From microfacets to participating media: A unified theory of light transport with stochastic geometry

DARIO SEYB, Dartmouth College, USA
EUGENE D'EON, NVIDIA, New Zealand
BENEDIKT BITTERLI, NVIDIA, USA
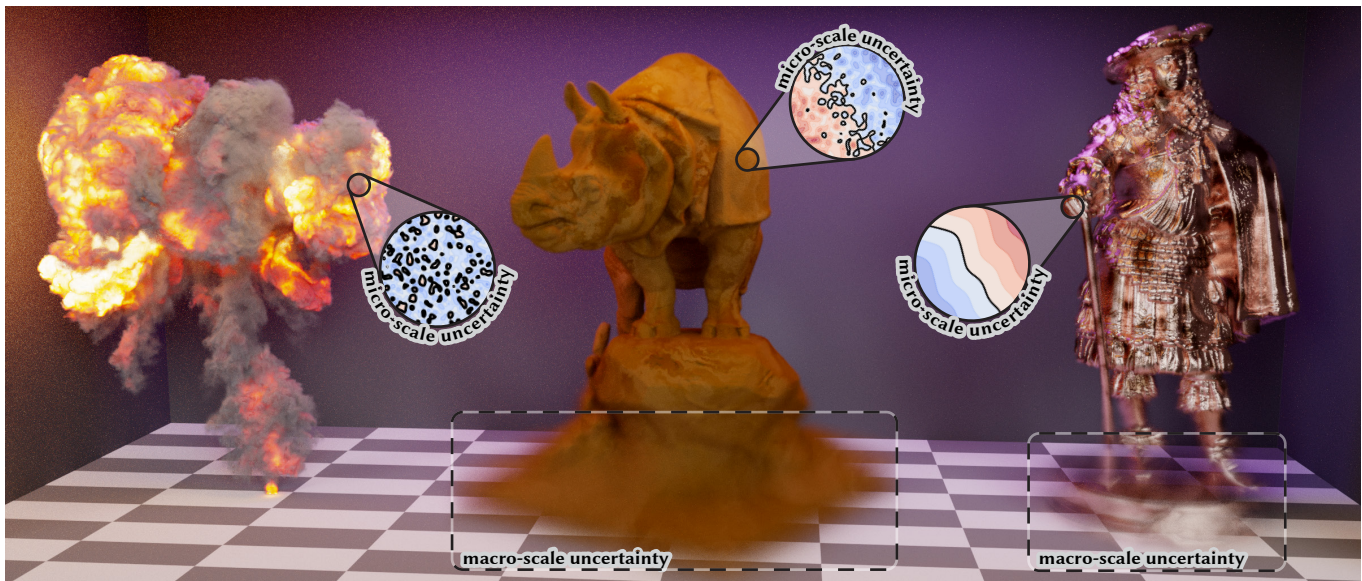WOJCIECH JAROSZ, Dartmouth College, USA

Fig. 1. We develop a theory of light transport for scenes with stochastic implicit surfaces [Dragiev et al. 2011; Sellán and Jacobson 2022; Williams and Fitzgibbon 2006] and show that this results in an expressive range of appearance behaviors that covers microfacet surfaces, classical participating media, and a novel continuum in between. Each object shown above is completely described by a 3D volume that encodes the mean and covariance kernel of a non-stationary Gaussian process and is rendered using the same rendering algorithm, agnostic of its "appearance type". The insets show example realizations of the process at the highlighted points in the scene. Note how the appearance of the stochastic geometry transitions from volumetric (left) to hard-surface (right) as the correlations in the process are strengthened. Increasing the *variance* of the process allows us to visualize uncertainty at the macro-scale (bottom).

Stochastic geometry models have enjoyed immense success in graphics for modeling interactions of light with complex phenomena such as participating media, rough surfaces, fibers, and more. Although each of these models operates on the same principle of replacing intricate geometry by a random process and deriving the average light transport across all instances thereof, they are each tailored to one specific application and are fundamentally distinct. Each type of stochastic geometry present in the scene is firmly encapsulated in its own appearance model, with its own statistics and light transport average, and no cross-talk between different models or deterministic and stochastic geometry is possible.

