SIAM/ASA J. UNCERTAINTY QUANTIFICATION
Vol. xx, pp. x

Power-Law Noises over General Spatial Domains and on Non-Standard Meshes*

Hans-Werner van Wyk†, Max Gunzburger †, John Burkhardt †, and Miroslav Stoyanov ‡

Abstract. Power-law noises abound in nature and have been observed extensively in both time series and spatially varying environmental parameters. Although, recent years have seen the extension of traditional stochastic partial differential equations to include systems driven by fractional Brownian motion, spatially distributed scale-invariance has received comparatively little attention, especially for parameters defined over non-standard spatial domains. This paper discusses the extension of power-law noises to general spatial domains by outlining their theoretical underpinnings as well as addressing their numerical simulation on arbitrary meshes. Three computational algorithms are presented for efficiently generating their sample paths, accompanied by numerous numerical illustrations.

Key words. Gaussian random fields, fractional Laplacian, fractional Brownian surfaces, power laws, self-similarity

AMS subject classifications. 65C02, 60G60

1. Introduction.

While the results of numerical simulations of physical systems may depend profoundly on the underlying model parameters, these can often only be observed partially and indirectly and arise within complex environments that cannot be described in full deterministic detail. In these cases, parameters are more appropriately modeled as random variables, -processes, or -fields. To be realistic, statistical models of these parameters should ideally not only be consistent with available measurements, but also incorporate broader, more qualitative information. Scale-invariance is one such property that has been observed widely both in natural time series, as well as spatially varying random fields. Traditionally, however, random perturbations appearing in stochastic differential equations (SDE's) or stochastic partial differential equations (SPDE's) were almost exclusively modeled as white noise. Composed in equal parts of random fluctuations at all length scales, white noise is relatively simple to generate and well understood, leading to its disproportionate use in light of available experimental evidence. Over the past two decades, a concerted effort was made to extend both the analysis and simulation of solutions of SDE's and SPDE's to include more realistic statistical models for *time-varying* parameters that exhibit scale invariance, such as fractional Brownian motion (fBm) [6]. Scale-invariant random parameters over *spatial domains* on the other hand have received comparatively little attention. They are difficult to analyze and simulate, partly due to fact that their covariance function can often not be expressed explicitly. This paper describes the theoretical underpinnings of scale-invariant-, or power-law noises and their numerical approximations, over arbitrary spatial domains, and proposes three computational algorithms to simulate

*Research partially supported by the US Department of Energy Advance Simulation Computing Research (ASCR) program under grant number DE-SC0010678.

†Department of Scientific Computing, Florida State University
‡Computational and Applied Mathematics Group, Oak Ridge National Laboratory

their sample paths over arbitrary meshes.

Power-law noise refers to a class of statistical signals whose power spectral density S, a measure of the power carried by the signal per unit frequency, satisfies the power law $S(\xi) \propto 1/\xi^{\alpha}$ for some $\alpha \geq 0$ and frequency $\xi > 0$ within some range $[\xi_{\min}, \xi_{\max}]$. Since the first description of so called 'pink noise', following an experiment designed to test Schottky's theory of 'shot noise' in vacuum tubes (see [28], [42]), power spectral densities exhibiting a power law decay have been observed experimentally in time series related to the voltages of diodes and transistors, the resistance of semiconductors and thin films, the average seasonal temperature, annual rainfall, traffic flows, action potentials in neurons, and the timbre of musical notes, to name but a few (see review articles [29],[31],[39]). The ubiquitous presence of power-law noise processes in such a wide range of natural- and man-made phenomena has led some researchers to suggest the existence of an underlying, characterizing statistical model for this type of noise that is independent of any specific physical mechanism producing the noisy signal. The development of such a statistical model would not only provide a foundation for understanding power-law noises, but would also contribute to the establishment of numerical algorithms for simulating their sample paths, which in turn can be used in aid of stochastic simulations. In light of the prevalence of power-law noises in nature, replacing more traditional white noise fields with power-law fields is likely to improve the validity of stochastic models and hence their predictive ability [45].

Although a canonical model remains elusive, the past century has seen a wealth of research devoted to the statistical description of power-law noise processes with a variety of different formulations, ranging from Poisson process models [36], fractional Brownian motion [3, 37], stochastic differential equation models [30], fractional differencing and ARFIMA (autoregressive fractional integration moving average) models [24], to approaches based on wavelets [49, 16] and system theory [29]. A central theme in the literature on $1/\xi^{\alpha}$-noise is that of scale invariance or statistical self-similarity, referring to the fact that essential statistical features remain unchanged as length scales vary. Indeed, for $1/\xi^{\alpha}$ power spectral densities, a change in the power per unit frequency due to a scaling of the frequency (i.e. 'zooming in') can be undone by simply scaling the signal itself by an appropriate constant. Since scale invariance occurs not only in *time* varying stochastic processes, but can also be observed in numerous *spatially* distributed phenomena, such as landscapes [38], glacial surface characteristics [1], fracture formation [23] and -surfaces [8], interface growth and roughening phenomena [2] (see also Barkhausen noise [43]), as well as various geophysical structures [40], including subsurface flow and transport parameters [46], it is useful to extend statistical models of power-law noise to more general domains.

Strictly speaking, no physical signal obeys the power-law $1/\xi^{\alpha}$ over its entire frequency spectrum $[0, \infty)$, since this would have unphysical implications for its energy, computed as the integral of its power spectral density over its frequency spectrum. For values of $\alpha \geq 1$, the energy in a $1/\xi^{\alpha}$-noise signal defined over the frequency range $[\xi_{\min}, \xi_{\max}]$ diverges as $\xi_{\min} \to 0^{+}$, as does that in the power-law noise signal with $0 \leq \alpha < 1$ when $\xi_{\max} \to \infty$.

In practice, restrictions on both the length of the signal and the sampling frequency, limit physical observations to a frequency interval bounded away from both zero and infinity. Periodogram estimates of S often look flat in the low frequency range, while taking the form $S(\xi) \propto 1/\xi^2$ for high frequencies (see [39]), although evidence of the persistence of power laws over large time periods, also known as infra-red divergence, has been reported in [31]. In light of the variety of physical systems that exhibit scale invariance, it is therefore important to choose a colored noise model that is both compatible with observations of the system as well as the scientific needs of the modeler. In the simulations of long memory processes such as network traffic trace processes [48] for example, it is important for the signal to adhere to the power law $1/\xi^\alpha$ in the low frequency spectrum, while numerical studies of local statistical self-similarity tend to focus on the high frequency range.

This paper discusses the numerical simulation of power-law noise fields over general multidimensional regions. We restrict ourselves to centered Gaussian noise fields, since these are used widely in practice and are uniquely determined by their covariance function. There are multiple possibilities for defining power laws over higher dimensional frequency domains, some of which are illustrated in Figure 1. The field may obey a different power law in each of its components (see Figures 1a and 1b), or the power law may be written in terms of the frequency vector's radial distance from the origin (see Figure 1c). Other forms are also possible. The fundamental challenge in generating power-law noises, even in one dimension, stems from the difficulty of translating requirements on the signal's power spectral density, i.e. the requirement $S(\xi) \propto 1/\xi^\alpha$, into quantifiable properties that are useful for construction of the signal, in this case the form of the covariance. This is partly due to the fact that different values of α give rise to signals with widely divergent statistical properties. While the Wiener-Khinchine Theorem [47, 32] relates auto-covariance functions of stationary processes to their power spectral densities via Fourier transform, this assumption does not hold for all $\alpha \geq 0$. For an intuition of the correlation structure of a general power-law noise process, it is useful to consider again the signal's energy. For small values of α, high frequencies contribute significantly to the total energy, suggesting erratic, uncorrelated behavior and hence stationarity, the extreme case of which is white noise ($\alpha = 0$) where all frequencies contribute equally to the total energy. As α increases, the contribution of the high frequency components is diminished, compared to that of the low frequencies and the signal tends to become smoother, more correlated, as well as non-stationary (see Brownian noise, $\alpha = 2$).

Considering the above-mentioned subtleties, many colored noise models are formulated by directly specifying the covariance structure of the field, instead of deducing its form based on specifications of the power law. Fractional Brownian motions, as well their multi-dimensional analogues, fractional Brownian surfaces and -sheets, form a widely used class of such noise models. In a sense they can be said exhibit power law behavior [15], in addition to having other desirable properties, such as Hölder continuous sample paths and stationary increments. Moreover, efficient algorithms, such as the circulant embedding method [10, 12], have been adapted [44] to generate fast, exact numerical simulations of fractional Brownian surfaces on equispaced rectangular grids defined over certain standard

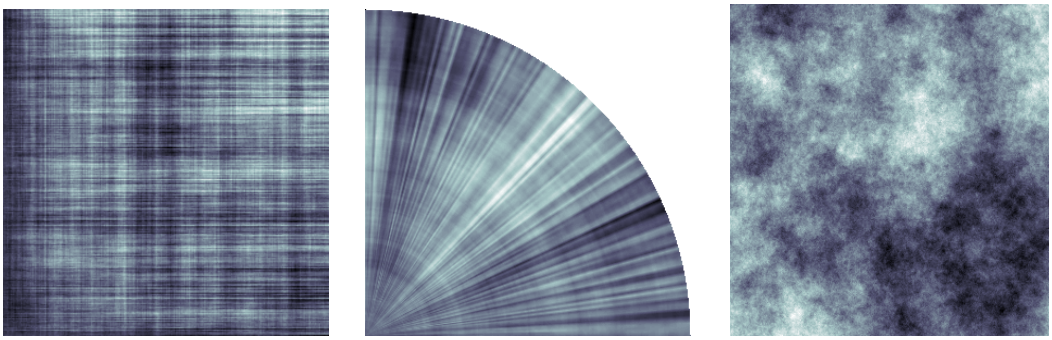

(a) Fractional Brownian sheet over a square with power law $S(\xi_1, \xi_2) = \xi_1^{-\alpha_1}\xi_2^{-\alpha_2}$.

(b) Fractional Brownian sheet over a quarter disc with power law $S(\xi_r, \xi_\theta) = \xi_r^{-\alpha_1}\xi_\theta^{-\alpha_2}$.

(c) Fractional Brownian surface over a square with power law $S(\xi_1, \xi_2) = \|\xi\|^{-\alpha}$.

Figure 1: Different power law noises over two dimensional domains

domains, such as rectangles or circles (see Figures 2a and 2b). The covariance function of the standard fractional Brownian surface, however, does not depend on the region's underlying geometry, which can lead to unrealistic models for fields defined over non-convex regions (see Figure 2c and Subsection 3.1). Moreover, if the field represents the input for a complex physical system that needs to be solved numerically, the meshes imposed on the region are often non-uniform for the sake of computational expediency. Depending on the application, it may also be more appropriate for the field to be stationary, in which case the power law no longer holds in the low frequency range. Elliptic Gaussian fields [5] generalize fractional Brownian motion, based on its spectral characterization in terms of the fractional Laplace operator. They are self-similar random fields (at least locally) and have Hölder continuous sample paths. Furthermore, this theoretical framework allows for the definition of colored noise fields with prescribed values in certain regions, and has even been used to define colored noise fields over manifolds ([18, 19]), by means of the Laplace-Beltrami operator.

After some preliminary remarks and establishing notation in Section 2), we review fractional Brownian surfaces and their generalizations, the elliptic Gaussian fields over Euclidean space $\mathbb{R}^d$ in Section 3. Section 4 treats the Riesz fields, generalizations of elliptic Gaussian fields to arbitrary bounded regions. The numerical simulation of Riesz fields is the subject of Section 5. We discuss numerical simulations based on the approximation of the discretized fractional Laplace operator, either through the eigen-decomposition of the finite element Laplace matrix or the contour integral method [9], as well as simulations based on a modified form of the Riesz potential, which allows for the simulation of multi-fractional fields for which the level of roughness may vary throughout the domain. Numerical illustrations accompany the discussion throughout. For the sake of visualization, all of our computational results are based on regions in $\mathbb{R}^2$. The algorithms discussed,

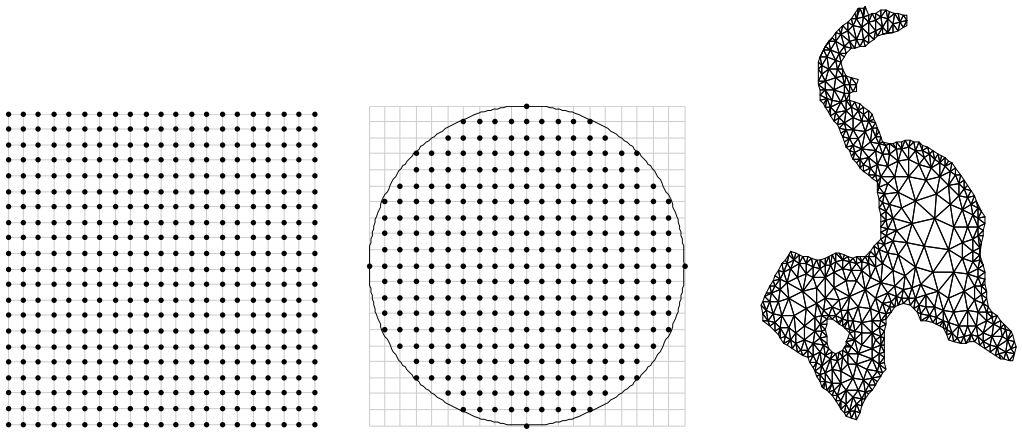

(a) A regular grid on a rectangular region

(b) A regular grid on a circular region

(c) A finite element mesh over a general region

Figure 2: Colored noise can be synthesized efficiently on regular grids over simple domains. However, non-uniform meshes over general domains are often preferred.

however, extend readily to three spatial dimensions. Finally, Section 6 contains concluding remarks. Since this paper focuses on the generation of sample paths over general regions, we do not discuss here the construction and analysis of scale-invariant noises by means of wavelets (see [16] and [49]).

