

Fractional Gaussian Fields for Modeling and Rendering of Spatially-Correlated Media

JIE GUO*, State Key Lab for Novel Software Technology, Nanjing University

YANJUN CHEN†, State Key Lab for Novel Software Technology, Nanjing University

BINGYANG HU†, State Key Lab for Novel Software Technology, Nanjing University

LING-QI YAN, University of California, Santa Barbara

YANWEN GUO*, State Key Lab for Novel Software Technology, Nanjing University

YUNTAO LIU, State Key Lab for Novel Software Technology, Nanjing University




Fig. 1. We render a complex scene containing several spatially-correlated media, demonstrating that our method is able to reproduce a wide range of appearances stemming from short-range to long-range correlations and support macroscopic heterogeneity (left). A reference generated by the classical transport theory is provided for a comparison (right).

Transmission of radiation through spatially-correlated media has demonstrated deviations from the classical exponential law of the corresponding uncorrelated media. In this paper, we propose a general, physically-based method for modeling such correlated media with non-exponential decay of transmittance. We describe spatial correlations by introducing the Fractional Gaussian Field (FGF), a powerful mathematical tool that has proven

* Corresponding authors

† Both authors contributed equally to the paper

Authors' addresses: Jie Guo, State Key Lab for Novel Software Technology, Nanjing University, guojie@nju.edu.cn; Yanjun Chen, State Key Lab for Novel Software Technology, Nanjing University, cujooyer@gmail.com; Bingyang Hu, State Key Lab for Novel Software Technology, Nanjing University, fhymyang@gmail.com; Ling-Qi Yan, University of California, Santa Barbara, lingqi@cs.ucsb.edu; Yanwen Guo, State Key Lab for Novel Software Technology, Nanjing University, ywguo@nju.edu.cn; Yuntao Liu, State Key Lab for Novel Software Technology, Nanjing University, windspectator@gmail.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

0730-0301/2019/7-ART1 \$15.00

<https://doi.org/10.1145/3306346.3323031>

useful in many areas but remains under-explored in graphics. With the FGF, we study the effects of correlations in a unified manner, by modeling both high-frequency, noise-like fluctuations and k -th order fractional Brownian motion (fBm) with a stochastic continuity property. As a result, we are able to reproduce a wide variety of appearances stemming from different types of spatial correlations. Compared to previous work, our method is the first that addresses both short-range and long-range correlations using physically-based fluctuation models. We show that our method can simulate different extents of randomness in spatially-correlated media, resulting in a smooth transition in a range of appearances from exponential falloff to complete transparency. We further demonstrate how our method can be integrated into an energy-conserving RTE framework with a well-designed importance sampling scheme and validate its ability compared to the classical transport theory and previous work.

CCS Concepts: • **Computing methodologies** → **Reflectance modeling**.

Additional Key Words and Phrases: Non-exponential transmittance, Volume rendering, Correlation, Random field, Importance sampling

ACM Reference Format:

Jie Guo, Yanjun Chen, Bingyang Hu, Ling-Qi Yan, Yanwen Guo, and Yuntao Liu. 2019. Fractional Gaussian Fields for Modeling and Rendering of Spatially-Correlated Media. *ACM Trans. Graph.* 38, 4, Article 1 (July 2019), 13 pages. <https://doi.org/10.1145/3306346.3323031>

1 INTRODUCTION

When propagating through a participating medium, light is scattered and absorbed in a very complicated way, heavily depending on the material properties. The classical *radiative transfer equation* (RTE) [Chandrasekhar 1960] is valid only when the independent scattering approximation (ISA) holds: the particles forming the medium are far apart from each other, and each particle scatters light independently. The assumption of statistical independent particle positions gives rise to the classical transport equation with exponential falloff of light.

However, most participating media found in nature are not entirely uncorrelated, and the resulting radiation transport will differ from classical transport theory. As non-classical radiation propagation has attracted much attention in many science and engineering fields, different types of models have been proposed to account for spatial correlations in a medium [Davis and Mineev-Weinstein 2011; Davis and Xu 2014; Davis et al. 2018]. In computer graphics, non-classical transport is initially studied and simulated in the context of discrete granular materials [Meng et al. 2015; Moon et al. 2007; Müller et al. 2016]. Recent work attempts to explore a general framework to handle spatially correlations in arbitrary random media [Bitterli et al. 2018; Jarabo et al. 2018]. Jarabo et al. [2018] introduce the Generalized Boltzmann Equation (GBE) proposed by Larsen and Vasques [2011] to the graphics realm, allowing a wide range of both positive and negative correlations. Although heterogeneity is theoretically supported by the GBE, implementing it in a renderer is still challenging. Bitterli et al. [2018] present a more convenient framework based on ensemble-averaged transport in stochastic media and adapt the classical RTE to satisfy energy conservation with non-exponential transmittance. However, only high-frequency (“noisy”) stochastic media are modeled in a physically plausible manner, which limits the range of spatial variability at the micro-scale.

In this paper, we demonstrate that spatial correlations can be represented by a powerful mathematical tool named *Fractional Gaussian Field* (FGF) [Lodhia et al. 2016] which not only includes the high-frequency noise (e.g., pink noise) as a special class, but also covers another well-known class: fractional Brownian motion (fBm) [Mandelbrot and Van Ness 1968]. A notable strength of the FGF is its capability of modeling a broad range of spatial variability from noise-like to completely smooth. This characteristic enables us to study and simulate the effects of both short-range and long-range correlations [Samorodnitsky 2007].

Based on the FGF, our goal is to provide a unified framework with a hand of physically meaningful and easily adjustable parameters for modeling random media with different types of spatial correlations. We use a *homogenization* approach [Davis and Mineev-Weinstein 2011] to derive analytical expressions of transmittance functions which exhibit heavy-tailed free path distributions of the power-law type instead of the exponential type. In our framework, long-range spatial correlations are properly modeled by random media with the micro-scale density of a fBm type and even a k th-order fBm type [Perrin et al. 2001]. We show that the synthesized media tend to be transparent when the order is high.

We implement our non-exponential transmittance functions in an energy-conserving RTE and show that a broad range of appearances

can be simulated with low variance benefiting from a well-designed importance sampling scheme. We demonstrate that the parameter tuning is convenient and intuitive since each parameter has a clear physical meaning. We also show that our method is easily extended to support macro-scale heterogeneity.

Note that our work focuses on deriving a physically-based transmittance model covering a wide range of spatial correlations. The presented model is complementary to existing non-exponential rendering frameworks such as the GBE [Jarabo et al. 2018; Larsen and Vasques 2011] and the framework of Bitterli et al. [2018].

In summary, our contributions are the following:

- We introduce a mathematical tool for physically-based modeling of a broad range of spatial correlations in random media.
- We generate long-range correlations using k -th order fBm and explain the effect of transparency that can be achieved by a high order of k -fBm.
- We demonstrate the practicality of the proposed non-exponential transmittance functions in an energy-conserving RTE framework.
- We design an effective and unified importance sampling scheme for all different types of transmittance functions modeled by FGFs.

2 RELATED WORK

2.1 Classical transport

An overwhelming body of literature exists on simulating the radiative transfer in uncorrelated media [Cerezo et al. 2005; Fong et al. 2017]. Conventionally, two main techniques have been proposed for solving the classical RTE: Monte Carlo path integration methods [Novák et al. 2018; Veach 1997] and photon density estimation methods [Hachisuka et al. 2013; Jarosz et al. 2011]. The former technique focuses on stochastic sampling of light paths connecting light sources and cameras while the latter technique is primarily based on photon mapping [Jensen and Christensen 1998]. It is also possible to unify these two techniques to combine their strengths in a robust estimator [Křivánek et al. 2014]. To significantly speed-up the light transport computation in optically dense media, the diffusion approximation is frequently adopted with strong constraints on the properties of the media [d’Eon and Irving 2011; Donner and Jensen 2005; Frisvad et al. 2014; Jensen et al. 2001]. Considering the versatility and robustness, our method currently builds upon Monte Carlo path integration.

2.2 Non-classical transport

Evidence for the correlation of scattering particles in many fields, including optics [Kostinski 2001; Rojas-Ochoa et al. 2004; Tsang et al. 2002], nuclear engineering [Larsen and Vasques 2011; Larsen and Clark 2014] and atmospheric science [Davis 2006; Davis and Marshak 2004; Doicu et al. 2014], has inspired researchers to investigate the mechanism of non-classical particle transport. In essence, three categories of models have been designed to deal with spatially-correlated media. The so-called homogenization approaches seek effective material properties of the medium to use in the solution of a classical transport problem [Bal and Jing 2011; Davis and Marshak 2004; Davis and Mineev-Weinstein 2011]. These approaches

are quite attractive since they retain the general form of the RTE by correcting conventional radiative parameters into effective parameters taking into account spatially-correlated effects. An alternative is to linearly combine solutions from uniform media and obtain approximated results for spatially-correlated media [Barker et al. 2008; Cahalan 1994]. In the last category, researchers seek to develop new RTEs amenable to either analytical or numerical solutions [Davis and Xu 2014; Davis et al. 2018; d'Eon 2018a; Larsen and Vasques 2011]. Although approaches in the third category should be more broadly applicable, they are rather difficult to be implemented in existing renderers.