In this paper, we derive a theory of light transport on *stochastic implicit surfaces*, a geometry model capable of expressing deterministic geometry,

microfacet surfaces, participating media, and an exciting new continuum in between containing aggregate appearance, non-classical media, and more. Our model naturally supports *spatial correlations*, missing from most existing stochastic models.

Our theory paves the way for tractable rendering of scenes in which all geometry is described by the same stochastic model, while leaving ample future work for developing efficient sampling and rendering algorithms.

CCS Concepts: • **Computing methodologies** → *Volumetric models*.

Additional Key Words and Phrases: volumetric light transport, stochastic processes, implicit surfaces

Authors' addresses: Dario Seyb, dario.r.seyb.gr@dartmouth.edu, Dartmouth College, USA; Eugene d'Eon, edeon@nvidia.com, NVIDIA, New Zealand; Benedikt Bitterli, bbitterli@nvidia.com, NVIDIA, USA; Wojciech Jarosz, wojciech.k.jarosz@dartmouth.edu, Dartmouth College, USA.

1 INTRODUCTION

Computer graphics has seen significant progress over the past few decades, with advances in light transport, material fidelity, and

scene complexity greatly expanding the set of images that can be rendered. However, scene representations have evolved only slowly since their inception. Today, there is still a dichotomy that permeates through light transport theory where scene components are treated either as participating media or as hard (deterministic) surfaces. Beyond being intellectually unsatisfying, this distinction creates real problems in practice. Many rendering concepts invariably create effects that lie somewhere between surfaces and media. For example, inverse rendering and real-world scene reconstruction generate surfaces that contain some level of uncertainty due to measurement errors. Filtering and level-of-detail (LOD) of complex scenes turn hard surfaces into volumetric effects that still feature surface-like reflectance and correlations that are poorly represented by classical media. Classic scene representations can handle these effects only via specialized models or not at all, and recent efforts to unify these models [Dupuy et al. 2016] have not come to clear conclusions.

In this paper, we investigate a principled treatment of light transport in scenes containing *stochastic implicit surfaces* (SIS)[Gamito 2009]—a powerful modeling framework that represents surfaces as the zero crossings of a random field. In particular, we identify Gaussian processes (GPs) as a useful building block. By manipulating the mean, variance, and correlation of a Gaussian field, we can represent surfaces, media, and a new range of in-between effects including non-exponential media with arbitrary heterogeneity and microfacets with correlated microstructures.

Our key contribution is the introduction of new algorithms for computing light transport in correlated disorder, addressing a long-standing challenge in computer graphics and related fields. Stochastic environments, characterized by non-uniform distributions of geometry or particles, are found in numerous settings like tissue [Tuchin 2000], leaves [Fukshansky et al. 1993], oceans [Kirk 1975], reactors [Williams 1974], and molecular clouds [Boissé 1990]. Hence, our algorithms, primarily designed for computer graphics applications, also hold potential for fields like reactor engineering and remote sensing. Unlike traditional media, our stochastic surfaces create oriented intersections, permitting the use of conventional BS-DFs for medium interactions, and can also delineate discrete phases of classical homogeneous material with stochastic boundaries.

Our rendering algorithms represent a unified and principled approach to simulating light transport in such complex scenes, but this currently comes at a cost—both in terms of computational resources, where they are more demanding than traditional graphics scene representations, and in terms of the availability of closed-form results. For example, we currently only support next-event estimation in a limited set of circumstances, preventing us from efficiently rendering scenes with sparse light sources when using mirror micro-BSDFs. Still, we believe that the method presented here will not only pave the way for novel applications in graphics but also extend its potential to scientific and industrial domains and we provide a reference implementation at https://github.com/daseyb/gpis-light-transport.