2. Notation and Preliminaries. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space and $D \subset \mathbb{R}^d$ a bounded region acting as the index set for our random field. Throughout, we use X and Y to denote generic random fields, while $\mathbf{X}$ and $\mathbf{Y}$ represent random vectors. A Gaussian random field $\{X(x)\}_{x \in D}$ is a collection of random variables so that for any finite subset $\{x_i\}_{i=1}^n \subset D$, the vector $\mathbf{X} = [X(x_1), ..., X(x_n)]^T$ has a joint Gaussian distribution. A random field is said to be centered if its expectation $\mathbb{E}[X(x)] = 0$ for all $x \in D$, an assumption we make throughout for the sake of simplicity. The covariance function $C_X : D \times D \to \mathbb{R}$ of a centered random field X is defined pointwise for any two points x and y by the covariance $C_X(x, y) := \mathrm{Cov}(X(x), X(y)) = E[X(x)X(y)]$. Fractional Brownian motions as well as -surfaces with Hurst parameter H are denoted by B_H. White noise fields B_0 are mostly denoted by W, especially when referring to the isotropic white noise process. The random vector $\mathbf{Z} = [Z_1, ..., Z_n]^T$ always represents a vector of identically distributed (i.i.d.) Gaussian random variables. To prevent unnecessary ambiguity, we use the variables t or s to index stochastic processes and x or y to index random fields over higher dimensional regions.

The power spectral density of a deterministic signal $\{x(t)\}_{t \geq 0}$ at a given frequency ξ, is simply the magnitude of its Fourier transform squared, i.e. $S(\xi) = |\mathcal{F}(x)(\xi)|^2$, where $\mathcal{F}(x)(\xi) := \int_0^\infty x(t) e^{-it\xi}\, dt$. Since the Fourier transform of a random signal $\{X(t)\}_{t \geq 0}$ over

$[0,\infty)$ does not exist for many signals of practical importance, the power spectral density is often defined via the truncated Fourier transform

$$\mathcal{F}_T(X)(\xi,\omega) := \frac{1}{\sqrt{T}}\int_0^T X(t,\omega)e^{-i\xi t}\,dt, \quad \text{for } \omega\in\Omega \text{ and } \xi\geq 0,$$

and hence

$$S(\xi) := \lim_{T\to\infty} \mathbb{E}\left[|\mathcal{F}_T(X)(\xi)|^2\right].$$

In our numerical calculations, we approximate the power spectrum by simply computing the discrete Fourier transform of its sample paths over a finite domain, using fast Fourier transforms, and averaging the square of their magnitudes. In multi-dimensional domains in $\mathbb{R}^d$, we use the multi-dimensional fast Fourier transform.

For a stationary time series $\{X(t)\}_{t\in\mathbb{R}}$, the Wiener-Khintchine Theorem can be used to directly relate its power spectrum to its auto-covariance function $\rho(\tau) := \mathrm{Cov}(X(t), X(t+\tau))$ via the Fourier transform, i.e. $S(\xi) = \hat{\rho}(\xi) := \int_{-\infty}^{\infty}\rho(\tau)e^{-i\xi\tau}\,d\tau$. Although no equivalent relation exists for non-stationary processes $\{X(t)\}_{t\in\mathbb{R}}$, the time dependent Wigner-Ville spectrum

$$S_X^{WV}(t,\xi) := \int_{-\infty}^{\infty}\mathrm{Cov}\left(X\left(t+\frac{\tau}{2}\right), X\left(t-\frac{\tau}{2}\right)\right)e^{-i\xi\tau}\,d\tau,$$

presents a convenient generalization. In [15], this version of the power spectrum is used to show that fractional Brownian motion is a $1/\xi^{\alpha}$-process.

It is often convenient to define a Gaussian random field as the convolution of a deterministic kernel function with white noise. To make this construction rigorous, we make use of the Skorokhod integral with respect to the isotropic white noise process W, defined on Hilbert space H as the mapping $W : H \to L^2(\Omega)$ so that i) for any $h\in H$, $W(h)\sim N(0,\|h\|^2)$ and ii) for any $h_1,h_2\in H$, $\mathbb{E}[W(h_1)W(h_2)] = \langle h_1,h_2\rangle_H$ (see [41]). Let $\mathcal{S}$ be the set of H-valued random variables of the form $F = f(W(h_1),...,W(h_n))$, where $h_1,...,h_n\in H$ and $f:\mathbb{R}^n\to\mathbb{R}$ is infinitely differentiable with partial derivatives that grow at most polynomially at infinity. Further, denote by $D(F)$ the Malliavin derivative of F, i.e. $D(F) = \sum_{i=1}^n \partial_i f(W(h_1),...,W(h_n))h_i$, which can be extended to a closed unbounded operator $D : L^2(\Omega)\to L^2(\Omega,H)$ with domain $\mathbb{D}^{1,2}$, the Sobolev-Watanabe space. Since $\mathbb{D}^{1,2}$ is dense in $L^2(\Omega)$, there exists a unique adjoint operator $\delta : \mathrm{dom}(\delta)\to L^2(\Omega)$ defined for each $u\in\mathrm{dom}(\delta)$ by the relation

$$\mathbb{E}[F\delta(u)] = \mathbb{E}[\langle D(F),u\rangle_H] \quad \text{for all } F\in\mathbb{D}^{1,2},$$

where $\mathrm{dom}(\delta) = \{u\in L^2(\Omega,H) : F\mapsto\mathbb{E}[\langle D(F),u\rangle_H] \text{ is bounded }\}$. This operator is also known as the Skorokhod integral, written as $\int u\,dW := \delta(u)$. For deterministic integrands $u\in H$, the isometry $\|\int u\,dW\|_{L^2(\Omega,H)} = \|u\|_H$ allows these integrals to be approximated by integrals of simpler functions. Let $V^n\subset H$ be a finite dimensional subspace, spanned

by basis functions $\{\phi_i\}_{i=1}^n$. Then

$$\int \sum_{i=1}^n c_i \phi_i \, dW = \sum_{i=1}^n c_i Z_i, \text{ where } [Z_1, ..., Z_n] \sim N(0, \Sigma), \text{ with } \Sigma_{ij} = \langle \phi_i, \phi_j \rangle_H$$

for $i, j = 1, ..., n$. In particular, when this basis consists of simple functions over a bounded domain, then the isotropic white noise process W takes the form of the white noise measure, defined on the set of Borel measurable subsets of $\mathbb{R}^d$ with finite volume (see [26]). For any such Borel set A with volume $|A| < \infty$, $W(A) \sim N(0, |A|)$, and for any finite collection $A_1, ..., A_n$ of disjoint sets, the $W(A_i)$ are independent and $W(\cup_{i=1}^n A_i) = \sum_{i=1}^n W(A_i)$. In this case, the white noise integral of a simple function $f_s(x) := \sum_{i=1}^n c_i \chi_{A_i}(x)$ $\int f_s(x) \, dW(x) := \sum_{i=1}^n c_i W(A_i)$.

3. Models for Power-Law Noises over $\mathbb{R}^d$. In certain simple cases, statistical models for power-law noises can be developed directly, based on specifications of their desired properties. In this section, we introduce the fractional Brownian motion (fBm) over the interval $[0, \infty)$ as the unique self-similar Gaussian process with stationary increments. Although the self-similarity required by Definition 3.2 does not involve the power spectral density, it can be shown that the Wigner-Ville spectra of fBm's do indeed exhibit a power law decay. The covariance function of fBm can directly be generalized to Euclidean space, giving rise to the fractional Brownian surface (fBs) over $\mathbb{R}^d$. This extension is too rigid, however, rendering fBs unsuitable to model spatially correlated noise in many practical applications. Their spectral characterization nevertheless identifies them as members of the family of elliptic Gaussian fields, Gaussian fields defined in terms of pseudo-differential operators with positive symbol. Using these operators (and in particular the fractional Laplace operator) as the basis for the generalization of power-law noises gives rise to models suitable for arbitrary domains that share the salient properties of fBs, such as Hölder continuity, while also allowing for additional flexibility, such the imposition of boundary conditions.

3.1. Fractional Brownian Motion. Fractional Brownian motion, a family of Gaussian random processes parameterized by the Hurst parameter $H \in (0, 1)$, has become a particularly widespread model for time dependent power-law noises. In the seminal paper [37], the authors introduce scale invariance directly through the definition of self-similarity with respect to a Hurst parameter H.

Definition 3.1. *A random process $X : [0, \infty) \to \mathbb{R}$ is said to be self-similar with Hurst parameter $H \in (0, 1)$ (H-s.s.) if for every $c > 0$, we have $X(ct) \overset{d}{=} c^H X(t)$, where '$\overset{d}{=}$' denotes equality in distribution.*

This definition generalizes the well-known self-similarity property of Brownian motion, whose Hurst parameter $H = 1/2$. For any $t \in [0, \infty)$, let $\triangle X(t, h) = X(t+h) - X(t)$ denote the h-increment process related to X. If in addition to H-s.s., we assume that $X(0) = 0$ almost surely (a.s.) and that the increments $\triangle X(t, h)$ are stationary (i.e. the distribution of the increment depends only on h), then we automatically arrive at the following explicit

form for its covariance function, namely

$$\begin{aligned}\mathbb{E}\left[X(t)X(s)\right] &= \frac{1}{2}\left(\mathbb{E}\left[X(t)^2\right] + \mathbb{E}\left[X(s)^2\right] - \mathbb{E}\left[\triangle X(s,t-s)^2\right]\right)\\ &= \frac{1}{2}\left(\mathbb{E}\left[X(t)^2\right] + \mathbb{E}\left[X(s)^2\right] - \mathbb{E}\left[(X(t-s))^2\right]\right)\\ &= \frac{1}{2}\left(t^{2H}\mathbb{E}\left[X(1)^2\right] + s^{2H}\mathbb{E}\left[X(1)^2\right] - (t-s)^{2H}\mathbb{E}\left[X(1)^2\right]\right)\\ &= \frac{\mathbb{E}\left[X(1)^2\right]}{2}(t^{2H} + s^{2H} - |t-s|^{2H}),\end{aligned}$$

for any points $s, t \in [0, \infty)$. The above equations also serve to prove the converse, namely that any zero mean stochastic process with this covariance matrix necessarily has stationary increments and is H-s.s.. For simplicity, we assume henceforth that $\mathbb{E}\left[X(1)^2\right] = 1$. The definition of fractional Brownian motion is given by the following.

Definition 3.2 (Fractional Brownian Motion). *A fractional Brownian motion $B_H(t)$ with Hurst parameter $H \in (0,1)$ is a continuous and centered Gaussian process, i.e. $B_H(0) = 0$ and $\mathbb{E}\left[B_H(t)\right] = 0$ for $t \geq 0$, with covariance function given by*

$$C_{B_H}(s,t) := \mathbb{E}\left[B_H(s)B_H(t)\right] = \frac{1}{2}\left(s^{2H} + t^{2H} - |t-s|^{2H}\right), \quad t, s \geq 0. \tag{3.1}$$

Although fractional Brownian motion is not mean square differentiable, it does admit a Hölder continuous modification [6]. Moreover, fBm is not a martingale for $H \neq \frac{1}{2}$, and since its variance $\mathrm{Var}(B_H(t)) = \mathbb{E}\left[B_H(t)^2\right] = t^{2H}$ for $t \geq 0$, it is non-stationary.