In computer graphics, non-classical transport was first explored by Moon et al. [2007] in the context of rendering granular materials. Later, this work was extended to support dynamic scenes [Meng et al. 2015] and heterogeneous mixtures of grains [Müller et al. 2016]. Recently, simulating non-exponential radiative transfer in arbitrarily stochastic media has become a topic of increasing interest, which largely enriches the range of achievable appearance [Bitterli et al. 2018; d'Eon 2014, 2018a,b; Jarabo et al. 2018; Wrenninge et al. 2017]. Based on the GBE proposed by Larsen and Vasques [2011], Jarabo et al. [2018] proposed a new solution to capture different types of correlations in the media. Similar to ours, Bitterli et al. [2018] also utilized a homogenization method to derive non-exponential transmittance in a physically-based manner. Their method built upon the Davis and Mineev-Weinstein model [2011] in which only high-frequency noise-type media could be simulated. Our goal in this paper also focuses on developing a general solution to solve the non-exponential radiation problem in spatially-correlated media. Compared with previous work, we significantly expand the admissible regime of spatially-correlated media benefiting from the FGF.

2.3 Volumetric representation of media

Spatial correlations are generally caused by unevenly distributed micro-structures inside the media. A straightforward way to model these micro-structures relies on constructing their geometric details explicitly [Meng et al. 2015; Moon et al. 2007; Müller et al. 2016; Schröder et al. 2014; Zhang et al. 2013]. Then, scattering parameters of the bulky behavior of the material are derived. In the context of rendering fibrous materials such as hair and fur, Kajiya and Kay [1989] pioneered the idea of using volumetric representations instead of a large quantity of explicit geometries to encode the micro-appearance of fibrous materials. The fabric volumes can be modeled procedurally [Schröder et al. 2011] or gathered from micro CT scans of real samples [Khungurn et al. 2015; Zhao et al. 2011]. After that, the classical RTE is still adopted by specifying the scattering parameters converted from the volumetric data. To account for angular structures in fibrous materials, the microflake model [Dupuy et al. 2016; Heitz et al. 2015; Jakob et al. 2010] is often used in the conversion which is able to reproduce the anisotropic behavior of light scattering. Although images rendered with these explicitly constructed volumetric data contain rich details, the data collection procedure is quite time-consuming. On the contrary, the model in this paper does not rely on any concrete form of media and only statistical quantities related to the FGF are required.

3 FRACTIONAL GAUSSIAN FIELDS

We now give a brief description of some mathematical terminologies and basic facts associated with fractional Gaussian fields (FGFs) [Lodhia et al. 2016]. FGFs are a set of special Gaussian random fields with a property of self-similarity, characteristic of a fractal behavior, which have been studied in many disciplines including hydrology, physics, oceanography and geography. Mathematically, the d -dimensional FGF indexed by a spectral parameter $s (\geq 0)$ is given by

$$M = (-\Delta)^{-\frac{s}{2}} W \quad (1)$$

where W represents a zero-mean Gaussian white noise on \mathbb{R}^d and $(-\Delta)^{-\frac{s}{2}}$ is the fractional Laplacian¹ [Kwaśnicki 2017]. Some well-known random fields are special cases of FGFs, such as white noise ($s = 0$), Brownian motion ($s = 1, d = 1$) and fractional Brownian motion ($1/2 < s < 3/2, d = 1$). Defining the Hurst parameter $H = s - d/2$, self-similarity (or scale invariant) means the field $M(ax)$ has the same law as $a^H M(x)$ for any $a > 0$ and $x \in \mathbb{R}^d$.

3.1 Fourier analysis

Eq. (1) can be explained in the Fourier space²:

$$\mathcal{F}[M] = (2\pi|\mathbf{v}|)^{-s} \mathcal{F}[W] \quad (2)$$

with $\mathbf{v} \in \mathbb{R}^d$ denoting the frequency variable. This equation allows us to interpret the FGF as a white noise filtered by a power function in the Fourier domain. Intuitively, any FGF can be generated by applying an inverse Fourier transform to the left-hand side of Eq. (2). However, this is not generally well-defined for $s \geq d/2$ [Duplantier et al. 2017; Lodhia et al. 2016]. When $s \in (0, d/2)$, the power spectral density (PSD) of the FGF is given by

$$S(\mathbf{v}) = (2\pi|\mathbf{v}|)^{-2s} S_w \quad (3)$$

where S_w is the PSD of white noise.

3.2 Autocovariance function

It is more convenient to study the characteristics of the FGF using the autocorrelation function $\text{cov}(\mathbf{x}, \mathbf{y}) = \langle M(\mathbf{x})M(\mathbf{y}) \rangle$. Here $\langle \cdot \rangle$ refers to the ensemble average over all possible realizations of the FGF. Obviously, the autocovariance function of white noise is a delta function $S_w \delta(\mathbf{x} - \mathbf{y})$, implying that white noise is uncorrelated in space and its variance is S_w . For $0 < s < d/2$ (i.e., $-d/2 < H < 0$), the autocovariance function is given by

$$\text{cov}(\mathbf{x}, \mathbf{y}) = C(H, d) S_w |\mathbf{x} - \mathbf{y}|^{2H} \quad (4)$$

in which $C(H, d)$ is a scaling term varying with respect to both H and d :

$$C(H, d) = \frac{2^{-2H-d} \Gamma(-H)}{\pi^{d/2} \Gamma(H + d/2)}. \quad (5)$$

Here $\Gamma(H) = \int_0^\infty x^{H-1} \exp(-x) dx$ is the Euler's Gamma function satisfying $\Gamma(H) = (H-1)\Gamma(H-1)$. Visualization of $C(H, d)$ with varied H and d is provided in Fig. 2 left. Conventionally, Eqs. (4) and (5) can be derived via applying inverse Fourier transform to the PSD

¹The fractional Laplacian generalizes the notion of derivatives to fractional powers.
²In this paper, the Fourier transformation of a function $f(\mathbf{x})$ is defined as $\mathcal{F}[f](\mathbf{v}) := \int_{\mathbb{R}^d} f(\mathbf{x}) e^{-2\pi i \mathbf{x} \cdot \mathbf{v}} d\mathbf{x}$.

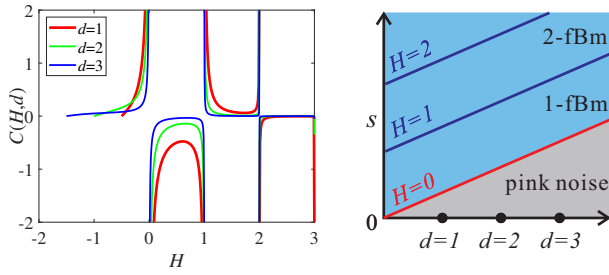


Fig. 2. Left: Visualization of the scaling term $C(H, d)$. Right: The family of FGFs.

of the FGF (i.e., Eq. (3)) according to the Wiener-Khinchin theorem³. $M(\mathbf{x})$ equipped with this type of autocovariance function is widely known as pink noise⁴. A notable property of the above autocovariance function is that it remains unchanged when shifted in space and is invariant under rotation. Therefore, the corresponding FGF is stationary and isotropic.

However, when $s > d/2$ (i.e., $H > 0$) the stationarity does not hold. For instance, the autocovariance function for $H \in (0, 1)$ is given by [Reed et al. 1995]

$$\text{cov}(\mathbf{x}, \mathbf{y}) = C(H, d)S_w(|\mathbf{x} - \mathbf{y}|^{2H} - |\mathbf{x}|^{2H} - |\mathbf{y}|^{2H}) \quad (6)$$

which clearly depends on the spatial positions of both \mathbf{x} and \mathbf{y} . These FGFs are actually multi-dimensional fractional Brownian motion (fBm) on \mathbb{R}^d .

The definition of fBm can be generalized to higher order. Precisely, the k th-order fBm (k -fBm) of H parameter in $(k-1, k)$, where k is a positive integer, has the autocovariance function [Perrin et al. 2001]:

$$\text{cov}(\mathbf{x}, \mathbf{y}) = C(H, d)S_w \left\{ |\mathbf{x} - \mathbf{y}|^{2H} - \sum_{j=0}^{k-1} (-1)^j \binom{2H}{j} \cdot \left[\left(\frac{|\mathbf{x}|}{|\mathbf{y}|} \right)^j |\mathbf{y}|^{2H} + \left(\frac{|\mathbf{y}|}{|\mathbf{x}|} \right)^j |\mathbf{x}|^{2H} \right] \right\} \quad (7)$$

with

$$\binom{2H}{j} = \frac{\Gamma(2H+1)}{\Gamma(j+1)\Gamma(2H-j+1)}. \quad (8)$$

Clearly, it is also non-stationary and includes fBm for the special case in which $k = 1$. Unlike pink noise (i.e., $-d/2 < H < 0$), PSDs of these non-stationary FGFs are difficult to interpret despite using some special tools [Flandrin 1989].

As summarized in Fig. 2 right, the boundary condition $H = 0$ divides the family of FGFs into two main groups. When $H \in (-d/2, 0)$ (gray shaded region) the FGF is defined as pink noise and is stationary. When $H \in (k-1, k)$ (blue shaded region) the FGF is defined as k -fBm and is non-stationary.