2 RELATED WORK

2.1 Stochastic Implicit Surfaces in Graphics and Beyond

Stochastic implicit surfaces have been used previously in graphics for terrain rendering [Gamito 2009] and procedural generation using e.g. Perlin noise [Ebert et al. 2003]. Here, a single realization of the surface is rendered, leading to purely solid surfaces. We consider the average light transport through all realizations, allowing us to express volumetric effects as well. Using a Gaussian process as a stochastic implicit surface has previously been used for surface reconstruction [Martens et al. 2017; Williams and Fitzgibbon 2006], where the GP is fitted to a set of point observations and the surface extracted from zero-crossings of the GP mean. (Co)variance in the GP has also been used to express reconstruction uncertainty to inform e.g. robot grasping [Dragiev et al. 2011]. Additionally, there has been a renewed interest in using GPs to quantify the uncertainty inherent in reconstructing surfaces from point clouds via Poisson surface reconstruction [Sellán and Jacobson 2022, 2023]. The posterior variance can then be used to e.g. optimize additional capture locations for maximum uncertainty reduction. Our method can use the retrieved posterior mean and covariance functions and faithfully incorporate the covariance during rendering, visualizing uncertainty as "fuzziness" in the surface.

2.2 Uncertainty Visualization

How to visualize uncertainty in volumetric datasets as "fuzziness" [Fout and Ma 2012] is a problem well explored in scientific data visualization and Brodlie et al. [2012] provide a comprehensive review. Closest to our method, Pfaffelmoser et al. [2011] compute "isosurface-first-crossing probabilities" (IFCPs) in Gaussian fields, but for efficiency reasons, they limit themselves to Markov processes. We will show (Section 5.3) that while Markov processes make IFCPs tractable, they produce "non-smooth" processes which do not allow for scattering and hence cannot be used as geometry models when the aim is to compute global illumination.

2.3 Volumetric Scene Representations

Volumetric scene representations have received increased interest in recent years because they are easily differentiable and thus amenable to inverse rendering. Methods based on Neural Radiance Fields (NeRF) [Barron et al. 2021; Mildenhall et al. 2020; Müller et al. 2022] encode a classical (emissive) volume using a neural network and train it from images. More recent works [Fridovich-Keil et al. 2022] drop the neural component and represent the volume entirely with an octree. These methods have proven very effective in novel view rendering, but struggle at scene reconstruction because hard surfaces must be approximated with volume densities. Methods such as NeuS [Wang et al. 2021] or NeUDF [Liu et al. 2023] instead choose implicit surfaces as the underlying representation, where surfaces are defined by zero-crossings of a coordinate network. The hard surface prior leads to higher quality surface reconstructions for solid objects, but these methods fail to encode both surface priors and uncertainty in their representation, leading to poor results on scenes that contain both poorly- and well-resolved detail.

2.4 Transport in Discrete Stochastic Media

Similar to our method, techniques that compute transport through collections of packed objects [Moon et al. 2007] or granular media [Meng et al. 2015; Müller et al. 2016] consider the *ensemble average light transport* over many possible configurations of objects or grains. Unlike us, they pre-compute *aggregate* transport — averaging not only over different geometry configurations but also all possible light paths connecting two points. Reusing the same aggregate transport tables in different regions of their medium makes their pre-computation tractable, but introduces approximations and hence they often only opt to use it for higher-order scattering where error is less visible. We do not rely on tabulated distributions and instead focus on principled derivations of ensemble average light transport in a specific class of stochastic geometry. Discrete stochastic media have a rich history outside of graphics [Estrade et al. 2012; Lu and Torquato 1992; Torquato and Lu 1993], and Gaussian processes have previously been used to represent microstructures formed by binary mixtures [Jiao et al. 2007, 2008; Torquato 1986, 2005]. To build practical transport algorithms in such media, strong assumptions about the chord lengths along a transect are often made, such as a renewal [Roberts and Torquato 1999] or Markov [Pomraning 1991] assumption, which we avoid making.

2.5 Non-exponential Radiative Transport

Similar to methods for discrete stochastic media, these works assume that the photons experience free-path lengths given by a *renewal process*, where correlations "reset" after each bounce, and collisions on different segments of a path are assumed independent, which can lead to significant error [d'Eon 2023]. While algorithms like the chord-length sampling [Zimmerman and Adams 1991], conditional point sampling [Vu and Olson 2021], and 1D-point-process sampling [d'Eon 2023] have improved results by providing additional memory along a particle's history, they are restricted to piecewise-stationary random systems and often make strong assumptions about the stochastic geometry. Our work transcends these limitations by leveraging the properties of Gaussian processes in the context of implicit surfaces. We show that our formalism can exhibit non-exponential attenuation as well but also allows for longer-reaching correlations across multiple segments or the entire path, achieving higher accuracy (but requiring higher cost). Additionally, work on non-exponential transport often only considers how correlation affects free-flights, while phase functions are assumed deterministic. In our method, appearance and transmittance are tightly coupled.