In [15], it was shown that the power spectrum $S_{B_H}(\xi)$ of fractional Brownian motion, defined in terms of the Wigner-Ville spectrum, satisfies $S_{B_H}(\xi) \propto 1/\xi^{2H+1}$, suggesting that the Hurst parameter H can be related to the power $\alpha > 0$ via $\alpha = 2H + 1$. A Hurst parameter $H \in (0,1)$ thus gives rise to power-law noises with $\alpha \in (1,3)$. Colored noise with $\alpha \in [0,1)$ can also be defined by making use of the increment process $\triangle B_H(t,h)$, also known as fractional Gaussian noise (fGn). In fact, it can be shown that the power spectrum of the increment process satisfies $S_{\triangle B_H}(\xi) \propto 1/\xi^{2H-1}$, independent of t, so that fractional Gaussian noises are power-law noises with $\alpha \in [0,1)$. The fGn process, while still correlated in general, is nevertheless stationary. Indeed, for $t, h \geq 0$ and $s = t + nh \geq 0$ for some $n \geq 0$, it can be shown that

$$\begin{aligned}\mathrm{Cov}(\triangle B_H(t,h), \triangle B_H(s,h)) &= \mathbb{E}\left[\triangle B_H(t,h)\triangle B_H(s,h)\right]\\ &= \frac{1}{2}h^{2H}[(n+1)^{2H} + (n-1)^{2H} - 2n^{2H}],\end{aligned}$$

which is independent of $s, t \geq 0$.

The covariance function defined in (3.1) readily suggests an extension of fractional Brownian motion over the index set $\mathbb{R}^d$.

Definition 3.3 (Fractional Brownian Surface). *A fractional Brownian surface $\{B_H(x)\}_{x\in\mathbb{R}^d}$ with Hurst parameter $H \in (0,1)$ and centered at the origin is defined to be the continuous,*

centered Gaussian random field with covariance function

$$C_{B_H}(x,y) = \mathbb{E}\left[B_H(x), B_H(y)\right] := \frac{1}{2}\left(\|x\|^{2H} + \|y\|^{2H} - \|x-y\|^{2H}\right) \tag{3.2}$$

The fractional Brownian surface retains most of the attractive features of fractional Brownian motion. It is self-similar in the sense that the field $\{B_H(cx)\}_{x\in\mathbb{R}^d}$ agrees with $\{c^H B_H(x)\}_{x\in\mathbb{R}^d}$ in law, for any scaling factor $c > 0$. Moreover, the covariance function (3.2) determines a probability measure over the set of Hölder continuous functions of degree H. The fBs defined above can readily be modified through translation to a fBs centered at an arbitrary point $x_0 \in \mathbb{R}^d$. It can also be conditioned on available measured data if the spatial domain contains points, or even whole regions, in which the field is known exactly (see e.g. [11, 33]). In this case, the conditional covariance function no longer takes the form (3.2), but the covariance matrix of the field at a finite number of points can be calculated by means of Schur complements.

The fact that (3.2) is defined independently of the underlying index set, however, makes this power-law noise model unsuitable for general regions. While the function given by (3.2) could plausibly be employed to measure the covariance between points in $\mathbb{R}^d$, or even points in a convex sub-domain, it cannot capture the covariance structure of fields defined over general non-convex regions, such as the one depicted in Figure 3, where there is no longer a direct correspondence between the length of the shortest path between two points and their Euclidean distance.

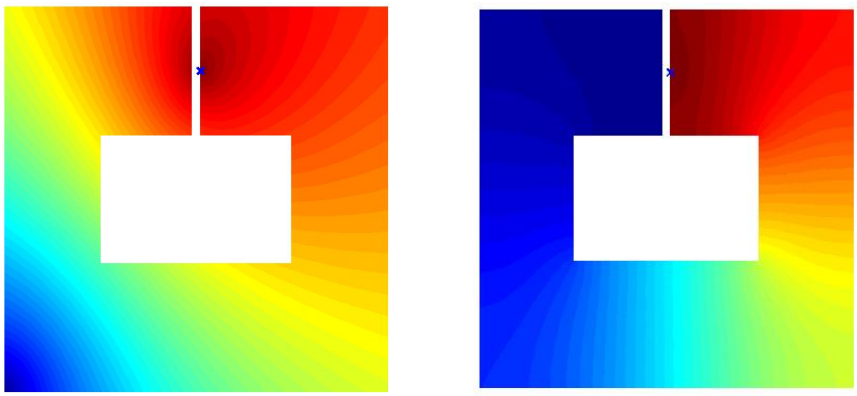

(a) Fractional Brownian Surface (b) Riesz random field

Figure 3: The covariance of colored noise fields at a point (cross).

Physically accurate colored noise models over general regions should therefore somehow incorporate the geometry of the underlying index set D. Similar issues arise in defining fractional Brownian fields over manifolds. One approach [25] is to modify the covariance function in (3.2) by replacing the Euclidean norm $\|x-y\|$ with a geodesic distance $d(x,y)$

for $x, y \in D$. It then remains to show that the resulting covariance function is positive definite, by proving for example that the metric $d(x, y)$ is of negative type. Here we follow another approach, based on the spectral characterization of fractional Brownian noise in terms of the Laplace operator on D. We begin by discussing the generation of colored noise through the fractional integration of white noise.

3.2. Hosking's Fractional Difference Model. In [24], the author introduces a discrete version of power-law noise by analogy with the random walk approximation $\{B_n\}_{n=0}^{\infty}$ of Brownian motion $\{B(t)\}_{t\geq 0}$, given by the convolution sum

$$B_n = \sum_{i=0}^{n} h_i Z_{n-i}, \ n = 1, 2, ...,$$

where $Z_i \sim N(0, 1)$ are identically distributed (iid) standard normal random variables and the discrete impulse response $h_i = 0$ for $i = 0$ and $h_i = 1$ for $i = 1, 2, ..., n$. If we let L denote the lag operator, i.e. $LB_n = B_{n-1}$ for $n = 1, 2, ...$ and $LB_0 = B_0 = 0$, then the first order difference satisfies $(1 - L)B_n = B_n - B_{n-1} = Z_n$ for $n = 1, 2, ...$ A power-law noise signal $\{B_n^{\beta}\}_{n=0}^{\infty}$ can now be defined as an ARFIMA (AutoRegressive Fractional Moving Average) process by specifying that its β^{th} fractional order difference is a white noise process, i.e.

$$(1 - L)^{\beta} B_n^{\beta} := \sum_{k=0}^{\infty} \binom{\beta}{k} (-L)^k B_n^{\beta} = Z_n, \text{ for } n = 1, 2, ...$$

The process B_n^{β} itself can therefore be regarded as a type of fractional cumulative sum of the white noise process Z_n. To obtain an explicit description of B_n^{β} as a discrete convolution with a white noise process, we note that the unilateral Z-transform H^{β} of its impulse response vector $\{h_i^{\beta}\}_{i=0}^{\infty}$ must be of the form $H^{\beta}(z) = \frac{1}{(1-z^{-1})^{\beta}}$, with $z > 1$. In [29] (see also [45]), the author arrives at the same transfer function H^{β}, by 'interpolating' between the transfer function of white noise ($\alpha = 0$), given by $H(z) = 1$, and that of Brownian noise ($\alpha = 2$), given by $H(z) = \frac{1}{1-z^{-1}}$. The parameter β is thus related to α via $\beta = \alpha/2$. The discrete impulse response function h_n^{β} can be computed by means of a simple recurrence relation and sample paths with the right power spectral decay can be generated efficiently through the use of fast Fourier transform.

The value of a Hosking colored noise process at any point of time t is thus determined by a collection of noise sources spread out over the domain (see Figure 4), whose influence decreases as their distance to t increases. Lower values of α imply steeper decline in their influence (indicated by the gray curves in 4), indicating that the influence of the noise terms are more local, while larger values of α allow the noise sources to have wider ranging influence. In anticipation of simulations of spatially varying noise, Figure 4 b) shows a modification of the Hosking process in which the noise sources are positioned on both sides of the current time point.

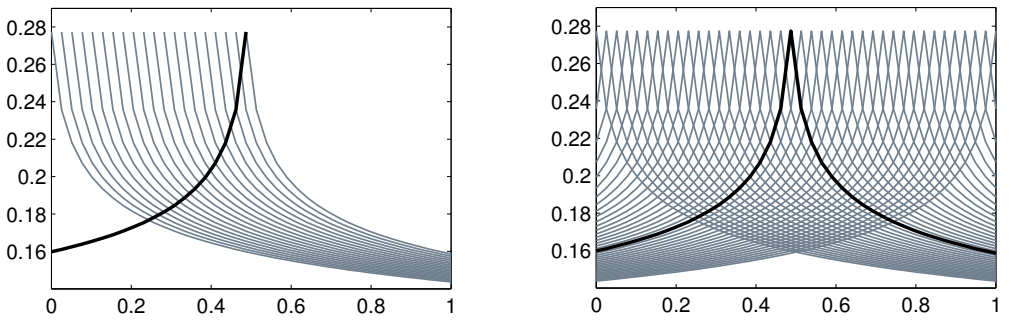

(a) Hosking Process ($\beta = 0.85$). (b) Modified Hosking Process ($\beta = 0.85$)

Figure 4: Schematic of the composition of a Hosking noise signal. The gray curves represent the influence of each white noise source term over the interval. The black curve attributes the contribution of the various noise sources to the signal's value at $t = 0.5$.

3.3. Elliptic Gaussian Fields. The idea of forming colored noise as the fractional integral of white noise dates back to 1953, when Paul Lévy [34] (see also [3]) commented on a stochastic process formed by the cumulative fractional Riemann-Liouville integral of white noise,

$$X(t) = \frac{1}{\Gamma(H + \frac{1}{2})} \int_0^t (t - s)^{H - \frac{1}{2}} dW(s),$$

where H may be any positive number. The Riesz potential I_s, defined for any locally integrable function $f : \mathbb{R}^d \to \mathbb{R}$ by

$$I_s[f](x) := c_s \int_{\mathbb{R}^d} \frac{f(y)}{\|x - y\|^{d-s}}\, dy, \quad \text{where } c_s = \frac{\Gamma(\frac{d-s}{2})}{\pi^{\frac{d}{2}} 2^s \Gamma(\frac{s}{2})}, \quad s \in (0, d) \tag{3.3}$$

is the multi-dimensional generalization of the Riemann-Liouville integral. The Fourier transform of the Riesz kernel, together with the convolution theorem reveal that $\mathcal{F}(I_s f)(\xi) = \|2\pi\xi\|^{-s}\mathcal{F}(f)(\xi)$ and hence I_s is a Fourier multiplier. In comparison, the Laplace operator satisfies $\Delta e^{2\pi i x\cdot\xi} = -4\pi^2\|\xi\|^2 e^{2\pi i x\cdot\xi}$, so that $\mathcal{F}(\Delta f)(\xi) = -4\pi^2\|\xi\|^2\mathcal{F}(f)(\xi)$. In a spectral sense, the Riesz potential therefore represents the fractional inverse $(-\Delta)^{-\frac{s}{2}}$ of the negative Laplace operator over $\mathbb{R}^d$. It is this connection with the fractional Laplace operator that allows for the development of models for power-law noise fields that preserve the essential properties of Euclidean fractional Brownian fields, such as scale invariance and the stationarity of its increments under translations and rotations, while also reflecting the underlying geometry of the index set.

More generally, let $\mathcal{A}$ be a pseudo-differential operator defined on an appropriate Hilbert

space $\mathcal{H}_{\mathcal{A}}$ of functions in terms of its Fourier transform, i.e.

$$(\mathcal{A}f)(x) := \frac{1}{(2\pi)^{\frac{d}{2}}} \int_{\mathbb{R}^d} e^{ix\cdot\xi} \sigma(x,\xi) \mathcal{F}(f)(\xi)\, d\xi,$$

with symmetric, positive symbol $\sigma : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$. In [5], the authors define a family of elliptic Gaussian fields $X_{\mathcal{A}}$ over $\mathbb{R}^d$ through the covariance function given by the integral kernel of the inverse operator $\mathcal{A}^{-1}$, over the appropriate function space, when it exists. The properties of the field X_A are directly related to those of the pair $(\mathcal{A}, \mathcal{H}_{\mathcal{A}})$. In particular, $X_{\mathcal{A}}$ has stationary increments if and only if σ does not depend on x.