³The Wiener-Khinchin theorem states that the PSD of a random field and its autocovariance function are Fourier transform pairs.

⁴The term pink noise has different explanations in literature and in this paper it refers to the random tempered distribution with a power law decay of its PSD.

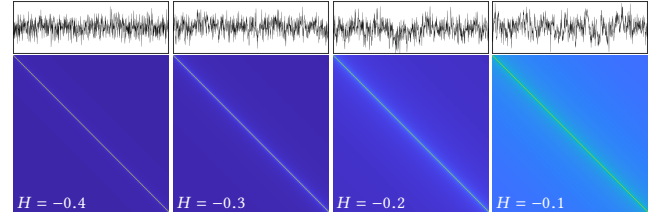


Fig. 3. Simulated realizations of 1D pink noise with different values of H (top) and visualization of their covariance matrices (bottom).

The above discussion assumes that H is not a nonnegative integer. When H is a strictly nonnegative integer (i.e., $s \in \{d/2, d/2+1, d/2+2, \dots\}$), there is a logarithmic correction in the correlation function [Duplantier et al. 2017]. These special cases are not further discussed in this paper. To avoid numerical errors around the integer values of H , we can simply clamp $C(H, d)$ to a predefined threshold in practice.

3.3 Long-range correlation

Physical phenomena modeled by k -fBm usually exhibit long-range correlation (or long memory) behavior that can hardly be obtained by stationary random fields (e.g., pink noise) [Davis and Marshak 2004; Samorodnitsky 2007]. It is well recognized in diverse fields that long-range correlation is the rule rather than the exception [Davis 2006; Davis and Marshak 2004; Davis and Xu 2014]. The Hurst parameter H , describing the raggedness of the resultant field, can serve as a measure of roughness and also a measure of the extent of long-range correlation.

Fig. 3 shows simulated realizations of pink noise with increasing values of H and visualization of the corresponding autocovariance functions. As expected, when $H < 0$, the FGF is noise-like and contains high-frequency fluctuations. The covariance in this case is very close to the delta function, implying that the correlation is relatively short. In comparison, the curves of k -fBm in Fig. 4 tend to be smooth as H increases and their autocovariance functions have remarkable long-range dependencies. In Fig. 5, 2D FGFs are visualized and sampled point sets are generated according to the 2D FGFs served as the density maps. Comparisons show that the point sets generated via white noise and pink noise are quite similar while as H goes up to the scope of fBm the patterns of the point sets have noticeable changes. In particular, aggregates are typically formed when H is large.

In conclusion, if H is large, the field varies very slowly and long-range correlation occurs. As H decreases, the dependencies lessen and the field fluctuates more and more rapidly. Additionally, we see that k -fBm has a never-ending growth of covariance with respect to the spatial location as predicted by Eq. (7) and evidenced in Fig. 4.

4 FLUCTUATION MODELING OF MEDIA

With FGFs at hand, we are able to model random media with different degrees of spatial correlations ranging from short-range to long-range in a physically plausible manner. We assume that the small-scale density of a medium fluctuates according to a FGF and

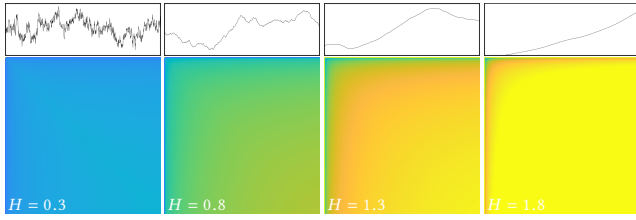


Fig. 4. Simulated realizations of 1D k -fBm with different values of H (top) and visualization of their covariance matrices (bottom).




Fig. 5. Simulated realizations of 2D FGFs. Here we show the density maps (top) of different distributions and the point sets generated according to the density maps (bottom).

only statistical parameters that provide a complete quantitative description of the medium inhomogeneity are required. Using a homogenization strategy focusing on the ensemble averages over all possible realizations, the transmittance functions are derived and used in an energy-conserving RTE.

4.1 Extinction field

We consider a medium that consists of a large quantity of tiny particles interacting with light either by absorption or scattering. The scattering coefficient σ_s and absorption coefficient σ_a respectively represent the fractions of light that are scattered and absorbed. They are routinely combined into the extinction coefficient $\sigma_t = \sigma_s + \sigma_a$ to give the total reduction per unit volume in radiance. The single-scattering albedo Λ is defined as the ratio of σ_s to σ_t , i.e., $\Lambda = \sigma_s/\sigma_t$.

Let us first investigate the behavior of light transport in a statistically homogeneous medium. It is worth noting that all physical media are intrinsically heterogeneous at some scale. We show that the micro-scale heterogeneities, which cannot be resolved deterministically, significantly influence the shape of transmittance.

In our model, the extinction of a medium is expressed as

$$\sigma_t(\mathbf{x}) = \sigma_m + \sigma_\mu(\mathbf{x}) \quad (9)$$

when viewed at the micro-scale. Here σ_m is a constant describing the overall density of the medium and $\sigma_\mu(\mathbf{x})$ is depicted by a random field whose 1D transects are from a 1D FGF with mean zero. Thus for a given position \mathbf{x} in the medium, $\sigma_t(\mathbf{x})$ defines a random variable

following the Gaussian distribution with its mean $\langle \sigma_t(\mathbf{x}) \rangle = \sigma_m$. Its variance is controlled by the FGF, and hence is influenced by its autocorrelation function.

Note that σ_μ is not modeled in 3D. This can simplify the calculation of statistical quantities for the FGF, especially for k -fBm, and is reasonable if we assume that the FGF is isotropic. In fact, any 1D transect of an isotropic 3D FGF is a 1D FGF with the same property (e.g., the Hurst parameter). In 1D, $C(H, d)$ reduces to

$$C(H, 1) = -\frac{1}{2\Gamma(2H+1)\sin(\pi H)}. \quad (10)$$

Please refer to the supplemental document for the detailed derivations. For brevity, we drop the dimension parameter (which is always 1) and rewrite $C(H, 1)$ as $C(H)$ henceforth.

4.2 Optical depth and line-averaged extinction

In simulating radiative transfer, we are primarily interested in the variability of the optical depth $\tau(\mathbf{x}, t)$ which is defined as

$$\tau(\mathbf{x}, t) = \int_0^t \sigma_t(\mathbf{x} + t'\boldsymbol{\omega}) dt' \quad (11)$$

where $\boldsymbol{\omega} \in S^2$ is a unit direction along which a photon is travelling and t denotes the free path between two interaction points. Holding t constant, $\tau(\mathbf{x}, t)$ is still a random variable following the Gaussian distribution.

Now consider the line-averaged extinction $\bar{\sigma}_t = [\int_0^t \sigma_t(\mathbf{x} + t'\boldsymbol{\omega}) dt']/t$ which is related to $\tau(\mathbf{x}, t)$ by

$$\tau(\mathbf{x}, t) = t\bar{\sigma}_t. \quad (12)$$

It's easy to verify that the ensemble-averaged $\bar{\sigma}_t$ over all possible realizations is

$$\langle \bar{\sigma}_t \rangle = \sigma_m. \quad (13)$$

The variance of $\bar{\sigma}_t$ is not trivially determined and varies significantly with respect to the Hurst parameter H . Defining $\mathbf{x}' = \mathbf{x} + t'\boldsymbol{\omega}$ and $\mathbf{x}'' = \mathbf{x} + t''\boldsymbol{\omega}$, the variance of $\bar{\sigma}_t$ is calculated as

$$\begin{aligned} \text{var}[\bar{\sigma}_t] &= \frac{1}{t^2} (\langle \tau(\mathbf{x}, t)^2 \rangle - \langle \tau(\mathbf{x}, t) \rangle^2) \\ &= \frac{1}{t^2} \int_0^t \int_0^t \langle \sigma_\mu(\mathbf{x}') \sigma_\mu(\mathbf{x}'') \rangle dt' dt'' \\ &= \frac{1}{t^2} \int_0^t \int_0^t \text{cov}(\mathbf{x}', \mathbf{x}'') dt' dt''. \end{aligned} \quad (14)$$

As seen, the variance of $\bar{\sigma}_t$ is tightly linked to the autocovariance function of the FGF.

Substituting the autocovariance function of pink noise (i.e., Eq. (4)) into the above formula allows us to obtain the variance of $\bar{\sigma}_t$ as (see the derivations in the supplemental document)

$$\text{var}_p[\bar{\sigma}_t] = -2C(H+1)S_w t^{2H}. \quad (15)$$

This reveals that the variance of $\bar{\sigma}_t$ has a power-law decay with respect to t when the media have a pink noise property at the micro-scale. This is evidenced in Fig. 7 where we conduct numerical experiments to show the correctness of Eq. (15). In the numerical experiments as sketched in Fig. 6, we first generate many realizations (1000 in our experiments) of 1D pink noise ($\sigma_m = 0$ and $S_w = 1$) for a given Hurst parameter H . Each realization contains $N = 2048$ samples and is defined on a scale from $t_{\min} = 0$ to $t_{\max} = L$ where L is the

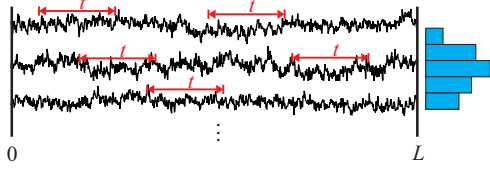


Fig. 6. Illustration of the numerical experiments. We collect line averages of length t from many realizations of an FGF and then build a histogram (blue bars) for the collected values.