3 BACKGROUND & NOTATION

In the following, we will present the definitions and introduce any notation required for the derivations in Section 4. Readers familiar with the theory are welcome to skip individual subsections.

3.1 Gaussian Processes & Fields

A Gaussian process $\text{GP}(\mu, k)_\Omega$ is a distribution over functions $f : \Omega \to \mathbb{R}$ such that for any finite set of locations $\mathbf{x}_1, \ldots, \mathbf{x}_n =: X \subseteq \Omega$, the evaluations of the function follow an n-dimensional Gaussian distribution $f_X \sim \mathcal{N}(\mu(X), k(X, X))$. Here, we abuse notation

slightly with $\mu(X) = [\mu(\mathbf{x}_1), \ldots, \mu(\mathbf{x}_n)]^\top$ being an n-dimensional *mean* vector and $k(X, X)$ an $n \times n$ *covariance matrix*, with entries $k(X, X)^{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)$. These are obtained by evaluating the mean function $\mu(\mathbf{x}) : \Omega \to \mathbb{R}$ and covariance kernel $k(\mathbf{x}, \mathbf{y}) : \Omega \times \Omega \to \mathbb{R}$ for each location or pair of locations in X, respectively. When the domain Ω is a subset of $\mathbb{R}^d$, we call the resulting Gaussian process a *Gaussian field* [Adler and Taylor 2009]. In this work, we will only investigate 3D Gaussian fields, but for other applications, for example, in machine learning, d can be very large. For a more thorough introduction to Gaussian processes we recommend Williams and Rasmussen [2006], whose notation we follow in this section.

3.1.1 Restrictions to sub-domains. We can restrict the input domain of a Gaussian process (e.g. from 3D space $\mathbb{R}^3$ to the points $\mathbf{x}_\mathbb{R} = \{\mathbf{x} + t\boldsymbol{\omega} | t \in \mathbb{R}\}$ along a line) with no change in the statistics. That is, if $f \sim \text{GP}(\mu, k)_{\mathbb{R}^3}$ and $g \sim \text{GP}(\mu, k)_{\mathbf{x}_\mathbb{R}}$, then $f(\mathbf{x}_\mathbb{R}) \overset{d}{=} g(\mathbf{x}_\mathbb{R})$, i.e. they are equal in distribution. This is a useful property because we will often look at lower dimensional "slices" of a Gaussian field, and we can do so without having to change the mean or covariance kernel.

3.1.2 Kernel Functions. Covariance kernels can take many forms and are the primary "design variable" for Gaussian processes. In Fig. 2 (left) we show realizations produced by a set of common kernels. Intuitively, the kernel prescribes how "similar" values are at nearby points in space, and small changes to the kernel can radically change the properties, such as differentiability, of the resulting sample functions. The only restriction on kernels is that they must be *positive semi-definite*. Kernels composed via multiplication and addition are again valid kernels, which allows the design of complex kernels out of simple building blocks.

We call a kernel *stationary* if $k(\mathbf{x}, \mathbf{y}) = k(\mathbf{y} - \mathbf{x})$, i.e. the kernel does not change shape depending on where we evaluate it. A further subset are *isotropic* kernels, $k(\mathbf{y} - \mathbf{x}) = k(\|\mathbf{y} - \mathbf{x}\|)$, which only depend on the *distance* between points. More general, non-stationary kernels for which $k(\mathbf{x}, \mathbf{y}) \neq k(\mathbf{y} - \mathbf{x})$ are not widely discussed in Gaussian process literature since their theoretical analysis is difficult. That said, the statistics of real-world datasets are very rarely truly stationary [Noack and Sethian 2022], and we will be using a
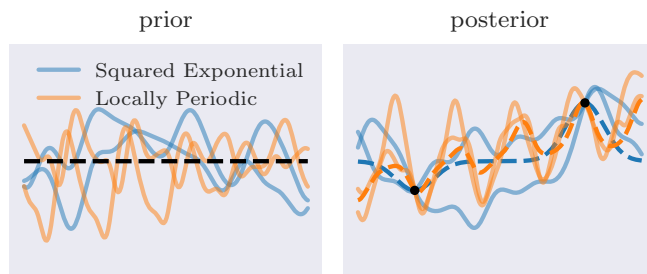