Since the set of self-similar Gaussian fields over $\mathbb{R}^d$ with stationary increments is restricted [13], the authors in [5] instead investigate a local form of self similarity. A random field X is said to be locally asymptotically self-similar of order $H \in (0,1)$ at the point $x_0 \in \mathbb{R}^d$ if the limit

$$\lim_{\rho \to 0} \frac{1}{\rho^H} (X(x_0 + \rho x) - X(x_0))$$

is non-trivial in law. Evidently, this condition generalizes H-s.s. in Definition 3.1. Moreover, by essentially disregarding the low frequency range, it affords the modeler a greater degree of flexibility in adapting the noise field to both the geometry of the underlying index region D as well as enforcing other more explicit conditions, such as boundary conditions. Standard Brownian motion and the Brownian bridge have the same local scaling properties, for example, although the bridge is not strictly self-similar in the sense of Definition 3.1.

The connection between elliptic Gaussian fields and fractional Brownian surfaces becomes evident when considering the pseudo-differential operator $\mathcal{A}$ with symbol $\sigma(x,\xi) = \|\xi\|^{d+2H}$, defined on $\mathcal{D}_0(\mathbb{R}^d) := \{f \in \mathcal{D}(\mathbb{R}^d) : f(0) = 0\}$. If we let H_A be the closure of $\mathcal{D}_0(\mathbb{R}^d)$ under the inner product $\langle \mathcal{A}f, g \rangle_{L^2}$, then (according to Lemma 1.1., [5]) it forms the reproducing kernel Hilbert space for the elliptic Gaussian field defined by

$$X_{\mathcal{A}}(x) = \int_{\mathbb{R}^d} \frac{e^{ix\cdot\xi} - 1}{\|\xi\|^{\frac{d}{2}+H}}\, d\widehat{W}(\xi), \tag{3.4}$$

where $\widehat{W}$ is a complex white noise measure. Apart from a scaling constant, this definition coincides precisely with the spectral characterization of fractional Brownian surfaces [50]. Moreover, by invoking the spectral definition of the fractional Laplacian, we can interpret the action of this integral operator on any function f in the dual Sobolev space $H^{-(\frac{d}{4}+\frac{H}{2})}(\mathbb{R}^d)$, as

$$\int_{\mathbb{R}^d} \frac{e^{ix\cdot\xi} - 1}{\|\xi\|^{\frac{d}{2}+H}} \hat{f}(\xi)\, d\xi = (-\Delta)^{-(\frac{d}{4}+\frac{H}{2})} f(x) - (-\Delta)^{-(\frac{d}{4}+\frac{H}{2})} f(0).$$

The addition of the term $(-\Delta)^{-(\frac{d}{4}+\frac{H}{2})} f(0)$ ensures that the random field $\{X_{\mathcal{A}}(x)\}_{x\in\mathbb{R}^d}$ is well-defined. Sample paths of fBs can therefore be formed by solving the fractional

Laplace equation with a white noise forcing term, i.e. $X_H(x) = (-\Delta)^{-(\frac{d}{4}+\frac{H}{2})}W(x) - (-\Delta)^{-(\frac{d}{4}+\frac{H}{2})}W(0)$, where $W(x)$ is a white noise field. The covariance of B_H is then given by the integral

$$C_{X_H}(x,y) = \int_{\mathbb{R}^d} \frac{e^{i(x-y)\cdot\xi} - e^{ix\cdot\xi} - e^{iy\cdot\xi} + 1}{\|\xi\|^{\frac{d}{2}+H}}\, d\xi.$$

Apart from classical fractional Brownian surfaces, this framework can also be used to define fields corresponding to values of $H \leq 0$, the case $H = 0$ coinciding with the well-known pink noise. The resulting field's sample paths are no longer continuous, but can nevertheless be analyzed in the sense of distributions. Multi-fractional Brownian surfaces represent another interesting class of elliptic fields, whose Hurst parameter is a spatially varying function.

4. Riesz Fields over Bounded Regions. The spectral characterization of fractional Brownian surfaces described above reveals how the fractional Laplace operator forms a natural point of departure for defining power law noises over more general index sets. Not only does the operator conform to the geometry of the underlying domain, but it also allows for the imposition of boundary conditions. In [18], the author defines the so-called 'Riesz fields' over Riemannian manifolds through the use of the Laplace-Beltrami operator. He shows that Riesz fields are generalizations of fractional Brownian surfaces, that they are Hölder continuous with a Hölder coefficient related to the field's Hurst parameter, and that these fields satisfy a form of self-similarity, after accounting for the effect of the Riemannian metric.

The fractional Laplace operator has been widely studied in fields such as physics, finance, and hydrology, where it is associated with models of anomalous diffusion. In this paper, we consider homogeneous Dirichlet, Neumann or Robin boundary conditions and treat fractional powers of the Laplace operator in terms of functional calculus related to its spectral decomposition. Indeed, if $D \subset \mathbb{R}^d$ is an open connected domain with piecewise smooth boundary ∂D, then the Laplace operator $-\Delta$, subject to the aforementioned boundary conditions, has a discrete, non-negative spectrum with eigenvalues $\{\lambda_k\}_{k=0}^{\infty}$ satisfying $0 \leq \lambda_0 < \lambda_1 \leq \lambda_2 \leq ... \uparrow \infty$ and eigenfunctions $\{\psi_k\}_{k=0}^{\infty}$ that form a complete basis in $L^2(D)$. The fractional power $(-\Delta)^s$ of the Laplacian $(-\Delta)$, applied to a function $f \in L^2(D)$ can then be expressed as

$$(-\Delta)^s f(x) := \sum_{k=0}^{\infty} \lambda_k^s \langle f, \psi_k\rangle \psi_k(x)$$

To define Riesz fields over D in accordance with the spectral definition of fractional Brownian surfaces over $\mathbb{R}^d$ requires the fractional inverse $(-\Delta)^{-(\frac{d}{2}+H)}$. For Dirichlet or Robin boundary conditions, the first eigenvalue λ_0 is strictly positive so that we can define the action of this fractional inverse in terms of the series

$$(-\Delta)^{-(\frac{d}{2}+H)} f(x) = \sum_{k=0}^{\infty} \lambda_k^{-(\frac{d}{2}+H)} \langle f, \psi_k\rangle \psi_k(x) = \int_D \sum_{k=0}^{\infty} \lambda_k^{-(\frac{d}{2}+H)} \psi_k(x)\psi_k(y) f(y)\, dy, \tag{4.1}$$

where Fubini's Theorem, together with the uniformly continuous convergence of the series, allows for the interchange of integration and summation. Just as for elliptic Gaussian fields, we now define the covariance function $C_{X_H} : D \times D \to \mathbb{R}$ of our Riesz field X_H over D as the kernel in (4.1), i.e.

$$C_{X_H}(x,y) = \mathbb{E}\left[X_H(x)X_H(y)\right] = \sum_{k=0}^{\infty} \lambda_k^{-(\frac{d}{2}+H)} \psi_k(x)\psi_k(y) \tag{4.2}$$

Weyl's law, prescribing the asymptotic growth rate of the eigenvalues $\{\lambda_k\}_{k=0}^{\infty}$ along the order of $O(k^{\frac{2}{d}})$, guarantees that this series converges for any $H \in (0,1)$ and hence X_H is well-defined as a Gaussian field. The Riesz field X_H itself can then be written in the form

$$X_H(x,\omega) = \sum_{k=0}^{\infty} \lambda_k^{-(\frac{d}{4}+\frac{H}{2})} \psi_k(x) Z_k(\omega), \quad \text{where } Z_k \sim N(0,1) \text{ i.i.d..} \tag{4.3}$$

For homogeneous Neumann boundary conditions, $\lambda_0 = 0$ and consequently the fractional inverse $(-\Delta)^{-(\frac{d}{4}+\frac{H}{2})}$ is not defined. Since the corresponding eigenfunction ψ_0 is constant, however, the field X_H can be nevertheless be constructed by letting

$$X_H(x) = \sum_{k=0}^{\infty} \lambda_k^{-(\frac{d}{4}+\frac{H}{2})} (\psi_k(x) - \psi_k(x_0)) Z_k, \tag{4.4}$$

where $x_0 \in D$ is some point serving as the origin. The covariance function is defined accordingly. This modification effectively eliminates the zeroth term in the series and therewith the singularity and amounts to imposing $X_H(x_0) = 0$. The same modification appears in the spectral definition (3.4) of fractional Brownian surfaces. Another possibility is to simply leave out the 0^{th} mode, resulting in a field that differs from the one above by an additive constant. In the special case when $D = [0,1] \subset \mathbb{R}$ and $H = \frac{1}{2}$, Definition (4.3) amounts to the well-known Fourier expansion of Brownian motion, if Dirichlet boundary conditions are imposed at $t = 0$ and Neumann conditions at $t = 1$, whereas letting Dirichlet conditions hold at both endpoints represents the Fourier expansion of the Brownian bridge (cf. [19]).

The Laplace operator's eigenvalues and eigenfunctions depend on the geometry of the region as well as on the imposed boundary conditions, but are invariant under rotations and shifts [20]. Moreover, when the domain is scaled by a factor $c > 0$ the eigenvalues are rescaled by $1/c^2$, with associated eigenfunctions $\psi_k(x/c)$ for $x \in cD$. In the case of Dirichlet- or Robin boundary conditions, we can relate the covariance function $C_{X_H}^D$ over D, i.e. the unique integral kernel for $(-\Delta)^{-(\frac{d}{2}+H)}$ over D, with the covariance $C_{X_H}^{cD}$ over

cD, by observing that for any $f \in L^2(cD)$, and $x, y \in cD$,

$$
\begin{aligned}
(-\Delta)^{-(\frac{d}{2}+H)} f(x) &= \int_{cD} C_{X_H}^{cD}(x,y) f(y)\, dy = \sum_{k=0}^{\infty} \left(\frac{\lambda_k}{c^2}\right)^{-(\frac{d}{2}+H)} \int_{cD} \psi_k(y/c) f(y)\, dy\, \psi_k(x/c) \\
&= c^{d+2H} \sum_{k=0}^{\infty} \lambda_k^{-(\frac{d}{2}+H)} \int_D \psi_k(\tilde{y}) f(c\tilde{y}) c^{-d}\, d\tilde{y}\, \psi_k(x/c) \\
&= \int_D c^{2H} \sum_{k=0}^{\infty} \lambda_k^{-(\frac{d}{2}+H)} \psi_k(x/c)\psi_k(\tilde{y}) f(c\tilde{y})\, d\tilde{y} = \int_D c^{2H} C_{X_H}^{D}(\tilde{x},\tilde{y}) f(c\tilde{y})\, d\tilde{y},
\end{aligned}
$$

where $\tilde{x} = x/c$ and $\tilde{y} = y/c$. This implies $C_{X_H}(cx, cy) = c^{2H} C_{X_H}(x, y)$ and therefore that the field $X_H(cx) \stackrel{d}{=} c^H X_H(x)$ for $x \in D$. Similar scale invariance holds for Riesz fields with Neumann conditions.

Figures 5 - 7 illustrate the effect of boundary conditions on realizations of the field, its covariance function and its power spectral density. In each case, the leftmost figure corresponds to homogeneous Neumann boundary conditions over the entire boundary, the middle figure corresponds to Neumann conditions on the left and right and Dirichlet conditions at the top and bottom, while the rightmost figure corresponds to homogeneous Dirichlet conditions. Evidently, the field's variance is considerably lower where Dirichlet boundary conditions are enforced, which manifests in both its sample path and covariance (see Figures 5 and 6).

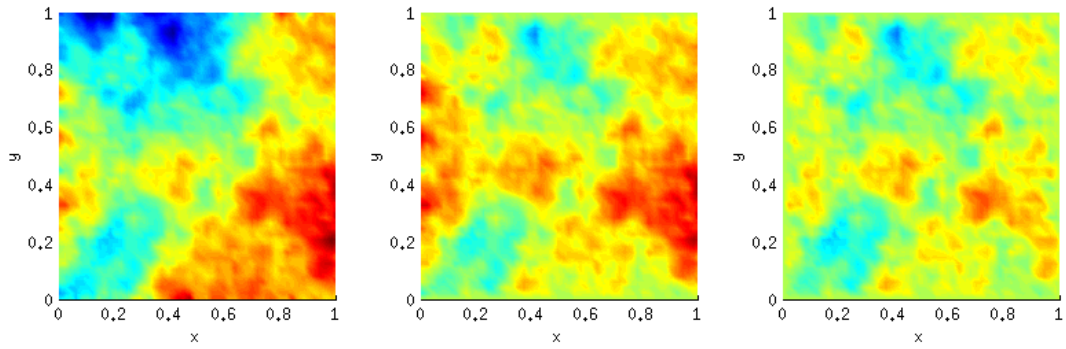

Figure 5: Realizations of the Riesz field $X_{0.25}$ over a square domain, using the same random seed, but with different boundary conditions.