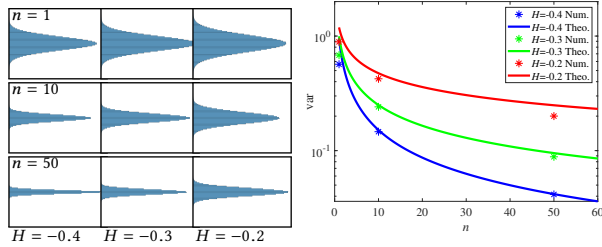


Fig. 7. Comparisons of $\text{var}_p[\bar{\sigma}_t]$ between numerical results (Num.) and theoretical results (Theo.). Left: Histograms of $\bar{\sigma}_t$ with different sample length n and Hurst parameter H collected in numerical experiments. Right: Plots of numerical results and theoretical results.

outer-scale of the field. We then collect the values of line-averaged extinction $\bar{\sigma}_t$ with sample length $n = 1, 10$, and 50 , respectively. The path length t and the sample length n are connected by $t = nL/N$. Histograms of different n and H are provided on the left panel of Fig 7. Variances of $\bar{\sigma}_t$ are estimated and plotted on the right panel. As seen, the theoretical results predicted by Eq. (15) achieve a reasonable agreement with numerical simulations.

For fBm, we may obtain a similar result using the same strategy. However, since it is non-stationary, the variance derived in this way will depend on the spatial position x . Please refer to the supplemental document for the detailed derivations and more discussions. A more reasonable expression of the variance can be derived according to the *one-point scale-independence* property of fBm. Specifically, if a random field satisfies one-point scale-independence, the variance of the line-averaged field and the variance of the field itself differ at most by a small amount on the order of a very small ratio [Davis and Marshak 2004]. With this property, we have

$$\frac{\text{var}[\bar{\sigma}_t]}{\text{var}[\sigma_t]} \approx 1. \quad (16)$$

Therefore, we can estimate the variance of $\bar{\sigma}_t$ in the case of fBm via the variance of σ_t itself. The variance of a field of fBm with an outer-scale L is calculated as $\text{var}_f = \int_0^L \text{cov}(x, x) dx / L$. Hence, we have

$$\text{var}_f[\bar{\sigma}_t] = \frac{1}{L} \int_0^L -2C(H)S_w|x|^{2H} dx = -\frac{2C(H)S_w}{2H+1}L^{2H}. \quad (17)$$

In this case, the variance keeps unchanged with respect to t and only varies according to H for a given outer-scale L . Note that an identical expression is derived in the supplemental document using the strategy of substituting Eq. (6) into Eq. (14).

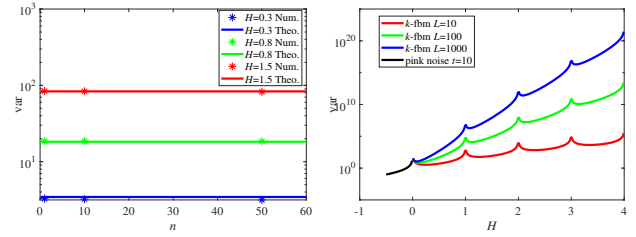


Fig. 8. Left: Comparisons of $\text{var}_{kf}[\bar{\sigma}_t]$ between numerical results (Num.) and theoretical results (Theo.). Right: Plots of $\text{var}_{kf}[\bar{\sigma}_t]$ with respect to the Hurst parameter H and outer-scale L .

The one-point scale-independence property also holds for k -fBm such that we can generalize the above calculation of variance to k -fBm (see the derivations in the supplemental document):

$$\text{var}_{kf}[\bar{\sigma}_t] = \frac{2C(H)S_w(-1)^k}{2H+1} \binom{2H-1}{k-1} L^{2H}. \quad (18)$$

Clearly, Eq. (17) is a special case of Eq. (18).

We verify the correctness of Eq. (17) and Eq. (18) by conducting the same numerical experiments as in the case of pink noise. The results (with $L = 10$ and $S_w = 1$) are shown in Fig. 8 left. Note that changing the sample length n and hence the path length t does not alter the value of $\text{var}_{kf}[\bar{\sigma}_t]$ for different H and the errors between numerical results and theoretical results are very subtle. This fulfills the property of one-point scale-independence for k -fBm. The comparison between Fig. 7 right and Fig. 8 left also tells that fields with a pink noise property (as well as a white noise property) are not one-point scale-independent.

All in all, the variance of $\bar{\sigma}_t$ is given by

$$\text{var}[\bar{\sigma}_t] = \begin{cases} S_w t^{-1} & H = -1/2 \\ S_p(H)t^{2H} & H \in (-1/2, 0) \\ S_{kf}(H)L^{2H} & H \in (k-1, k) \end{cases} \quad (19)$$

in which

$$S_p(H) = \frac{S_w}{\Gamma(2H+3)|\sin(\pi H)|} \quad (20)$$

$$S_{kf}(H) = \frac{S_w}{\Gamma(2H+2)|\sin(\pi H)|} \binom{2H-1}{k-1}. \quad (21)$$

Generally, $\text{var}[\bar{\sigma}_t]$ becomes large as the Hurst parameter H grows; see Fig. 8 right. Moreover, we see in Fig. 8 that the outer-scale L has a great impact on $\text{var}[\bar{\sigma}_t]$ when the field is modeled as k -fBm. Since the k -fBm has a never-ending growth of variance with respect to the spatial location, a large L will inevitably cause a large variance. In practice, L should be chosen larger than the extent of the medium.

4.3 Transmittance

We already know that traditional transmittance for participating media without spatial correlation is exponential. In this section, we derive transmittance for spatially-correlated media modeled by FGFs and show that their expressions of transmittance deviate greatly from the exponential function.

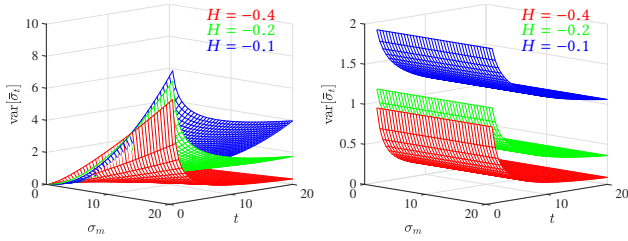


Fig. 9. The differences of $\text{var}[\bar{\sigma}_t]$ between the Bitterli model (left) and our model (right).

In our model of random media, we attempt to capture the ensemble-averaged transmittance:

$$\langle \text{Tr}(t) \rangle = \langle e^{-t\bar{\sigma}_t} \rangle = \int_0^\infty e^{-t\bar{\sigma}_t} p(\bar{\sigma}_t) d\bar{\sigma}_t \quad (22)$$

in which $p(\bar{\sigma}_t)$ is the probability density function (PDF) of the random variable $\bar{\sigma}_t$. Similar to previous work [Bitterli et al. 2018; Davis and Marshak 2004; Davis and Mineev-Weinstein 2011], $\langle \text{Tr}(t) \rangle$ can be evaluated by the characteristic function of $p(\bar{\sigma}_t)$, as this form of calculation in the above equation resembles the computation of the characteristic function. Moreover, we assume $\bar{\sigma}_t$ follows a strictly non-negative gamma distribution in calculating the ensemble-averaged transmittance. This avoids the problem that the extinction becomes negative in the Gaussian distribution. Based on the characteristic function of the gamma distribution, we derive analytic expressions of transmittance for media with different types of spatial correlations. Mathematically, we have

$$\varphi_{\bar{\sigma}_t}(t) = \int_{\mathbb{R}} e^{it\bar{\sigma}_t} p_{\Gamma(\alpha)}(\bar{\sigma}_t) d\bar{\sigma}_t = \left(1 - \frac{it\langle\bar{\sigma}_t\rangle}{\alpha}\right)^{-\alpha} \quad (23)$$

and the transmittance is given by

$$\text{Tr}(t) = \varphi_{\bar{\sigma}_t}(it) = \left(1 + \frac{\langle\bar{\sigma}_t\rangle}{\alpha(t)}\right)^{-\alpha(t)} \quad (24)$$

with

$$\alpha(t) = \frac{\langle\bar{\sigma}_t\rangle^2}{\text{var}[\bar{\sigma}_t]}. \quad (25)$$

Substituting $\langle\bar{\sigma}_t\rangle = \sigma_m$ and the expressions of $\text{var}[\bar{\sigma}_t]$ into Eq. (24) and Eq. (25), we obtain the transmittance for different types of spatial correlations. The expressions of various transmittance functions are listed in Table 1.

The transmittance functions derived above bear some similarity to the fractal noise model of Bitterli et al. [2018] which utilizes the Davis and Mineev-Weinstein model [2011] to derive closed-form expressions of non-exponential transmittance. Unlike our model in which the mean extinction σ_m and the variance of $\bar{\sigma}_t$ are decoupled, the mean extinction in their models has a great impact on the variance of $\bar{\sigma}_t$ except for some special cases. As shown in Fig. 9 left, $\text{var}[\bar{\sigma}_t]$ would grow dramatically as σ_m increases when H is relatively high. However, this is not the case in our model (Fig. 9 right). As an advantage, the decoupling of σ_m and $\text{var}[\bar{\sigma}_t]$ facilitates independent parameter tuning.

