbol{\omega}$ is

$$p(F(\boldsymbol{\omega})) \propto \exp\left(-\frac{1}{2}|G(\boldsymbol{\omega})|^2 |F(\boldsymbol{\omega})|^2\right).$$

Therefore, $F(\boldsymbol{\omega})$ is a random Gaussian variable with variance $|G(\boldsymbol{\omega})|^{-2}$, and the signal $f(\mathbf{u})$ is correlated Gaussian noise with a spectral distribution

$$S_f(\boldsymbol{\omega}) = |G(\boldsymbol{\omega})|^{-2}. \tag{22}$$

From this analysis, we can conclude that using a controlled-continuity spline is equivalent to using a correlated Gaussian field as the Bayesian prior. The spectral characteristics of this Gaussian field are determined by the choice of spline parameters. For the membrane and the thin plate models, we have

$$S_{\text{membrane}}(\boldsymbol{\omega}) \propto |\boldsymbol{\omega}|^{-2} \tag{23}$$

and

$$S_{\text{thin-plate}}(\boldsymbol{\omega}) \propto |\boldsymbol{\omega}|^{-4}. \tag{24}$$

These equations are interesting because they correspond in form to the spectra of Brownian fractals, which can be characterized by the power law

$$S_v(\omega) \propto \omega^{-\beta}. \tag{25}$$

This spectral density characterizes a fractal Brownian function $v_H(u)$ with $2H = \beta - E$, whose fractal dimension is $D = E + 1 - H$ (where E is the dimension of the Euclidean space) [19]. Comparing (23) or (24) to (25), we can conclude that a random sample drawn from a Boltzmann distribution constructed using the energy model of a membrane or a thin plate is indeed fractal [13].

[7] While this assumption does not strictly apply to general piecewise continuous interpolation, it provides an approximation to its local behavior away from boundaries and discontinuities.

References

[1] J. H. Ahlberg, E. N. Nilson, and J. L. Walsh. *The Theory of Splines and their Applications.* Academic Press, New York, 1967.

[2] R. H. Bartels, J. C. Beatty, and B. A. Barsky. *An Introduction to Splines for use in Computer Graphics and Geometric Modeling.* Morgan Kaufmann, Los Altos, CA, 1987.

[3] T. A. Foley. Weighted bicubic spline interpolation to rapidly varying data. *ACM Transactions on Graphics*, 6(1):1–18, January 1987.

[4] A. Fournier, D. Fussel, and L. Carpenter. Computer rendering of stochastic models. *Communications of the ACM*, 25(6):371–384, 1982.

[5] H. Fuchs, Z. M. Kedem, and S. P. Uselton. Optimal surface reconstruction from planar contours. *Communications of the ACM*, 20(10):693–702, October 1977.

[6] D. Geman and Hwang C.-R. Diffusions for global optimization. *SIAM Journal of Control and Optimization*, 24(5):1031–1043, September 1986.

[7] S. Geman and D. Geman. Stochastic relaxation, Gibbs distribution, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6):721–741, November 1984.

[8] B. K. P. Horn. *Robot Vision.* MIT Press, Cambridge, Massachusetts, 1986.

[9] J. P. Lewis. Generalized stochastic subdivision. *ACM Transactions on Graphics*, 6(3):167–190, July 1987.

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

[11] D. G. Schweikert. An interpolation curve using spline in tension. *J. Math. and Physics*, 45:312–317, 1966.

[12] R. Szeliski. *Bayesian Modeling of Uncertainty in Low-Level Vision.* PhD thesis, Carnegie Mellon University, August 1988.

[13] R. Szeliski. Regularization uses fractal priors. In *AAAI-87: Sixth National Conference on Artificial Intelligence*, pages 749–754, Morgan Kaufmann Publishers, Seattle, Washington, July 1987.

[14] R. Szeliski and D. Terzopoulos. Constrained fractals using stochastic relaxation. Submitted to *ACM Transactions on Graphics*, 1989.

[15] D. Terzopoulos. Multilevel computational processes for visual surface reconstruction. *Computer Vision, Graphics, and Image Processing*, 24:52–96, 1983.

[16] D. Terzopoulos. Regularization of inverse visual problems involving discontinuities. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(4):413–424, July 1986.

[17] D. Terzopoulos and K. Fleischer. Deformable models. *The Visual Computer*, 4(6):306–331, December, 1988.

[18] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer. Elastically deformable models. *Computer Graphics (SIGGRAPH'87)*, 21(4):205–214, July 1987.

[19] R. F. Voss. Random fractal forgeries. In R. A. Earnshaw, editor, *Fundamental Algorithms for Computer Graphics*, Springer-Verlag, Berlin, 1985.