The sample paths, however, seem to have the same degree of 'roughness'. This is confirmed by log plots of the appropriate periodograms (see Figures 7 and 8), displaying comparable power spectral decay rates in the high frequency regions.

5. Numerical Simulations of Riesz Fields. This section discusses three classes of methods for generating numerical simulations of power-law noises on non-standard grids over

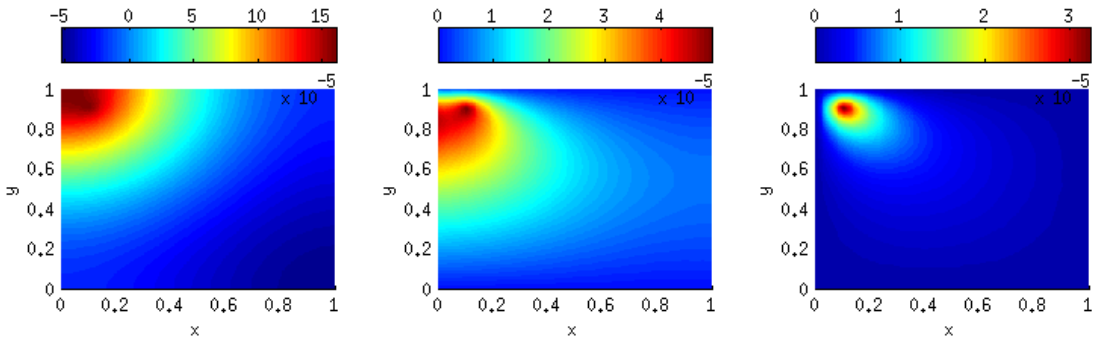

Figure 6: The covariance of the field $X_{0.25}$ at the point $(0.1, 0.9)$ with other points in the domain, for different boundary conditions.

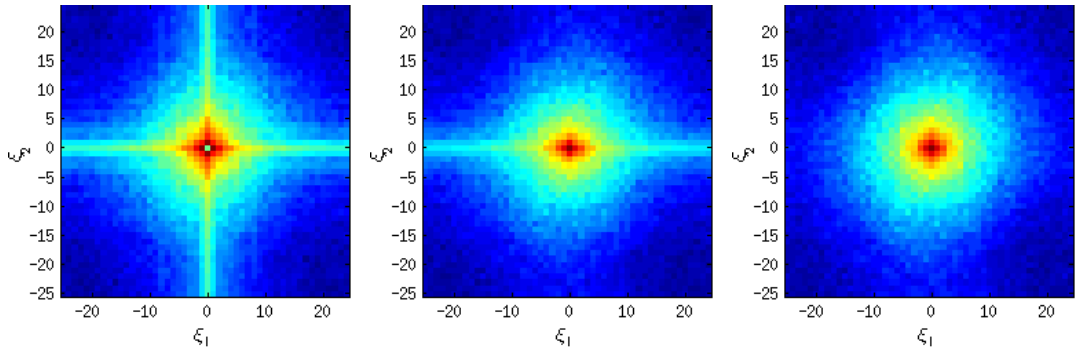

Figure 7: Estimated log power spectral densities of the Riesz field $X_{0.25}$ for different boundary conditions.

arbitrary spatial domains. Subsection 5.1 treats approximations of Riesz field sample paths based on the series representations (4.3) and (4.4), which may require computing the full eigen-decomposition of the discretized Laplace matrix. Subsection 5.2 constructs these sample paths as the solution of a fractional-in-space diffusion equation with white noise forcing term. The contour integral method [21] is a parallelizable algorithm that allows accurate approximations of these solutions to be computed more cheaply than using the eigen-decomposition of the discrete Laplacian. Finally, Subsection 5.3 treats the generation of sample paths as the integral of the Riesz potential with respect to the white noise measure.

5.1. Spectral Representation. Equations (4.3) and (4.4) express the sample paths of Riesz fields over a bounded domain $D \subset \mathbb{R}^d$ in terms of the eigenvalues and eigenfunctions of the negative Laplacian. These are uniquely determined by the geometry of the underlying domain, together with the imposed boundary conditions. Explicit formulae for them exist over a variety of simple domains, including intervals, hyper-rectangles, parallelepipeds,

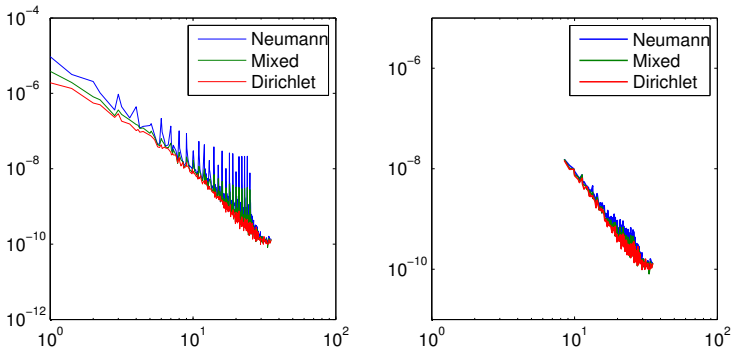

Figure 8: Log-log plot of the radial frequency against the azimuthal average of the log power spectral density for different boundary conditions, for all observed frequencies (left) and when frequencies lower than 5Hz are discarded (right).

disks, sectors, spheres and spherical shells, ellipses and elliptical annuli, as well as triangles (see [20]). In these cases, approximate sample paths $\hat{X}_H$ can be generated at an arbitrary set of points $\{x_i\}_{i=1}^n$, by truncating the sum in (4.3) or (4.4) up to the K^{th} term, yielding

$$X_{H,K}^{\text{eig}}(x_i) = \sum_{k=0}^{K} \lambda_k^{-(\frac{d}{4}+\frac{H}{2})} \psi_k(x_i) Z_k, \quad \text{for } i = 1, 2, ..., n.$$

This spectral approximation has the benefit of being applicable on an arbitrary set of points, in contrast to simulations generated by the circulant embedding technique, where points are required to form a uniformly spaced rectangular grid. Because the Laplacian eigenfunctions are orthogonal in $L^2(D)$, Weyl's law provides an asymptotic error estimate for the spectral approximation of C_{X_H} in terms of the sum of negative powers of the neglected eigenvalues. Since eigenfunctions corresponding to high eigenvalues are often highly oscillatory, however, the accuracy of this spectral approximation is conditional on using meshes with a fine enough resolution, so as to avoid additional errors due to aliasing.

5.1.1. Finite Element Approximations of the Laplace Operator. In general, the fractional Laplacian must be approximated. This can be done by numerically solving the eigenvalue problem for a matrix representation A of the Laplacian $(-\Delta)$, based on the mesh induced by the nodes and on the given boundary conditions. Various methods are available for discretizing the Laplace operator on a mesh, including finite differences or finite volumes, but we focus here on finite element approximations, since they are well-known for their efficacy in the presence of complex geometries. Let $\{x_i\}_{i=1}^n$ be finite element nodes associated with a regular triangulation $\tau_h = \{\triangle_i\}_{i=1}^m$ of the domain D with maximum mesh spacing parameter h, and let $\{\phi_i\}_{i=1}^n$ be the corresponding piecewise polynomial basis functions. The finite element Laplace matrix A can now be formed by letting $A = M^{-1}L$, where M is the mass matrix and L the stiffness matrix with appropriate boundary conditions,

defined respectively by

$$M_{ij} = \int_D \phi_i(x)\phi_j(x)\,dx \quad \text{and} \quad L_{ij} = \int_D \nabla\phi_i(x)\cdot\nabla\phi_j(x)\,dx, \quad \text{for} \quad i,j = 1,2,\dots n.$$

The matrix M is strictly positive definite, while L may only be positive semi-definite if pure Neumann boundary conditions are imposed. Both matrices are therefore unitarily diagonalizable with a non-negative spectrum. A on the other hand, is not symmetric in general. However, since A is similar to $M^{1/2}AM^{-1/2} = M^{-1/2}LM^{-1/2}$, which is symmetric semi-positive definite, it is diagonalizable with non-negative spectrum $\sigma(A)$. Therefore, let $A = V\Lambda V^{-1}$, where $\Lambda = \operatorname{diag}(\lambda_1^h, \dots, \lambda_n^h)$ so that $0 \le \lambda_1^h \le \lambda_2^h \le \dots \le \lambda_n^h$, and the columns of V form the coefficients in the finite element approximation of the associated eigenfunctions. If L is strictly positive definite then the sample path X_H can be approximated by means of the spectral expansion

$$\widehat{X}^{\text{eig}}_{H,h}(x) = \sum_{k=1}^{n} (\lambda_k^h)^{-(\frac{d}{4}+\frac{H}{2})}\psi_k^h(x)Z_k,$$

where $\psi_k^h(x) = \sum_{i=1}^n V_{ik}\phi_i(x)$ and $\mathbf{Z} = [Z_1, ..., Z_n]^T$ is a standard normal random vector in $\mathbb{R}^n$. Since this method requires the full eigen-decomposition of A, the computational cost involved can be considerable, scaling as $O(n^3)$. The spatial accuracy of $\widehat{X}^{\text{eig}}_{H,h}$ is determined by the accuracy with which the discrete eigenvalues λ_k^h and eigenfunctions $\psi_k^h(x)$ approximate the true spectrum. In [7] the finite element approximation of the Laplacian eigenvalue problem is discussed at length. In particular, it is shown that for smooth eigenfunctions (such as those arising when D is convex), $\lambda_k^h \to \lambda_k$ at the rate $O(h^2)$, while $\psi_k^h \to \psi_k$ at the rate $O(h^2)$ in the $L^2(D)$ norm and $O(h)$ in the $H^1(D)$ norm. The corresponding rate coefficients depend on the particular eigenspace being approximated and can be larger if ψ_k is close to singular.

5.2. Fractional Powers of the Discrete Laplacian. The discussion in Section 4 suggests that in light of the finite element discretization A of the Laplacian, approximations $\hat{X}_{H,h}$ of Riesz sample paths can also be obtained by computing the finite element coefficient vector $\hat{\mathbf{X}}_{H,h}$ as the solution of the discretized fractional diffusion equation, i.e. $A^{\frac{d}{4}+\frac{H}{2}}\hat{\mathbf{X}}_{H,h} = \mathbf{Z}$, where $\mathbf{Z} = [Z_1, ..., Z_n]^T$ is a standard normal vector and $A^{-(\frac{d}{4}+\frac{H}{2})} = V\Lambda^{-(\frac{d}{4}+\frac{H}{2})}V^{-1}$. Indeed, the right hand side of the finite element discretization of the Poisson problem with white noise forcing term is given by $M\mathbf{Z}$, since $\int \phi_i\,dW = W(\phi_i)$ for all test functions $\{\phi_i\}_{i=1}^n$, while the stiffness matrix L remains unchanged. Since the solution of the fractional diffusion equation can also readily be expressed in terms of the spectral decomposition of A, the spatial accuracy of $\hat{\mathbf{X}}_{H,h}$ is determined by the accuracy of the discretized spectrum, discussed in Section 5.1.1.

The contour integral method presents a more efficient way of computing fractional powers of A without resorting to its full eigen-decomposition.