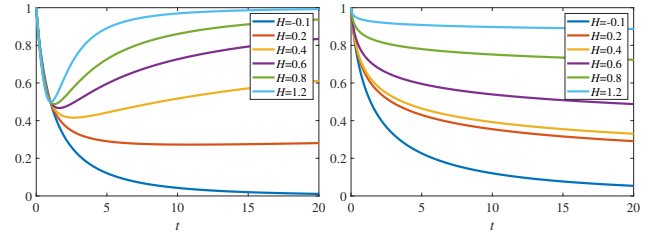


Fig. 10. Ensemble-averaged transmittance functions of the Bitterli model [Bitterli et al. 2018] (left) and our model (right). The Bitterli model will produce non-physical transmittance functions when $H > 0$.

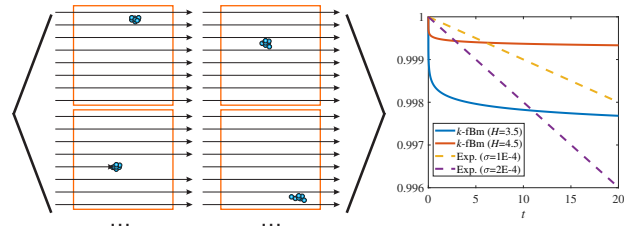


Fig. 11. Explanation of transparency as an extreme case of our model. Left: Illustration of light interacting with a random medium with a large H value. Right: The transmittance functions between our model (k -fBm) and the traditional model (Exp.) with low extinctions are quite different although similar transparent effects can be achieved.

Another main difference between the Bitterli model and our model is that we support a much wider scope of spatial correlations, ranging from short-range to long-range. In this context, the Bitterli model can be viewed as a subset of our model which only handles random media with a pink noise property ($-1/2 < H < 0$). Unnatural transmittance will be generated if H in the Bitterli model goes beyond the scope $-1/2 < H < 0$, as shown in Fig. 10.

4.4 Transparency

The tendency in Fig. 10 right also indicates that the transmittance function will degrade into a constant function $\text{Tr} = 1$ when H is sufficient large. In this extreme case, we actually get a transparent medium. This phenomenon can be explained physically in Fig. 11. A very large value of H implies that primary particles (blue dots) in a random medium start to tightly bound together, forming aggregates or even agglomerates. Obviously, this medium has a very strong spatial correlation, i.e., most of the particles in the medium share the same position. When a light beam hits such a medium, it rarely hits an aggregated scatter in the medium, leading to an almost constant transmittance function after performing ensemble averaging. Although setting the extinction coefficient to a small value may also achieve the transparent effect, the variation trends of transmittance decay between these two models are quite different as illustrated in Fig. 11 right. Visual comparisons of synthesized images provided in Fig. 15 and Fig. 16 further showcase the differences between these two settings.

Table 1. Expressions of transmittance functions and the corresponding derivatives for different types of spatial correlations at the micro-scale.

	H	$\text{var}[\bar{\sigma}_t]$	Transmittance $\text{Tr}(t)$	$\partial\alpha(t)/\partial t$	$\partial\beta(t)/\partial t$
White noise	$-1/2$	$S_w t^{-1}$	$\left(1 + \frac{S_w}{\sigma_m}\right)^{-\frac{\sigma_m^2}{S_w} t}$	$\frac{\sigma_m^2}{S_w}$	0
Pink noise	$(-1/2, 0)$	$S_p(H)t^{2H}$	$\left(1 + \frac{S_p(H)}{\sigma_m} t^{2H+1}\right)^{-\frac{\sigma_m^2}{S_p(H)} t^{-2H}}$	$\frac{-2H\sigma_m^2}{S_p(H)} t^{-2H-1}$	$\frac{(2H+1)S_p(H)}{\sigma_m} t^{2H}$
k -fBm	$(k-1, k)$	$S_{kf}(H)L^{2H}$	$\left(1 + \frac{S_{kf}(H)}{\sigma_m} L^{2H} t\right)^{-\frac{\sigma_m^2}{S_{kf}(H)} L^{-2H}}$	0	$\frac{S_{kf}(H)}{\sigma_m} L^{2H}$

5 RENDERING TECHNIQUES

So far, we have derived the expressions of transmittance for various spatially-correlated media modeled by FGFs. We now use these transmittance functions in a physically-based rendering framework. Classical RTE suffers from unacceptable non-conservation of energy when these non-exponential transmittance functions are used directly. To alleviate this problem, an energy-conserving RTE framework is proposed by Bitterli et al. [2018].

Similar to the classical RTE, the integral form of this framework, i.e., the *Volume Rendering Equation* (VRE), is expressed as

$$L_o(\mathbf{x}, \omega) = \int_0^t T(\mathbf{x}, \mathbf{x}') \Lambda \sigma_t(\mathbf{x}') L_i(\mathbf{x}', \omega) dt' + T(\mathbf{x}, \mathbf{x}_s) L_s(\mathbf{x}_s, \omega) \quad (26)$$

where

$$L_i(\mathbf{x}', \omega) = \int_{S^2} L(\mathbf{x}', \omega') f_p(\omega, \omega') d\omega' \quad (27)$$

is the in-scattered radiance, \mathbf{x}_s is a point on the surface, and f_p denotes the phase function. To ensure energy conservation, the transport kernel $T(\mathbf{x}, \mathbf{x}')$ should be set to

$$T(\mathbf{x}, \mathbf{x}') = \begin{cases} \text{Tr}(\mathbf{x}, \mathbf{x}') & \text{if } \mathbf{x}' \text{ on surface} \\ -\frac{1}{\sigma_t(\mathbf{x}')} \frac{\partial \text{Tr}(\mathbf{x}, \mathbf{x}')}{\partial t} & \text{if } \mathbf{x}' \text{ in media.} \end{cases} \quad (28)$$

It is easy to verify that if the medium fluctuates as white noise at the micro-scale the transport kernel is exponential and the energy-conserving VRE converges to the classical VRE. See more discussions in the supplemental document.

Although energy-conserving, this form of VRE is not currently reciprocal. To ensure reciprocity, some term rearrangements and approximations are required in the path integral formulation [Bitterli et al. 2018].

Eq. (28) shows that the transport kernel relies on the partial derivative of the transmittance. Since every transmittance function in our model has the form $\text{Tr}(t) = \beta(t)^{-\alpha(t)}$, its partial derivative with respect to t is calculated as

$$\frac{\partial \text{Tr}(t)}{\partial t} = \text{Tr}(t) \left[-\frac{\partial \alpha(t)}{\partial t} \ln \beta(t) - \frac{\partial \beta(t)}{\partial t} \frac{\alpha(t)}{\beta(t)} \right]. \quad (29)$$

The expressions of $\partial\alpha(t)/\partial t$ and $\partial\beta(t)/\partial t$ are also listed in Table 1.

5.1 Importance sampling

When solving the above VRE via Monte Carlo-based integration, it is central to design a proper importance sampling scheme for sampling the free path t [Novák et al. 2018]. For the case of k -fBm,

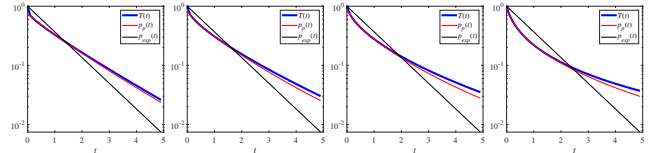


Fig. 12. Sampling PDF comparisons of our method against the traditional exponential distribution. From left to right, $H = -0.4, -0.3, -0.2,$ and $-0.1,$ respectively.




Fig. 13. Equal-sampling-rate (64 spp) comparisons of our analytical path sampling method against traditional path sampling method using an exponential distribution on rendering a homogeneous medium with fixed mean extinction $\sigma_m = 2$ and different degrees of correlations. From top left to bottom right, $H = -0.3, -0.2, -0.1, 0.5, 1.5,$ and $2.5,$ respectively. Quantitative measures in terms of RMSE are shown in the brackets.

the sampling PDF is relative easy to obtain since the transport kernel reduces to

$$T(\mathbf{x}, \mathbf{x}') = \begin{cases} \left(1 + \frac{\sigma_m}{\alpha} t\right)^{-\alpha} & \text{if } \mathbf{x}' \text{ on surface} \\ \left(1 + \frac{\sigma_m}{\alpha} t\right)^{-\alpha-1} & \text{if } \mathbf{x}' \text{ in media} \end{cases} \quad (30)$$

with $\alpha = L^{-2H} \sigma_m^2 / S_{kf}(H)$ which is actually a constant. In this case, we sample the free path t with the PDF:

$$p_{kf}(t) = \sigma_m \left(1 + \frac{\sigma_m}{\alpha} t\right)^{-\alpha-1} \quad (31)$$

according to the transport kernel in media.