Fig. 2. Gaussian processes provide a very flexible way to define sets of functions (each being a *realizations* of the process), which can have wildly different properties based on the *covariance kernel* of the process. On the left we draw realizations from two *prior* (i.e. unconditioned) processes with different covariance kernels. On the right, we show realizations from the corresponding *posterior* process, conditioned on the two black points (dashed lines are the posterior means of the same hue kernels).

non-stationary kernel to allow for spatially varying geometry and material properties such as surface roughness.

3.1.3 Conditioned Processes.
Gaussian processes can be *conditioned* on observations, and the conditioned *posterior* process is again Gaussian. For example, given measurements m at locations C, we can derive the *conditioned* $\text{GP}(\mu(x), k(x, y) \mid \zeta_m)$ with the condition $\zeta_m := f(C) = m$ as

$$f \sim \text{GP}(\mu_{|\zeta_m}(x), k_{|\zeta_m}(x, y)), \quad \text{with} \tag{1}$$

$$\mu_{|\zeta_m}(x) = \mu(x) + k(x, C)k(C, C)^{-1}(m - \mu(C)) \tag{2}$$

$$k_{|\zeta_m}(x, y) = k(x, x) - k(x, C)k(C, C)^{-1}k(C, x) \tag{3}$$

Sampling from the conditioned GP then produces realizations that pass through the values m, but interpolate according to the chosen covariance kernel in-between as shown in Fig. 2 (right). Unfortunately, computing the posterior (i.e. conditioned) mean and covariance involves inverting the $n \times n$ matrix $k(C, C)$, where $n = |C|$. This has $O(n^3)$ time complexity and becomes intractable for large n.

3.1.4 Sampling from Gaussian Processes.
Efficiently generating large numbers of correlated samples from a Gaussian process also poses a challenge. Samples can be drawn directly from $\text{GP}(\mu, k)$ via

$$f(X) = \mu(X) + k(X, X)^{1/2}\eta, \tag{4}$$

where $\eta \sim \mathcal{N}(0, 1)$ and $A^{1/2}$ is the matrix square root s.t. $A = A^{1/2}A^{1/2}$. In either case, the time complexity is $O(p^3)$ with $p = |X|$, and drawing correlated samples from a Gaussian process quickly becomes intractable for large p. This method of directly drawing values of a sample function is called the *function-space view* of Gaussian processes [Williams and Rasmussen 2006]. There are other methods of simulating Gaussian processes, such as the *weight-space view* [Rahimi and Recht 2007], but these are usually limited in the types of processes they support (e.g. only ones with *stationary* kernels). Still, the weight-space view will be valuable when validating our method, even though we cannot use it in more complex applications.

3.1.5 Derivatives of Gaussian Processes and "Multi-Task" GPs.
Due to the linearity of the derivative operator, the derivative of a Gaussian process is again a Gaussian process and takes the form

$$\text{GP}'(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{y})) = \text{GP}(\mu'(\mathbf{x}), k_{\mathbf{x}, \mathbf{y}}(\mathbf{x}, \mathbf{y})) \tag{5}$$

where $k_{\mathbf{x}, \mathbf{y}}(\mathbf{x}, \mathbf{y}) = \frac{\partial^2 k(\mathbf{x}, \mathbf{y})}{\partial \mathbf{x} \partial \mathbf{y}}$. Additionally, we can reason about the joint value-derivative distribution

$$\begin{bmatrix} f(X) \\ f'(Y) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mu(X) \\ \mu'(Y) \end{bmatrix}, \begin{bmatrix} k(X, Y) & k_{\mathbf{y}}(X, Y) \\ k_{\mathbf{x}}(Y, X) & k_{\mathbf{x}, \mathbf{y}}(Y, Y) \end{bmatrix}\right). \tag{6}$$

This will come in useful when we reason about distributions of normals on Gaussian process implicit surfaces in Section 4, since the normals of a GPIS are aligned with the spatial derivative of the underlying GP.