5.2.1. The Contour Integral Method. The contour integral method (CIM) is based on the representation of matrices of the form $f(A)$ as contour integrals around the spectrum $\sigma(A)$ of A and allows for the computation of fractional powers of the discrete Laplace operator A without first computing its entire spectrum. Specifically, let f be an analytic function in a region containing $\sigma(A)$ and let Γ be a contour lying within this region and winding once around $\sigma(A)$ in a counter-clockwise direction. Then (see [22], Definition 1.11 and Theorem 1.12)

$$f(A) = \frac{1}{2\pi i} \int_{\Gamma} f(z)(zI - A)^{-1}\, dz. \tag{5.1}$$

In practice, this matrix-valued integral must be approximated by numerical quadrature, giving rise to the weighted sum

$$f_N(A) = \sum_{i=1}^{N} w_i f(\xi_i)(\xi_i I - A)^{-1},$$

where w_i and ξ_i are a set of weights and nodes. In general, the evaluation of this sum requires computing N resolvent matrices, although these computations can be done completely in parallel. For the purposes of simulating Riesz sample paths, we require only matrix-vector products of the form $A^{-(\frac{d}{4}+\frac{H}{2})}\mathbf{Z}$, where $\mathbf{Z}$ is the standard normal vector, provided A is non-singular. Since $A = M^{-1}L$, the sum above takes the form

$$\hat{\mathbf{X}}^{\text{cim}}_{H,h} := A^{-(\frac{d}{4}+\frac{H}{2})}\mathbf{Z} \approx \sum_{i=1}^{N} w_i \xi_i^{-(\frac{d}{4}+\frac{H}{2})}(\xi_i M - L)^{-1} M\mathbf{Z},$$

requiring N system solves, so that neither A nor its inverse need be formed explicitly. Moreover, these systems are sparse and can therefore be solved efficiently (both in terms of storage and the number of floating point operations) through the use of sparse linear solvers, whose computational cost typically scales as $O(n)$ so that the total cost scales as $O(nN)$. To compute the associated covariance matrix, however, the full matrix inverse is required. In the case of pure Neumann boundary conditions, it was shown (c.f. Theorem 4.1 [9]) that $A_{\text{neu}}^{-(\frac{d}{4}+\frac{H}{2})} := V_{\text{neu}}\Lambda^{-(\frac{d}{4}+\frac{H}{2})}V_{\text{neu}}^{-1}$, where $\Lambda_{\text{neu}} = \text{diag}(\lambda_2^h, ..., \lambda_n^h)$ and V_{neu}, is formed from all but the first columns of V, can be computed by the contour integral

$$A_{\text{neu}}^{-(\frac{d}{4}+\frac{H}{2})} = \frac{1}{2\pi i} \int_{\Gamma_2} z^{-(\frac{d}{4}+\frac{H}{2})}(zI - A)^{-1}\, dz,$$

where Γ_2 contains $\lambda_2, ..., \lambda_n$, but not λ_1.

Unfortunately, the accuracy of conventional approaches, such as applying the trapezoidal rule to the circular contour enclosing $\sigma(A)$, deteriorates as the condition number $\kappa(A)$ of A grows, with a convergence rate that depends linearly on $\kappa(A)$. For the finite element Laplacian, $\kappa(A) = \lambda_n^h/\lambda_1^h$ in turn grows as the mesh is refined. In [21], the authors develop a numerical quadrature scheme whose accuracy deteriorates only logarithmically in terms of $\kappa(A)$ for functions f that are analytic in the slit complex plane

$\mathbb{C}\backslash(-\infty, 0]$, by first mapping the region $\mathbb{C}\backslash((-\infty, 0] \cup \sigma(A))$ conformally onto an annulus and applying the trapezoidal rule there. It can then be shown (see [21] Theorem 2.1) that $\|f(A) - f_N(A)\| = O(e^{-\pi^2 N/(\log(\kappa(A)+3)})$. This quadrature scheme was used in [9] to compute the fractional FEM Laplace operator A in aid of approximating the solution of a fractional-in-space reaction diffusion equation.

To investigate the accuracy and efficiency of the CIM method, we compute the finite element coefficients $\hat{\mathbf{X}}_{H,h}^{\text{cim}}$ of a single sample path of the Riesz field satisfying homogeneous Dirichlet conditions for different values of the Hurst parameter, different levels of mesh refinement, and different numbers of quadrature nodes. The quadrature nodes and weights were computed using Algorithm 1 in [9] (see also `method1.m` in [21]), based on Driscoll's Schwarz-Christoffel Toolbox [14]. For the first and second mesh refinement levels, the reference path was computed using Matlab's `mpower` command and the backslash operation '\'. For the finer meshes, we used the CIM method with 100 quadrature points. Figure 9 shows the convergence rate of the CIM method for different meshes. As expected, the method converges exponentially, but the convergence rate deteriorates as the mesh becomes finer, giving rise to a higher condition number. To assess the efficiency of the CIM, we

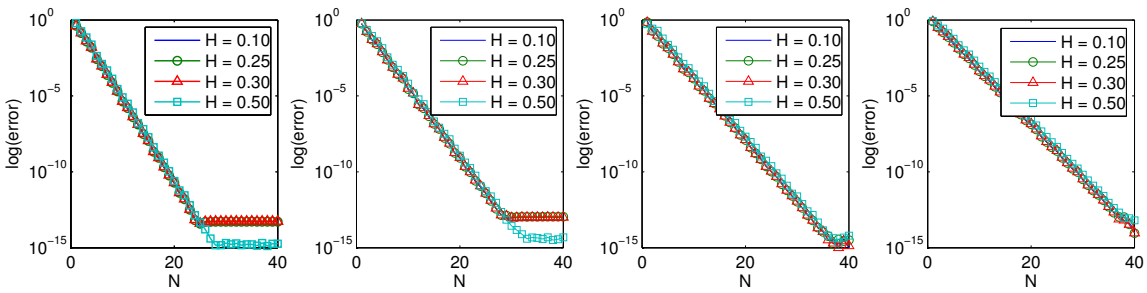

(a) 144 spatial nodes. (b) 529 spatial nodes. (c) 2025 spatial nodes. (d) 7921 spatial nodes.

Figure 9: The relative L^∞-error of the CIM approximation of the sample path $\mathbf{X} = A^{-(\frac{d}{4}+\frac{H}{2})}\mathbf{Z}$ for successive refinements of the spatial mesh.

compare its CPU time, using $N = 40$ quadrature nodes, with that of forming the fractional inverse of A by computing its full eigen-decomposition, using the `eig` function in Matlab. We computed both the timings and relative errors in each case, averaged over different values of H. We ran our computations on a Intel Core i5-2520M CPU @ 2.50GHz x 4, running Matlab R2012 without parallelization. Table 1 clearly shows the advantage in computational cost of the CIM over using the eigen-decomposition of A. For a comparable (or even better) relative error, the CIM is an order of magnitude faster, especially for finer meshes.

5.3. Discretization of the Riesz Kernel. In this section we propose a method for simulating power law noises as convolutions of the Riesz potential (3.3) with white noise fields, as an alternative to using the inverse fractional Laplace operator. This formulation

Level	Nodes	$\kappa(A)$	CIM time (sec)	eig time (sec)	CIM error	eig error
1	144	148	0.1056	0.0068	3.5431e-14	2.2629e-14
2	529	623	0.3449	0.2196	6.9600e-14	6.3948e-14
3	2025	2525	1.4320	11.5975	2.9589e-15	4.6633e-13
4	7921	10133	10.3513	645.7518	2.4304e-14	1.0503e-11

Table 1: Condition numbers of the discrete Laplacian as well as computational times and relative errors for both the CIM ($N = 40$) and the eigenvalue method for different spatial refinement levels.

does not allow for the explicit enforcement of boundary conditions, although it is simpler to implement than previously discussed methods, since it does not require the inversion of fractional Laplace operators. To prevent fields from exhibiting spurious correlations over non-convex domains (see Figure 3), we replace the Euclidean distance appearing in the Riesz kernel by the distance $d_D(x, y)$ of the shortest path in D between the points, i.e.

$$d_D(x, y) := \min\{\text{length}(\gamma) : \gamma \text{ is a path from } x \text{ to } y\}, \tag{5.2}$$

giving rise to the modified kernel $k_H(x, y) := d_D(x, y)^{-(\frac{d}{2}+H)}$. At any point $x \in D$, the value of the random field $X_H^{\text{riesz}}(x)$ is then given by

$$X_H^{\text{riesz}}(x) = c_{H+\frac{d}{2}} \int_D k_H(x, y)\, dW(y),$$

where $c_{H+\frac{d}{2}}$ is given in (3.3) and W is the white noise measure defined in Section 2.

As before, let $\tau_h = \{\triangle_i\}_{i=1}^m$ be a regular triangulation of the region D. For any $x \in D$, we approximate the integral kernel $k_H(x, \cdot)$ by the piecewise constant function

$$\hat{k}_H(x, y) := \sum_{i=1}^m k_H(x, y_i^*) \mathbb{1}_{\triangle_i}(y),$$

where $\mathbb{1}$ is the indicator function and $y_i^* \in \triangle_i$ is a representative point in the interior of the i^{th} element $\triangle_i$. In our computations, we take y_i^* to be the centroid of the i^{th} element. Both the computational and storage cost of this method are dominated by the computation and storage of the shortest distances between any two finite element nodes. In our computations, we make use of Floyd's algorithm [17] (see Figure 10), with a total computational cost of $O(n^3)$ and a storage cost of $O(n^2)$. However, for finite element meshest which are sparse graphs, the Johnson algorithm [27] is more efficient, with a storage cost of $O(n)$ and a computational cost of $O(n^2 \log(n))$.

Consequently, the random field X_H^{riesz} can be approximated by

$$\hat{X}_{H,h}^{\text{riesz}}(x) := \sum_{i=1}^m k_H(x, y_i^*) W(\triangle_i) = \sum_{i=1}^m k_H(x, y_i^*) |\triangle_i| Z_i, \text{ with } Z_i \sim N(0, 1) \text{ i.i.d..} \tag{5.3}$$

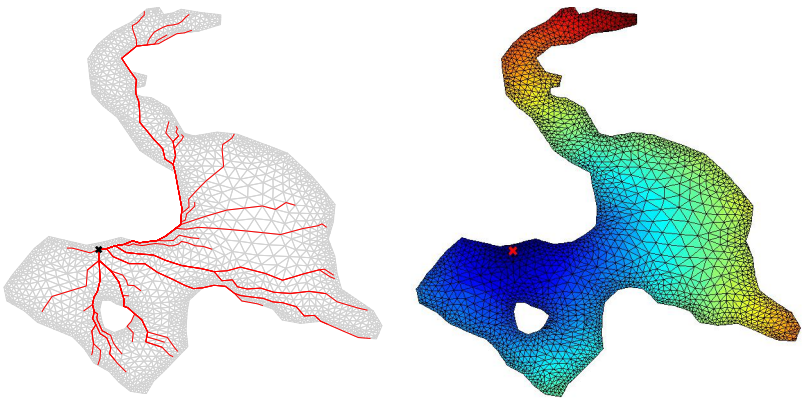

(a) Shortest paths from point marked by x to 40 other random points in the domain.

(b) Distances from point marked by x to all other points in the domain

Figure 10: Illustrations of the modified metric over the lake domain.

Like Hosking's model, this formulation expresses the field at a given point $x \in D$ as a linear combination of white noise sources located in each of the elements, so that both the size of the element and the distance from x to the element's centroid determine the strength of the influence of the noise source. For points lying close together, the field is largely determined by the same random disturbances leading to a higher level of correlation, the strength of which depends on the decay rate of the kernel, i.e. on H. Figure 11c shows the correlation of the field $\hat{X}^{\text{riesz}}_{0.25,h}$, whose sample path is depicted in Figure 11a, at a given point. As expected, the correlation decreases as we move away from the point. Unlike elliptic Gaussian fields with Neumann boundary conditions (see Figure 6), however, the correlation is always positive. An unusual property of this field is that the correlation of the point 'x' with points on the boundary is slightly higher than with points that are closer but that lie in the interior. This is due to the fact that points on the boundary are influenced by fewer noise sources than their interior neighbors, rendering them more correlated. Figure 11b shows that both the distance of the noise source from a point x, as well as the size of the element determine its contribution to the field's variance at x.

Spatially Varying Hurst Parameters. Although $\hat{X}^{\text{riesz}}_{H,h}$ cannot be used when it is necessary to enforce boundary conditions, the explicit appearance of the Hurst parameter in Equation (5.3) allows us to model fields whose Hurst parameter H is spatially varying, by letting

$$\hat{X}^{\text{riesz}}_{H,h}(x) := \sum_{i=1}^{m} d_D(x, y_i^*)^{-(\frac{d}{2}+H(x))} W(\triangle_i).$$

Figure 12b shows the sample path of a field $\hat{X}^{\text{riesz}}_{H,h}(x)$ with Hurst parameter that increases along the x-direction from left to right.

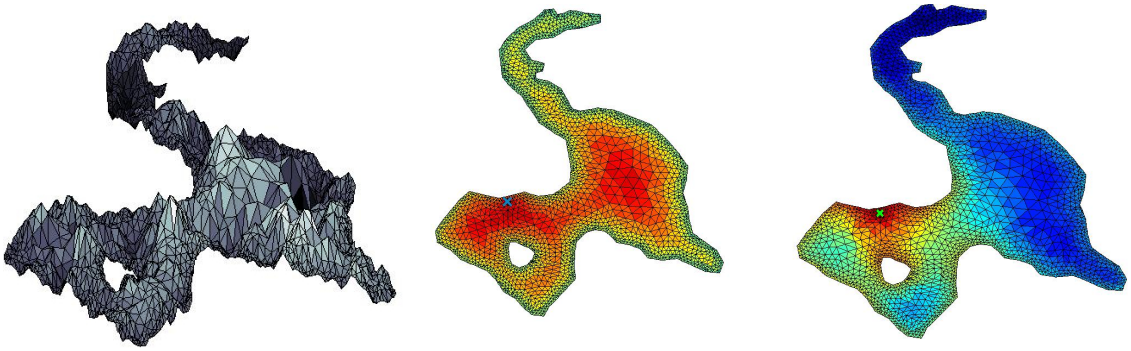

(a) Sample path of the power-law noise $\hat{X}^{\text{riesz}}_{0.25,h}$ over the lake region.