For the case of pink noise, sampling according to the transport kernel is impractical since the inverse of the cumulative distribution function (CDF) lacks a closed-form expression. Observing that the

transport kernel of pink noise lies between that of white noise and k -fBm, we suggest use the following sampling PDF:

$$p_p(t) = w_1 c_1 \left(1 + \frac{S_p(H)}{\sigma_m} t\right)^{-\frac{\sigma_m^2}{S_p(H)} - 1} + w_2 c_2 \left(1 + \frac{S_p(H)}{\sigma_m} t\right)^{-\frac{\sigma_m^2}{S_p(H)} t} \quad (32)$$

in which $c_1 = \sigma_m$ and $c_2 = \frac{\sigma_m^2}{S_p(H)} \ln\left(1 + \frac{S_p(H)}{\sigma_m}\right)$ are two scaling constants ensuring that the PDF is properly normalized. The weights w_1 and w_2 are set to $(2H + 1)$ and $-2H$, respectively. Fig. 12 verifies that this sampling PDF (red curves) tightly matches the transport kernel in media (blue curves). As this sampling PDF comprises two components, an additional random variable is required to select the component first according to the weights w_1 and w_2 .

Fig. 13 compares our sampling method to the traditional sampling method using an exponential function on rendering a homogeneous medium with various spatial correlations. The sampling rate for all algorithms is fixed to 64 spp. As expected, our method produces visibly smoother images containing less firefly artifacts for all types of microscopic fluctuations of media. Quantitative comparisons in terms of root mean squared error (RMSE) further prove the effectiveness of the proposed sampling method.

5.2 Heterogeneous media

Currently, we borrow the same idea from the Bitterli model [Bitterli et al. 2018] to support macro-scale heterogeneity. Assuming a constant correlation throughout the medium, this strategy scales the free path t to match the macroscopic density locally. Specifically, denoting σ_M as the average extinction of the heterogeneous medium at the macro-scale, the transmittance function is modified as

$$\text{Tr}(\mathbf{x}, t) = \langle \text{Tr}(\tau(\mathbf{x}, t) / \sigma_M) \rangle. \quad (33)$$

When handling heterogeneous media, we first sample an optical depth τ using the above sampling strategy and then search for the corresponding free path t via regular tracking or ray marching. A prevailing alternative is to use delta tracing [Novák et al. 2018; Raab et al. 2008] which is more efficient and unbiased in general. Unfortunately, delta tracing, which is based on rejection sampling, requires that the transmittance to be exponential [Bitterli et al. 2018]. Although our transmittance functions do not satisfy this requirement, we can still perform delta tracing for k -fBm since the PDF of $\bar{\sigma}_t$ is independent with t in this case. Looking back to the calculation of ensemble-averaged transmittance in Eq. (22), it is possible to approximate $\langle \text{Tr}(t) \rangle$ with the Monte Carlo estimator $[\sum_i^N \exp(-t\bar{\sigma}_{t_i})] / N$ where the samples $\bar{\sigma}_{t_i}$ are drawn from the gamma distribution:

$$\bar{\sigma}_t \propto \Gamma\left(\frac{\sigma_m^2}{S_{kf}(H)L^{2H}}, \frac{\sigma_m}{S_{kf}(H)L^{2H}}\right). \quad (34)$$

Now, traditional delta tracing applies for each sample $\bar{\sigma}_{t_i}$.

6 RESULTS

We have implemented the proposed technique on top of the Mitsuba physically based renderer [Jakob 2010] and successfully used it to simulate different types of spatially-correlated media. The degree

of correlations of the media is mainly determined by the Hurst parameter H . For other parameters, L is set to 10, S_w is set to 1 and the phase function is set to isotropic unless mentioned otherwise. All synthesized images are created by a PC with Intel Core i7-6900K CPU and 16G RAM.

Rendering of homogeneous media. In Fig. 14 we render a homogeneous medium with the same overall density σ_m but different spatial correlations. The Hurst parameter H ranges from -0.4 (pink noise) to 1.5 (2-fBm) which means the correlations become longer gradually. Compared with the reference generated by the classical exponential transmittance, spatial correlations, either a pink noise type or a k -fBm type, yield brighter appearances due to slower-than-exponential attenuations. It achieves a near transparent effect when H is very large (e.g., $H = 1.5$) since the transmittance function tends to be unity in this case. Oppositely, when H is very close to $-1/2$, e.g., $H = -0.4$, the result is almost identical to that of the classical transport with exponential falloff. This implies that our method converges to the classical transport when the FGF is white noise. Here, we set the phase function to isotropic. For the results of anisotropic phase functions, please refer to the supplemental document.

Rendering of heterogeneous media. To test the versatility of our technique, we render three different forms of heterogeneous clouds under side lighting, front lighting and back lighting, respectively. Irrespective of lighting condition, the appearance becomes softer as H increases and approaches transparency steadily. Due to the short-range correlation of pink noise, the difference between the exponential transmittance and the non-exponential transmittance is not that remarkable. For k -fBm, media become uneven as shown in Fig. 5, i.e., regions where particles have an excess of close neighbors are surrounded by regions devoid of particles. This allows light penetrate further on average, leading to transparency in some extent. Although it is also possible to get transparent effects via scaling the original density with a very small value in the classical RTE, the results are quite different in these two cases as evidenced in Fig. 16. As seen, lowering the overall densities only makes the media thin globally while increasing H leads a local change of transparency. Fig. 17 further shows that significant appearance changes exist in different spatially-correlated media even after normalization such that the transmittances share the same mean free path. In addition, as H increases the media tend to create less scattering interactions, which reduces the total rendering cost.

Comparison to previous work. In Fig. 18 we compare to the GBE proposed by Jarabo et al. [2018] in simulating non-exponential transmittance of a homogeneous medium. Although their transmittance is different from ours, a similar appearance can be obtained with carefully parameter adjusting. Unfortunately, the GBE currently does not support macroscopic heterogeneous in practice, which significantly limits its applicability. We also compare to the Bitterli model [Bitterli et al. 2018] in Fig. 19. Fixing other parameters, the appearance changes are subtle with respect to H in that model, since they only support random media of a pink noise type, i.e., $H \in (-1/2, 0)$. Because of the long-range correlation of k -fBm, a much wider range of appearances can be achieved by our model through adjusting the Hurst parameter H only.

Impact of L and S_w . In Fig. 20, we analyze the impact of L and S_w on the effects of correlations by rendering a simple object under




Fig. 14. Effects of spatial correlations in homogeneous media. The Hurst parameter H ranges from -0.4 (short-range correlation) to 1.5 (long-range correlation) which yields a wide range of appearances. Other medium parameters are $\sigma_m = 1.5$ and $\Lambda = 1$.

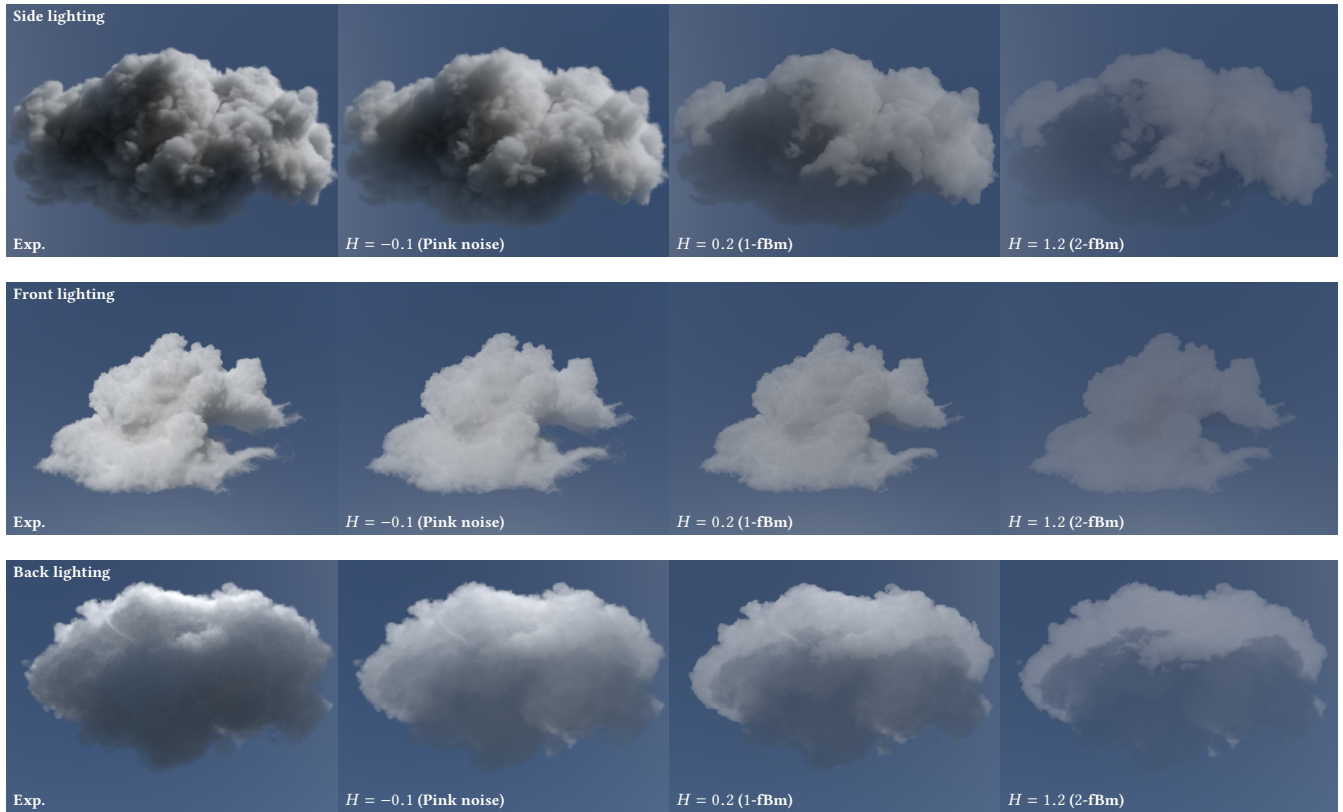


Fig. 15. Effects of spatial correlations in heterogeneous media under different lighting configurations. In general, the appearance becomes softer as H increases and approaches transparency steadily. The albedo Λ is set to 1 in all scenes.