3.2 Implicit Surfaces

An implicit surface is defined as the level-set of a scalar function $f : \Omega \rightarrow \mathbb{R}$, $\{\mathbf{x} \in \Omega \mid f(\mathbf{x}) = l\}$. Without loss of generality, we consider the *zero level-set*, where $l = 0$. In addition to the surface, we can partition Ω into *inside* ($\{f(\mathbf{x}) < 0\}$) and *outside* ($\{f(\mathbf{x}) > 0\}$) regions. Computing the intersection of a ray with an implicit surface can be
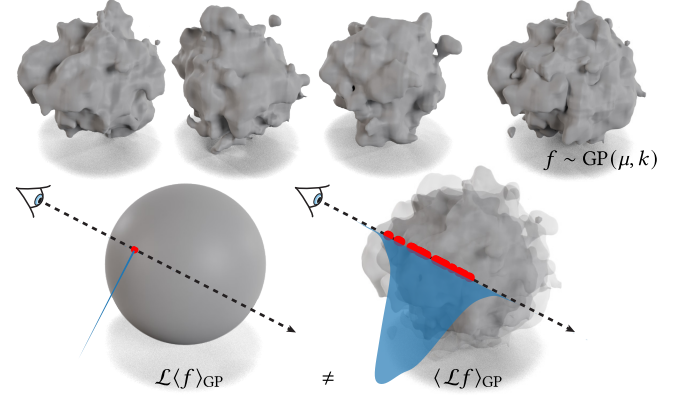
Fig. 3. Each realization of a 3D Gaussian process forms a corresponding implicit surface (top). The effect of geometry is highly non-linear, and light transport ($\mathcal{L}$) on the *mean implicit surface* (bottom left, prior work) is drastically different to the *mean light transport* over all realizations of implicit surfaces (bottom right, ours). For the mean surface, free-flights are deterministic (peak, bottom left), while they form a distribution in the ground truth transport (distribution, bottom right).

done via *root-finding*, i.e. finding distance s along ray $\mathbf{x} + s\omega$ such that $f(\mathbf{x}_s) = 0$, where $\mathbf{x}_t = \mathbf{x} + t\omega$, i.e. a point which is a distance t along a ray $(\mathbf{x}, \omega)$. Many root-finding methods are available, depending on which properties of f are known [Hart 1996; Stol and De Figueiredo 1997]. The intersection point $\mathbf{x}_s$ with $s = \arg\min_{t \in \mathbb{R}^+} f(\mathbf{x}_t) = 0$ always depends on a particular ray and a particular implicit surface f. When we want to make this clear, we indicate the dependence on the surface as $\mathbf{x}_s^f$. In addition to intersections, we are often interested in the *normal* $\mathbf{n}_s$ of the surface at a point $\mathbf{x}_s$. Implicit surfaces make it easy to calculate the normal vector as $\mathbf{n}_s = \bar{\nabla} f(\mathbf{x}_s)$, where $\bar{\nabla}$ computes the normalized gradient $\nabla f(\mathbf{x}_s)/\|\nabla f(\mathbf{x}_s)\|$. Just as before, we sometimes will highlight the dependence of the normal at the intersection on a particular implicit surface as $\mathbf{n}_s^f$.

3.2.1 Stochastic Implicit Surfaces.
A stochastic implicit surface (SIS) is the distribution of level sets of a stochastic process. In this paper, our main focus is on *Gaussian process implicit surfaces* (GPIS), where the stochastic process is Gaussian. Each realization $f \sim \text{GP}$ of the process generates an implicit surface at the zero level set $f(\mathbf{x}) = 0$.
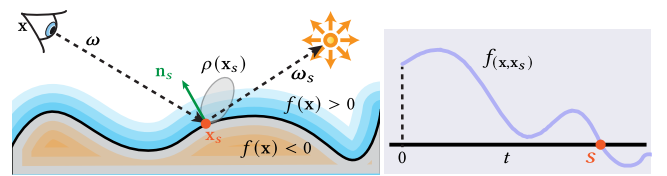