(b) Contribution of each noise term to the variance at the point marked by 'x'.

(c) The correlation at the point marked by 'x'.

Figure 11: Sample paths and point-wise covariance structure of $\hat{X}^{\text{riesz}}_{0.25,h}$.

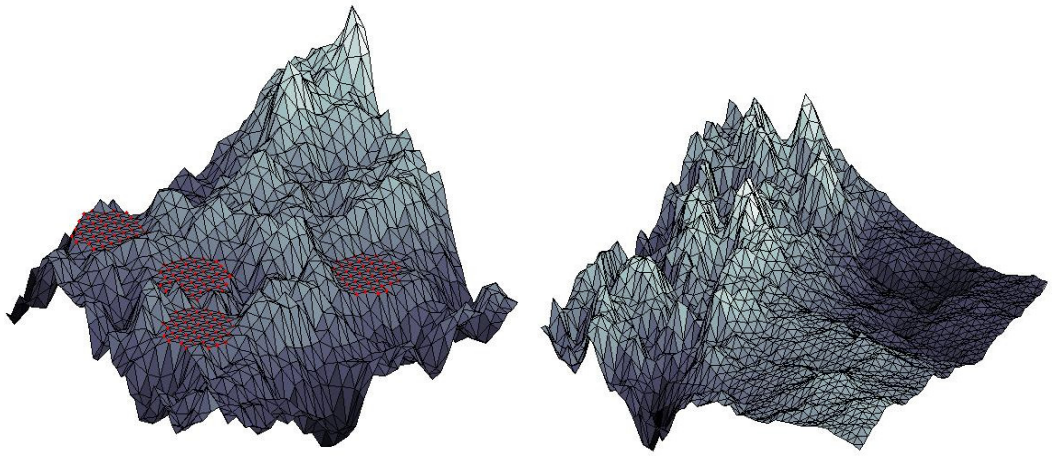

(a) $\hat{X}^{\text{cim}}_{H,h}$ with homogeneous Dirichlet conditions on the red circles.

(b) $\hat{X}^{\text{riesz}}_{H,h}$ with Hurst parameter $H(x,y) = 0.5 + 0.5/\pi * \arctan(10 * (x - 0.5))$.

Figure 12: Sample paths of non-standard scale-invariant fields over a square domain.

6. Conclusion. Statistical self-similarity is fundamental property of random fields, observable under various guises in many physical-, biological-, and man-made systems. Due to the multitude of ways in which self-similarity can be manifested, especially over general, multi-dimensional index sets, statistical models of self-similar or power-law noises should be sufficiently flexible to accommodate both constraints imposed by observations, and the needs of the modeler. In this paper we showed that the elliptic Gaussian fields form a wide class of locally self-similar random fields of known smoothness that can be generated over

non-standard spatial domains and on arbitrary meshes. Defining the random field as the solution of an SPDE has several advantages, such as the ability to seamlessly incorporate observations though the imposition of interior Dirichlet conditions (see Figure 12a). We proposed three algorithms for generating numerical simulations of power-law noise and discussed and compared their properties, strengths and limitations.

The theory of Gaussian models for spatially varying power-law noises, also known as elliptic Gaussian processes [5] [4], or Riesz fields [18, 19], is fairly recent and there is a need for a more complete understanding of the nature of solutions of SDE's and SPDE's and their approximations, when the underlying parameters that are spatially varying power-law noises. This includes a convergence theory for approximations based on finite elements, or finite differences. Another direction of future research involves the quantification of more specific features of the field, such as the presence of a grain, as well as the incorporation of these into the random field model.

There is also an interesting connection between Riesz fields and Matérn random fields, a family of Gaussian random fields whose power spectral densities exhibit a power-law decay in the intermediary frequency range. In particular, Matérn fields Y can be shown to satisfy the equation $(\kappa^2 - \Delta)^{\alpha/2} Y(\mathbf{x}) = W(\mathbf{x})$ over $\mathbb{R}^d$, where $\alpha > 0, \kappa > 0$ are parameters related to the covariance function of Y. In [35], this relation was used to recursively construct a Gaussian Markov Random field representation for Matérn random fields, when $\alpha \in \mathbb{N}$. This representation leads to considerably more efficient approximations than traditional covariance factorization, due to the relative sparsity of the precision matrix. A possible avenue for future research would be to investigate this relation for non-integer values of α and to consider the limiting case as $\kappa \to 0$.

REFERENCES

[1] N. Arnold and W. Rees, *Self-similarity in glacier surface characteristics*, Journal of Glaciology, 49 (2003), pp. 547–554.

[2] A. Barabási, *Fractal concepts in surface growth*, Cambridge university press, 1995.

[3] J. Barnes and D. Allan, *A statistical model of flicker noise*, Proceedings of the IEEE, 54 (1966), pp. 176–178.

[4] A. Benassi, S. Cohen, and J. Istas, *Local self-similarity and the hausdorff dimension*, Comptes Rendus Mathematique, 336 (2003), pp. 267–272.

[5] A. Benassi, S. Jaffard, and D. Roux, *Elliptic gaussian random processes.*, Revista matemática iberoamericana, 13 (1997), pp. 19–90.

[6] F. Biagini, Y. Hu, B. Øksendal, and T. Zhang, *Stochastic calculus for fractional Brownian motion and applications*, Springer, 2008.

[7] D. Boffi, *Finite element approximation of eigenvalue problems*, Acta Numerica, 19 (2010), pp. 1–120.

[8] E. Bouchaud, *Scaling properties of cracks*, Journal of Physics: Condensed Matter, 9 (1997), pp. 4319–4344.

[9] K. Burrage, N. Hale, and D. Kay, *An efficient implicit fem scheme for fractional-in-space reaction-diffusion equations*, SIAM Journal on Scientific Computing, 34 (2012), pp. A2145–A2172.

[10] G. Chan and A. T. Wood, *Simulation of multifractional brownian motion*, in COMPSTAT, Springer, 1998, pp. 233–238.

[11] N. Cressie, *Fitting variogram models by weighted least squares*, Journal of the International Associ-

ation for Mathematical Geology, 17 (1985), pp. 563–586.
[12] C. Dietrich and G. N. Newsam, *Fast and exact simulation of stationary gaussian processes through circulant embedding of the covariance matrix*, SIAM Journal on Scientific Computing, 18 (1997), pp. 1088–1107.
[13] R. Dobrushin, *Gaussian and their subordinated self-similar random generalized fields*, The Annals of Probability, (1979), pp. 1–28.
[14] T. A. Driscoll, *Algorithm 756: A matlab toolbox for schwarz-christoffel mapping*, ACM Trans. Math. Software, (1996), pp. 168–186.
[15] P. Flandrin, *On the spectrum of fractional brownian motions*, Information Theory, IEEE Transactions on, 35 (1989), pp. 197–199.
[16] ———, *Wavelet analysis and synthesis of fractional brownian motion*, Information Theory, IEEE Transactions on, 38 (1992), pp. 910–917.
[17] R. W. Floyd, *Algorithm 97: Shortest path*, Commun. ACM, 5 (1962), pp. 345–.
[18] Z. Gelbaum, *Fractional brownian fields over manifolds*, arXiv preprint arXiv:1207.6419, (2012).
[19] Z. Gelbaum and M. Titus, *Simulation of fractional brownian surfaces via spectral synthesis on manifolds*, arXiv preprint arXiv:1303.6377, (2013).
[20] D. S. Grebenkov and B.-T. Nguyen, *Geometrical structure of laplacian eigenfunctions.*, SIAM Review, 55 (2013), pp. 601–667.
[21] N. Hale, N. J. Higham, and L. N. Trefethen, *Computing a^α,\log(a), and related matrix functions by contour integrals*, SIAM Journal on Numerical Analysis, 46 (2008), pp. 2505–2523.
[22] N. J. Higham, *Functions of matrices: theory and computation*, Siam, 2008.
[23] T. Hirata, T. Satoh, and K. Ito, *Fractal structure of spatial distribution of microfracturing in rock*, Geophysical Journal International, 90 (1987), pp. 369–374.
[24] J. R. Hosking, *Fractional differencing*, Biometrika, 68 (1981), pp. 165–176.
[25] J. Istas, *Spherical and hyperbolic fractional brownian motion*, Elec. Comm. Prob, 10 (2005), p. 266.
[26] S. Janson, *Gaussian hilbert spaces*, vol. 129, Cambridge university press, 1997.
[27] D. B. Johnson, *Efficient algorithms for shortest paths in sparse networks*, J. ACM, 24 (1977), pp. 1–13.
[28] J. B. Johnson, *The schottky effect in low frequency circuits*, Physical review, 26 (1925), p. 71.
[29] N. J. Kasdin, *Discrete simulation of colored noise and stochastic processes and 1/f α power law noise generation*, Proceedings of the IEEE, 83 (1995), pp. 802–827.
[30] B. Kaulakys, *Ruseckas j. gontis v. alaburda m.(2006) nonlinear stochastic models of 1/f noise and power-law distributions*, Physica A, 365 (2006), pp. 217–221.
[31] M. S. Keshner, *1/f noise*, Proceedings of the IEEE, 70 (1982), pp. 212–218.
[32] A. Khintchine, *Korrelationstheorie der stationären stochastischen prozesse*, Mathematische Annalen, 109 (1934), pp. 604–615.
[33] D. Krige, *A statistical approach to some mine valuations and allied problems at the witwatersrand*, master's thesis, University of Witwatersrand, 1951.
[34] P. Lévy, *Random functions: general theory with special reference to Laplacian random functions*, vol. 1, University of California Press, 1953.
[35] F. Lindgren, H. Rue, and J. Lindström, *An explicit link between gaussian fields and gaussian markov random fields: the stochastic partial differential equation approach*, Journal of the Royal Statistical Society: Series B (Statistical Methodology), 73 (2011), pp. 423–498.
[36] S. B. Lowen and M. C. Teich, *Fractal-based point processes*, vol. 366, John Wiley & Sons, 2005.
[37] B. B. Mandelbrot and J. W. Van Ness, *Fractional brownian motions, fractional noises and applications*, SIAM review, 10 (1968), pp. 422–437.
[38] B. T. Milne, *Spatial aggregation and neutral models in fractal landscapes*, American Naturalist, (1992), pp. 32–57.
[39] E. Milotti, *1/f noise: a pedagogical review*, arXiv preprint physics/0204033, (2002).
[40] A. Nieto-Samaniego, S. Alaniz-Alvarez, G. Tolson, K. Oleschko, G. Korvin, S. Xu, and J. Pérez-Venzor, *Spatial distribution, scaling and self-similar behavior of fracture arrays in the los planes fault, baja california sur, mexico*, pure and applied geophysics, 162 (2005), pp. 805–826.
[41] D. Nualart, *The Malliavin calculus and related topics*, Springer, 2006.

[42] W. Schottky, *Small-shot effect and flicker effect*, Physical Review, 28 (1926), p. 74.
[43] D. Spasojević, S. Bukvić, S. Milošević, and H. E. Stanley, *Barkhausen noise: Elementary signals, power laws, and scaling relations*, Phys. Rev. E, 54 (1996), pp. 2531–2546.
[44] M. L. Stein, *Fast and exact simulation of fractional brownian surfaces*, Journal of Computational and Graphical Statistics, 11 (2002), pp. 587–599.
[45] M. Stoyanov, M. Gunzburger, and J. Burkardt, *Pink noise, $1/f^{\alpha}$ noise, and their effect on solutions of differential equations*, International Journal for Uncertainty Quantification, 1 (2011).
[46] D. L. Turcotte, *Fractals and chaos in geology and geophysics*, Cambridge university press, 1997.
[47] N. Wiener, *Generalized harmonic analysis*, Acta mathematica, 55 (1930), pp. 117–258.
[48] W. Willinger, V. Paxson, and M. S. Taqqu, *Self-similarity and heavy tails: Structural modeling of network traffic*, A practical guide to heavy tails: statistical techniques and applications, 23 (1998), pp. 27–53.
[49] G. W. Wornell, *Wavelet-based representations for the 1/f family of fractal processes*, Proceedings of the IEEE, 81 (1993), pp. 1428–1450.
[50] A. M. Yaglom, *Correlation theory of stationary and related random functions*, Springer, 1987.