a variety of configurations. Recall that in our model L is the outer-scale of the random field while S_w is the PSD of the underlying white noise. L only influences the k -fBm type media and has no effect on the pink-noise type media since pink noise is stationary. To make our model physically plausible, we ensure that L is larger than the extent of the media. However, if L is very large, the medium will approach complete transparency very fast as H increases. The parameter S_w is similar to the parameter “ C ” in the Bitterli model

which generally determines the energy level of the field. As S_w increases, the variance of line-averaged extinction increases and the attenuation of light tends to be small.

Rendering of a complex scene. In Fig. 1, we render a complex scene containing several spatially-correlated media, either homogeneous or heterogeneous, using our method (left). The Hurst parameter H ranges from -0.4 (pink noise) to 1.8 (2-fBm). As expected, the appearance is quite different from the one generated by the



Fig. 16. The nearly transparent effects achieved by lowering the overall densities of the media. In these scenes, we set the density of each cloud to one percent of the original density. Note the differences compared against our model with a large value of H (the right-most column of Fig. 15).



Fig. 17. The same cloud as in the last row of Fig. 15 is rendered with different degrees of spatial correlations, but normalized such that all three transmittances have the same mean free path.



Fig. 18. Comparing to the GBE [Jarabo et al. 2018] in rendering a spatially-correlated, homogeneous medium. To closely match the appearance generated by the GBE, we choose $H = 1.2$ in this scene. A reference rendered with the traditional exponential transmittance is also provided.

classical exponential light transport (right). Since long-range spatial correlations allow the light penetrate deeper, the appearance will be softer and more details in the media may become clear, especially for the heterogeneous media. In addition, spatial correlations also enhance the effect of caustics.

7 LIMITATIONS AND FUTURE WORK

Negative correlations. The correlations simulated by our technique using Hurst parameter $H > -1/2$ are widely recognized as positive correlations. In this vein, many particles huddle together, leaving empty regions in the media and hence letting more light pass through. On the opposite side, there are negative correlations [Shaw et al. 2002] in which the particles distribute more uniform than white noise⁵, leading to faster-than-exponential attenuations. Actually, our model can achieve negative correlations with super-exponential extinction when setting H to a value smaller than $-1/2$. However, this is empirical and lacks a mathematically rigorous explanation. See more discussions in the supplemental document.

⁵We actually get a blue noise type in this case.




Fig. 19. Compared with the Bitterli model [Bitterli et al. 2018] (top row), our model (bottom row) can achieve a much wider range of appearances through adjusting the Hurst parameter H only.

	$S_w = 0.5$	$S_w = 1.0$	$S_w = 1.5$	$S_w = 2.0$	$S_w = 2.5$
$L = 10$					
$L = 50$					
$L = 100$					

Fig. 20. Impact of L and S_w on the appearance of spatially-correlated media. Other medium parameters are $H = 0.5$, $\sigma_m = 5$ and $\Lambda = 1$.

Unbiased delta tracking. Although unbiased delta tracking works for the case of k -fBm in our model, it can not be easily extended to the general case. Delta tracking requires that the transmittance is exponential such that it is multiplicative along a ray. This is not satisfied by general transmittance functions. Considering the efficiency of unbiased delta tracking, it is quite attractive to develop proper importance sampling solutions for handling arbitrary non-exponential transmittance functions, and we leave it as future work.

Visible aggregates. As we mentioned earlier in Sec. 4, the appearance of spatially-correlated media become transparent as H increases gradually. This is due to the fact that tiny particles in a medium form larger aggregates, leading to the almost unity transmittance. If the original medium is dense enough, the aggregates will be so large that their shapes are even visible. In this case, it is possible that geometric optics have to be employed to capture the appearance of each bulky particles. However, to ensure conservation of mass, the relative densities of bulky particles and original

tiny particles should be taken into consideration. How to model the relative densities and how to enable smooth transitions from media to visible aggregates require further investigation.

8 CONCLUSION

We have presented a new framework to model and render spatially-correlated media in a physically-based manner. Built upon the fractional Gaussian field (FGF), our modeling strategy of the spatially-correlated media is capable of reproducing a much wider variety of media exhibiting diverse appearances. We have derived the transmittance functions for different types of spatial correlations from short-range to long-range, and demonstrated their practicality in an energy-conserving RTE framework. We also have discussed a special transparent effect that our model converges to as the Hurst parameter of the FGF continuously increases. Experiments and results show that our method is able to significantly expand the admissible regime of spatially-correlated media benefitting from the powerful functionality of the FGF. We believe that our method is another important step towards fully understanding, representing and connecting the micro and macro worlds, and we also hope this paper will increase the exposure of the FGF in the rendering community.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their valuable comments. This work was supported in part by the National Key Research and Development Program of China (2018YFB1004901) and NSFC (61502223 and 61772257).