Fig. 4. We can solve the recursive integral formulation of light transport in Eq. (7) in an implicit surface f by intersecting rays with the surface, computing the normal at the intersection point, and picking a scattering direction according to the BSDF (left). Finding the first intersection point requires us to find the first zero crossing of the implicit surface evaluated along the ray (right).
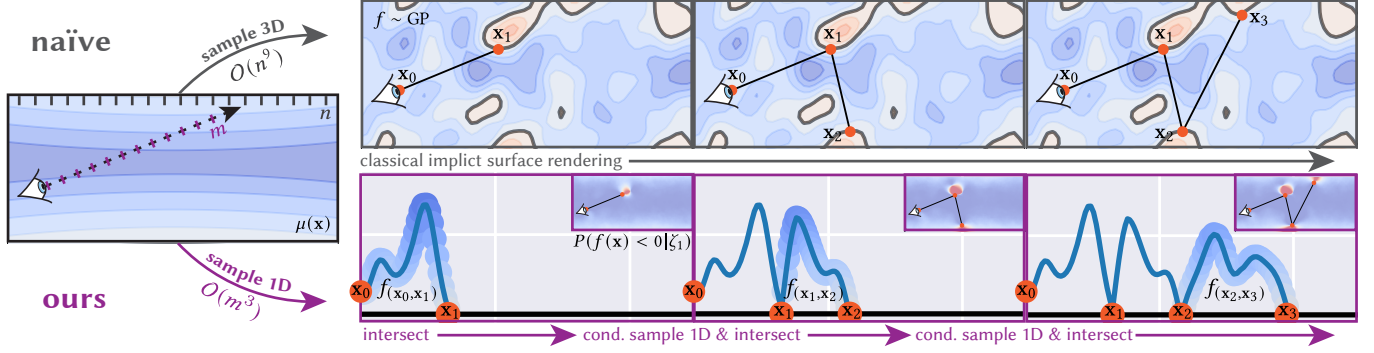
Fig. 5. We show the difference between the "naïve" method in Eq. (8) (top), which requires global realization sampling, and the progressive sampling method we present in Eq. (14), which only samples realizations along path segments as needed (bottom). Both compute the ground truth ensemble average light transport in a given GPIS (μ shown on the left). In the naive method, tracing is simple since each path sees a deterministic f (red: inside, blue: outside, black: zero-level set), but sampling large numbers of correlated samples from a general Gaussian process is prohibitively expensive even for modest resolutions (scaling with $O(n^9)$ where n is the discretization resolution). In our method, we only ever sample 1D slices of the process, which scales with $O(m^3)$, where m is the number of steps taken along each ray. While still expensive compared to classical rendering techniques, it makes it feasible, for the first time, to compute GPIS light transport in a principled manner.

In Fig. 3, we show how the deterministic intersection distance and normal in implicit surface rendering turn into distributions when considering SISes.

3.3 Light Transport

We write the surface rendering equation [Kajiya 1986] for light transport in a scene defined by the implicit surface f as

$$L^f(\mathbf{x}, \boldsymbol{\omega}) = \int_{S^2} \rho(\mathbf{x}_s^f) L^f(\mathbf{x}_s^f, \boldsymbol{\omega}_s^f) \, d\boldsymbol{\omega}_s \qquad (7)$$

where $\mathbf{x}_s^f$ and $\mathbf{n}_s^f$ are defined as in Section 3.2 and, for brevity, we use the shorthand $\rho(\mathbf{x}_s) \coloneqq \rho(\mathbf{x}_s, -\boldsymbol{\omega}, \boldsymbol{\omega}_s, \mathbf{n}_s)|\mathbf{n}_s \cdot \boldsymbol{\omega}_s|$ for the cosine-weighted BSDF of the implicit surface, which we assume to be a deterministic function of position $\mathbf{x}_s$, normal $\mathbf{n}_s$ and incoming and outgoing directions $\boldsymbol{\omega}_s, -\boldsymbol{\omega}$ respectively. We illustrate this form of the rendering equation in Fig. 4.