Referee 1 (Remarks to the Author):

The manuscript presents models for power-law noises and also numerical methods for simulating them on general domains and meshes. The authors draw the connection with fractional Laplacians and thus introduce two simulation methods based on a discrete approximation of the Laplacian. Both methods can be summarized as computing the matrix-vector product $A * Z$, where A is a fractional power of the Laplacian and Z is a random vector drawn from the standard Gaussian. Yet a third method is to resort to a convolution with the Riesz kernel. Hence, the simulation amounts to computing the convolution in a discrete/approximate form. This method does not incorporate boundary conditions, but it has an interesting extension for varying the Hurst parameter spatially. The paper is well written and it merits publication in a UQ journal, wherein simulation of complex physical and stochastic systems is a focused interest.

I would like to comment on the computational costs of the three simulation methods. The authors may consider arranging a discussion of them for enhancing the paper. The first method (spectral representation) is to compute the spectral decomposition of the discrete Laplacian and to raise a certain power on the eigenvalues. The spectral decomposition has an $O(n^3)$ cost, where n is the number of nodes in the mesh. In physical simulations, such an n may be very large, far larger than 7921 in the demonstrated examples. Even though the Laplacian is "sparse" (because both the mass and the stiff matrices are) and sparse eigen-solvers may apply, they do not help too much when all eigenvalues and vectors are needed.

Added the following to Section 5.1.1

Since this method requires the full eigen-decomposition of A, the computational cost involved can be considerable, generally scaling at $O(n^3)$, although it may be partially mitigated through the use of sparse eigen-solvers.

The second method (contour integral) appears to be the most efficient, both in memory and storage. It requires solving a number of linear systems, but the systems are sparse and thus sparse linear solvers can be of good help. When n is 7921, one can afford using the Matlab backslash (which is a direct solver). But for larger n, one probably needs to consider using an iterative solver. Generally, we consider that the cost of solving one system is O(n), although the hidden prefactor can sometimes be large.

Added the following sentence to the middle of Section 5.2.1

Moreover, these systems are sparse and can therefore be solved efficiently (both in terms of storage and the number of floating point operations) through the use of sparse linear solvers, whose computational cost typically scales as $O(n)$ so that the total cost scales as $O(nN)$.

The third method (Riesz kernel) is simpler to implement, but its cost is nontrivial. The reason is that the method requires computing all-pairs shortest paths on a mesh. The Floyd's algorithm used in the paper requires storing the whole $n * n$ distance matrix and the algorithm takes $2 * n^3$ operations. The Johnson's algorithm, on the other hand, may be more efficient, because the mesh is a sparse graph. It requires only $O(n)$ storage for a finite-element mesh, but still, the time cost is $O(n^2 * log(n))$.

Added the following sentences to the paragraph preceding equation 5.3

Both the computational- and storage cost of this method are dominated by the computation and storage of the shortest distances between any two finite element nodes. In our computations, we make use of Floyd's algorithm [17] (see Figure 10), with a total computational cost of $O(n^3)$ and a storage cost of $O(n^2)$. It has since come to our attention that for finite element meshes, which are sparse graphs, the Johnson algorithm [27] is more efficient, with a storage cost of $O(n)$ and a computational cost of $O(n^2 \log(n))$.

A few typos:

1. Page 13, formula below (3.4). The integral should integrate over ξ, not x. 2. Page 20, caption of Figure 10. The exponent of A should be $-(d/4 + H/2)$, rather than $-d/4 + H/2$. Such a missing of parentheses in the exponent also appears in the definition of the modified kernel $k_H(x, y)$ below (5.2) and in the formula at the top of page 23.

Corrected typos.

Referee 2 (Remarks to the Author):

Referee Report for "Power law noises over general spatial domains and on non-standard meshes" by Hans-Werner van Wyk and Max Gunzburger

This paper reviews some theory for random field models, specifically models for which the spectra can be described by a power law. The authors propose three methods for drawing approximate simulations from power law models on non-Euclidean spaces, and they make heavy use of triangularizations to aid in the approximations.

I believe that this paper can be useful to practitioners in many fields who desire computationally efficient methods for simulating random fields. However, it would be helpful to know some more details about the computations and the quality of the approximations.

Overlap, see corrections for reviewer 1.

I realize that the present paper is quite general, but it would also be helpful to have some discussion of how the present paper fits into the well-developed area of Markov random fields, where finite element methods are common, as well as some references to this work, such as the book "Markov Random Fields, Theory and Applications" by Havard Rue and Leonhard Held (2005), and the article "An explicit link between Gaussian fields and Gaussian Markov random fields: The SPDE approach" by Finn Lindgren, Havard Rue, and Johann Lindstrom (2011, JRSS-B).

1. In light of the discussion of periodogram estimates of S at the top of page 3, it seems worth mentioning that the Matern process has precisely the power spectrum that you describe-flat in the low frequency range and a power law at high frequencies.

Added the following at the very end.

There is also an interesting connection between Riesz fields and Matérn random fields, a family of Gaussian random fields whose power spectral densities exhibit a power-law decay in the intermediary frequency range. In particular, Matérn fields Y can be shown to satisfy the equation $(\kappa^2 - \Delta)^{\alpha/2} Y(\mathbf{x}) = W(\mathbf{x})$ over $\mathbb{R}^d$, where $\alpha > 0, \kappa > 0$ are parameters related to Y's covariance function. In [35], this relation was used to recursively construct a Gaussian Markov Random field representation for Matérn random fields, when $\alpha \in \mathbb{N}$. This representation leads to considerably more efficient approximations than traditional covariance factorization, due to the relative sparsity of the precision matrix. A possible avenue for future research would be to investigate this relation for non-integer values of α and to consider the limiting case as $\kappa \to 0$.

2. In the last paragraph on page 4, you mention that all algorithms extend to arbitrary dimensions. This seems believable, but do you have a sense of how the quality of the approximations behaves in higher dimensions? For example, it is well known that discrete Fourier approximations to stationary random fields on regular lattices degrade in quality as dimension increases, due to increased importance of edge effects.

Changed sentence to:

For the sake of visualization, all of our computational results are based on regions in $\mathbb{R}^2$. The algorithms discussed, however, extend readily to three spatial dimensions.

3. On the bottom of page 4, you write "Considering the vastness of this research field, it is inevitable that this paper omits many important approaches, ..." While I sympathize that it is difficult to be aware of all relevant research, especially in other disciplines, I find

this statement to be unsatisfying, especially when there are research fields, such as spatial statistics, that are almost entirely devoted to the problem of studying and making use of random field models.

Replaced this sentence with:

Since this paper focuses on the generation of sample paths over general regions, we do not discuss here the construction and analysis of scale-invariant noises by means of wavelets (see [49] and [16]).

4. In the middle of page 8 you write, "Although fractional Brownian motion, unlike standard Brownian motion, is neither mean square differentiable nor a martingale for $H \neq 1/2, ...$" This sentence seems to imply that standard Brownian motion is mean square differentiable, which I believe not to be true. Please correct this sentence if I am in fact correct.

Changed sentence to:

Although fractional Brownian motion is not mean square differentiable, it does admit a Hölder continuous modification [6]. Moreover, fBm is not a martingale for $H \neq \frac{1}{2}$, and since its variance $\mathrm{Var}(B_H(t)) = \mathbb{E}\left[B_H(t)^2\right] = t^{2H}$ *for $t \geq 0$, it is also non-stationary.*

5. If it is necessary to conserve space, Figure 3 could be omitted.

Omitted Figure 3

6. Definition 3.3 on page 9 implies that $Var(B_H(0)) = 0$, whereas the following paragraph discusses an arbitrary point of origin x_0. For consistency, I suggest either changing Equation (3.2) to reflect arbitrary origin x_0 or explicity stating $x_0 = 0$ in the following paragraph.

7. On page 9, you state that the power law noise model may no longer be relevant in the spatial domain due to the fact that, in many applications, the value of the process may be known at several locations. I disagree with the premise that this case makes such models irrelevant in the spatial domain. It is quite common to assume that the data arose from a realization of a random field that happened to be observed at several locations. The distribution of the process at locations where it was not observed is then quantified as the conditional distribution of the process at the unobserved locations given the values at the observed locations.

Kept Definition 3.3 intact and deleted the argument that fractional Brownian surfaces are not relevant in the context of conditioning on data, but added the following paragraph:

The fBs defined above can readily be modified through translation to a fBs centered at an arbitrary point $x_0 \in \mathbb{R}^d$. It can also be conditioned on available measured data if the spatial domain contains points, or even whole regions, in which the field is known exactly (see e.g. [11, 33]). In this case, the conditional covariance function no longer takes the form (3.2), but the covariance matrix of the field at a finite number of points can be calculated by means of Schur complements.

8. Can you add axis labels in Figures 6, 7, and 8? I'm not sure I understand Figure 7.

Added axes and labels to Figures (now 5,6,7) and rewrote the caption for figure 7 (now 6).

9. In Section 5.1.1, there appears to be a typo in the second displayed equation. Should

$\lambda be \lambda_k^h$?

Corrected

10. In Figure 10, can you add axis labels? It took me quite some time to decide that I think the horizontal axis is N, but I'm still not sure. Do the four plots refer to the four levels in Table 1?

Added axis labels and captions (number of spatial nodes) for each subplot. Added a column to Table 1, showing the relative error of the CIM with $N = 40$. Elaborated on Table 1:

To assess the efficiency of the CIM, we compare its CPU time, using $N = 40$ quadrature nodes, with that of forming the fractional inverse of A by computing its full eigen-decomposition, using the `eig` *function in Matlab. We computed both the timings and relative errors in each case, averaged over different values of H. We ran our computations on a Intel Core i5-2520M CPU @ 2.50GHz x 4, running Matlab R2012 without parallelization. Table 1 clearly shows the advantage in computational cost of the CIM over using the eigen-decomposition of A. For a comparable (or even better) relative error, the CIM is an order of magnitude faster, especially for finer meshes.*

11. You write that the CIM method is compared to the sample path $X = A^{-d/4+H/2}$ (should the exponent be $-d/4 - H/2$?). Isn't this itself a finite element approximation to the true process? Do you have a sense of how well any of these methods approximate the true process?

Corrected the exponent. Added the following paragraph, discussing the spatial accuracy of the finite element approximation, at the end of Section 5.1.1:

The spatial accuracy of $\widehat{X}_{H,h}^{\text{eig}}$ is determined by the accuracy with which the discrete eigenvalues λ_k^h and eigenfunctions $\psi_k^h(x)$ approximate the true spectrum. In [7] the finite element approximation of the Laplacian eigenvalue problem is discussed at length. In particular, it is shown that for smooth eigenfunctions (such as those arising when D is convex) and for piecewise linear finite element approximations, $\lambda_k^h \to \lambda_k$ at the rate $O(h^2)$, while $\psi_k^h \to \psi_k$ at the rate $O(h^2)$ in the $L^2(D)$ norm and $O(h)$ in the $H^1(D)$ norm. The corresponding rate coefficients depend on the particular eigenspace being approximated and can be larger if ψ_k is close to singular.

Also added the following sentence to the first paragraph of Section 5.2:

Since the solution of the fractional diffusion equation can also readily be expressed in terms of the spectral decomposition of A, the spatial accuracy of $\hat{\mathbf{X}}_{H,h}$ is determined by the accuracy of the discretized spectrum, discussed in Section 5.1.1.

12. In Table 1, are the times in seconds? And what is meant by "eig error"?

See corrections for comment 10.

13. Do you have a sense of how the computational effort, in both memory and floating point operations, scales with the number of nodes? I suspect that you stopped at $n = 7921$ nodes due to an $O(n^2)$ memory requirement.

Overlap with comments of reviewer 1. Corrected

14. Can you elaborate on the simulation in Figure 13(a)? This is particularly interesting because it suggests a way to perform conditional simulations of the process given that it

is known at several locations, which is a common application with spatial data.

Added the following sentence to the first paragraph of the conclusion:

Defining the random field as the solution of an SPDE has several advantages, such as the ability to seamlessly incorporate observations though the imposition of interior Dirichlet conditions (see Figure 12a).

15. The first sentence of the conclusion seems out of place. Its inclusion makes it seem like you have just undergone a thorough evaluation of the hypothesis that self-similarity is a natural and universal property.

Removed this sentence, replacing it with:

Statistical self-similarity is fundamental property of random fields, observable in various guises in many physical- biological and man-made systems.