REFERENCES

- Guillaume Bal and Wenjia Jing. 2011. Fluctuation theory for radiative transfer in random media. *Journal of Quantitative Spectroscopy and Radiative Transfer* 112, 4 (2011), 660 – 670.
- H. W. Barker, J. N. S. Cole, J.-J. Morcrette, R. Pincus, P. Räisänen, K. von Salzen, and P. A. Vaillancourt. 2008. The Monte Carlo Independent Column Approximation: an assessment using several global atmospheric models. *Quarterly Journal of the Royal Meteorological Society* 134, 635 (2008), 1463–1478.
- Benedikt M. Bitterli, Srinath Ravichandran, Thomas Müller, Magnus Wrenninge, Jan Novák, Steve Marschner, and Wojciech Jarosz. 2018. A radiative transfer framework for non-exponential media. *ACM Trans. Graph.* (2018).
- Robert Cahalan. 1994. Bounded cascade clouds: Albedo and effective thickness. *Nonlinear Processes in Geophysics* 1 (01 1994).
- Eva Cerezo, Frederic Pérez, Xavier Pueyo, Francisco J. Seron, and François X. Sillion. 2005. A survey on participating media rendering techniques. *The Visual Computer* 21, 5 (01 Jun 2005), 303–328.
- Subrahmanyan. Chandrasekhar. 1960. *Radiative transfer*. Dover.
- Anthony B. Davis. 2006. Effective Propagation Kernels in Structured Media with Broad Spatial Correlations, Illustration with Large-Scale Transport of Solar Photons Through Cloudy Atmospheres. In *Computational Methods in Transport*, Frank Graziani (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 85–140.
- Anthony B. Davis and Alexander Marshak. 2004. Photon propagation in heterogeneous optical media with spatial correlations: enhanced mean-free-paths and wider-than-exponential free-path distributions. *Journal of Quantitative Spectroscopy and Radiative Transfer* 84, 1 (2004), 3 – 34.
- Anthony B. Davis and Mark Mineev-Weinstein. 2011. Radiation propagation in random media: From positive to negative correlations in high-frequency fluctuations. *Journal of Quantitative Spectroscopy and Radiative Transfer* 112 (2011), 632–645.
- Anthony B. Davis and Feng Xu. 2014. A Generalized Linear Transport Model for Spatially Correlated Stochastic Media. *Journal of Computational and Theoretical Transport* 43 (2014), 474–514.
- Anthony B. Davis, Feng Xu, and David J. Diner. 2018. Generalized radiative transfer theory for scattering by particles in an absorbing gas: Addressing both spatial and spectral integration in multi-angle remote sensing of optically thin aerosol layers. *Journal of Quantitative Spectroscopy and Radiative Transfer* 205 (2018), 148 – 162.
- Eugene d'Eon. 2014. Computer graphics and particle transport: our common heritage, recent cross-field parallels and the future of our rendering equation. In *Digital Production Symposium 2014*.
- Eugene d'Eon. 2018a. A reciprocal formulation of non-exponential radiative transfer. 1: Sketch and motivation. arXiv:1803.03259 <https://arxiv.org/abs/1803.03259>.
- Eugene d'Eon. 2018b. A reciprocal formulation of non-exponential radiative transfer. 2: Monte Carlo estimation and diffusion approximation. arXiv:1809.05881 <https://arxiv.org/abs/1809.05881>.
- Eugene d'Eon and Geoffrey Irving. 2011. A Quantized-diffusion Model for Rendering Translucent Materials. In *ACM SIGGRAPH 2011 Papers (SIGGRAPH '11)*, 56:1–56:14.
- Adrian Doicu, Dmitry S. Efremenko, Diego Loyola, and Thomas Trautmann. 2014. Approximate models for broken clouds in stochastic radiative transfer theory. *Journal of Quantitative Spectroscopy and Radiative Transfer* 145 (2014), 74 – 87.
- Craig Donner and Henrik Wann Jensen. 2005. Light Diffusion in Multi-layered Translucent Materials. In *ACM SIGGRAPH 2005 Papers (SIGGRAPH '05)*, 1032–1039.
- Bertrand Duplantier, Rémi Rhodes, Scott Sheffield, and Vincent Vargas. 2017. *Log-correlated Gaussian Fields: An Overview*. Springer International Publishing, 191–216.
- Jonathan Dupuy, Eric Heitz, and Eugene d'Eon. 2016. Additional Progress Towards the Unification of Microfacet and Microflake Theories. In *Eurographics Symposium on Rendering - Experimental Ideas & Implementations*, Elmar Eisemann and Eugene Fiume (Eds.). The Eurographics Association.
- P. Flandrin. 1989. On the spectrum of fractional Brownian motions. *IEEE Transactions on Information Theory* 35, 1 (1989), 197–199.
- Julian Fong, Magnus Wrenninge, Christopher Kulla, and Ralf Habel. 2017. Production Volume Rendering: SIGGRAPH 2017 Course. In *ACM SIGGRAPH 2017 Courses (SIGGRAPH '17)*, 2:1–2:79.
- Jeppe Revall Frisvad, Toshiya Hachisuka, and Thomas Kim Kjeldsen. 2014. Directional Dipole Model for Subsurface Scattering. *ACM Trans. Graph.* 34, 1, Article 5 (Dec. 2014), 12 pages.
- Toshiya Hachisuka, Wojciech Jarosz, Iliyan Georgiev, Anton Kaplanyan, and Derek Nowrouzezahrai. 2013. State of the Art in Photon Density Estimation. In *ACM SIGGRAPH Asia Courses*.
- Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. 2015. The SGGX Microflake Distribution. *ACM Trans. Graph.* 34, 4 (July 2015), 48:1–48:11.
- Wenzel Jakob. 2010. Mitsuba renderer. <http://www.mitsuba-renderer.org>.
- Wenzel Jakob, Adam Arbree, Jonathan T. Moon, Kavita Bala, and Steve Marschner. 2010. A Radiative Transfer Framework for Rendering Materials with Anisotropic Structure. *ACM Trans. Graph.* 29, 4 (July 2010), 53:1–53:13.
- Adrian Jarabo, Carlos Aliaga, and Diego Gutierrez. 2018. A Radiative Transfer Framework for Spatially-correlated Materials. *ACM Trans. Graph.* 37, 4 (July 2018), 83:1–83:13.
- Wojciech Jarosz, Derek Nowrouzezahrai, Iman Sadeghi, and Henrik Wann Jensen. 2011. A Comprehensive Theory of Volumetric Radiance Estimation Using Photon Points and Beams. *ACM Trans. Graph.* 30, 1 (Feb. 2011), 5:1–5:19.
- Henrik Wann Jensen and Per H. Christensen. 1998. Efficient Simulation of Light Transport in Scenes with Participating Media Using Photon Maps. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*, 311–320.
- Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. 2001. A Practical Model for Subsurface Light Transport. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01)*, 511–518.
- J. T. Kajiya and T. L. Kay. 1989. Rendering Fur with Three Dimensional Textures. In *Proceedings of the 16th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '89)*, 271–280.
- Pramook Khungurn, Daniel Schroeder, Shuang Zhao, Kavita Bala, and Steve Marschner. 2015. Matching Real Fabrics with Micro-Appearance Models. *ACM Trans. Graph.* 35, 1 (Dec. 2015), 1:1–1:26.
- Alexander B. Kostinski. 2001. On the extinction of radiation by a homogeneous but spatially correlated random medium. *J. Opt. Soc. Am. A* 18, 8 (Aug 2001), 1929–1933.
- Jaroslav Krivánek, Iliyan Georgiev, Toshiya Hachisuka, Petr Vévoda, Martin Šik, Derek Nowrouzezahrai, and Wojciech Jarosz. 2014. Unifying Points, Beams, and Paths in Volumetric Light Transport Simulation. *ACM Trans. Graph.* 33, 4 (July 2014), 103:1–103:13.
- Mateusz Kwaśnicki. 2017. Ten equivalent definitions of the fractional laplace operator. *Fractional Calculus and Applied Analysis* 20, 1 (2017).
- Edward W. Larsen and Richard Vasques. 2011. A generalized linear Boltzmann equation for non-classical particle transport. *Journal of Quantitative Spectroscopy and Radiative Transfer* 112, 4 (2011), 619 – 631.
- Michael L. Larsen and Aaron S. Clark. 2014. On the link between particle size and deviations from the Beer Lambert Bouguer law for direct transmission. *Journal of Quantitative Spectroscopy and Radiative Transfer* 133 (2014), 646 – 651.
- Asad Lodhia, Scott Sheffield, Xin Sun, and Samuel S. Watson. 2016. Fractional Gaussian fields: A survey. *Probability Surveys* 13 (2016), 1–56.
- B. Mandelbrot and J. Van Ness. 1968. Fractional Brownian Motions, Fractional Noises and Applications. *SIAM Rev.* 10, 4 (1968), 422–437.
- Johannes Meng, Marios Papas, Ralf Habel, Carsten Dachsbacher, Steve Marschner, Markus Gross, and Wojciech Jarosz. 2015. Multi-scale Modeling and Rendering of

- Granular Materials. *ACM Trans. Graph.* 34, 4 (July 2015), 49:1–49:13.
- Jonathan T. Moon, Bruce Walter, and Stephen R. Marschner. 2007. Rendering Discrete Random Media Using Precomputed Scattering Solutions. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques (EGSR'07)*. 231–242.
- Thomas Müller, Marios Papas, Markus Gross, Wojciech Jarosz, and Jan Novák. 2016. Efficient Rendering of Heterogeneous Polydisperse Granular Media. *ACM Trans. Graph.* 35, 6 (Nov. 2016), 168:1–168:14.
- Jan Novák, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz. 2018. Monte Carlo Methods for Volumetric Light Transport Simulation. *Computer Graphics Forum (Proceedings of Eurographics - State of the Art Reports)* 37, 2 (may 2018).
- E. Perrin, R. Harba, C. Berzin-Joseph, I. Iribarren, and A. Bonami. 2001. n -th-order fractional Brownian motion and fractional Gaussian noises. *IEEE Transactions on Signal Processing* 49, 5 (May 2001), 1049–1059.
- Matthias Raab, Daniel Seibert, and Alexander Keller. 2008. Unbiased Global Illumination with Participating Media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*, Alexander Keller, Stefan Heinrich, and Harald Niederreiter (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 591–605.
- I. S. Reed, P. C. Lee, and T. K. Truong. 1995. Spectral representation of fractional Brownian motion in n dimensions and its properties. *IEEE Transactions on Information Theory* 41, 5 (1995), 1439–1451.
- L. F. Rojas-Ochoa, J. M. Mendez-Alcaraz, J. J. Sáenz, P. Schurtenberger, and F. Scheffold. 2004. Photonic Properties of Strongly Correlated Colloidal Liquids. *Phys. Rev. Lett.* 93 (2004), 073903. Issue 7.
- Gennady Samorodnitsky. 2007. Long Range Dependence. *Foundations and Trends in Stochastic Systems* 1, 3 (2007), 163–257.
- Kai Schröder, Reinhard Klein, and Arno Zinke. 2011. A Volumetric Approach to Predictive Rendering of Fabrics. In *Proceedings of the Twenty-second Eurographics Conference on Rendering (EGSR '11)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 1277–1286.
- Kai Schröder, Arno Zinke, and Reinhard Klein. 2014. Image-Based Reverse Engineering and Visual Prototyping of Woven Cloth. *IEEE Transactions on Visualization and Computer Graphics* PP, 99 (2014).
- Raymond A Shaw, Alexander B Kostinski, and Daniel D Lanterman. 2002. Super-exponential extinction of radiation in a negatively correlated random medium. *Journal of Quantitative Spectroscopy and Radiative Transfer* 75, 1 (2002), 13 – 20.
- Leung Tsang, Jin Au Kong, Kung-Hau Ding, and Chi On Ao. 2002. *Scattering of Electromagnetic Waves: Numerical Simulations*. John Wiley & Sons.
- Eric Veach. 1997. *Robust Monte Carlo Methods for Light Transport Simulation*. Ph.D. Dissertation. Stanford, CA, USA.
- Magnus Wrenninge, Ryusuke Villemin, and Christophe Hery. 2017. *Path Traced Sub-surface Scattering using Anisotropic Phase Functions and Non-Exponential Free Flights*. Pixar Technical Memo 17-07. Pixar Inc.
- Jiahua Zhang, G. Baci, Dejun Zheng, Cheng Liang, Guiqing Li, and Jinlian Hu. 2013. IDSS: A Novel Representation for Woven Fabrics. *IEEE Transactions on Visualization and Computer Graphics* 19, 3 (2013), 420–432.
- Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2011. Building Volumetric Appearance Models of Fabric Using Micro CT Imaging. *ACM Trans. Graph.* 30, 4 (July 2011), 44:1–44:10.