4 ENSEMBLE-AVERAGED LIGHT TRANSPORT IN GPISES

The radiance L^f introduced in the previous section describes light transport for a fixed implicit surface f (Eq. (7)). We now consider the case when f is a realization of a Gaussian process (a GPIS), and will derive methods for computing the average transport over all realizations. Formally, we investigate the *ensemble averaged* light transport over all realizations f of the Gaussian process $\mathrm{GP}(\mu, k \mid \zeta)$,

$$\langle L^f(\mathbf{x}, \boldsymbol{\omega}) \rangle_\zeta = \int_{\mathrm{GP}} L^f(\mathbf{x}, \boldsymbol{\omega}) \, d\gamma_{\mu,k}(f \mid \zeta), \qquad (8)$$

where $\gamma_{\mu,k}(f \mid \zeta)$ is the classical Wiener measure [Taylor 2006, Ch. 16] of f with respect to the conditioned $\mathrm{GP}(\mu, k \mid \zeta)$, i.e. the probability density of sampling $f \sim \mathrm{GP}(\mu, k \mid \zeta)$. We will use the condition ζ as an "editing" tool in Section 6, since it allows us to only consider realizations that, e.g., pass through a certain point in space. For brevity, the dependency on μ, k is made implicit going

forward. We can form a straightforward Monte Carlo of Eq. (8) via

$$\langle L^f(\mathbf{x}, \boldsymbol{\omega}) \rangle_\zeta \approx \frac{1}{N} \sum_{j=1}^{N} L^{f_j}(\mathbf{x}, \boldsymbol{\omega}), \quad f_j \sim \mathrm{GP}(\mu, k \mid \zeta), \qquad (9)$$

where we simply average the light transport across N independent realizations of the Gaussian process. Within each realization, L^{f_j} may in turn be estimated with Monte Carlo using e.g. path tracing.

While Eq. (9) produces correct results on average, it is wholly impractical: For each sample, an entire 3D realization f_j must be constructed. If the GPIS is discretized into a volume of sidelength $O(n)$, constructing each realization takes $O(n^9)$ time, making it infeasible. In the following we will derive an alternative form of Eq. (8) that interweaves the Monte Carlo rendering process with the random sampling of realizations. As a result, we will be able to construct realizations in a "just in time" fashion, only sampling the parts of f that are required to continue the light transport path.

Explicit delta free-flights in fixed f. As a first step we make it explicit that free-flight distributions are *delta functions* in fixed implicit surfaces, rewriting Eq. (7) as

$$L^f(\mathbf{x}, \boldsymbol{\omega}) = \int_0^\infty \iint_{S^2} \rho(\mathbf{x}_t) \delta^f(\mathbf{x}_t, \mathbf{n}) I^f(0, t) \, L^f(\mathbf{x}_t, \boldsymbol{\omega}_t) \, d\boldsymbol{\omega}_t \, d\mathbf{n} \, dt,$$
$$(10)$$

where $\delta^f(\mathbf{x}_t, \mathbf{n}) = \delta(f(\mathbf{x}_t) - 0) \cdot \delta(\bar{\nabla} f(\mathbf{x}_t) - \mathbf{n})$ and

$$I^f(0, t) = \begin{cases} 1 & \forall s \in (0, t) : f(\mathbf{x}_s) > 0 \\ 0 & \text{otherwise} \end{cases}. \qquad (11)$$

Intuitively, $\delta^f(\mathbf{x}_t, \mathbf{n}) I^f(0, t)$ is non-zero when $\mathbf{x}_t$ is the intersection point of the ray in realization f and $\mathbf{n}$ is the normal at the intersection point. Expanding L^f as defined in Eq. (10) and then pulling the averaging over GP realizations into the integrals over intersection
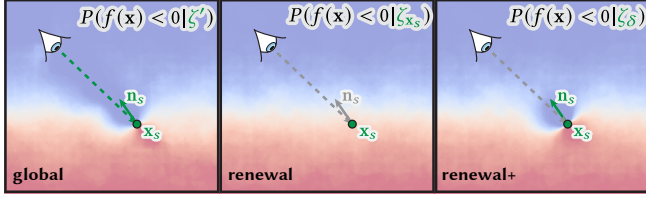
Fig. 6. We show the conditional probability of being inside a surface (blue: low, red: high) for every point in space, after a collision at $\mathbf{x}_s$. The ground truth (left, our Global model) requires *segment* correlations. This makes sampling free-flights after a collision expensive but preserves information along the whole path. The Renewal model (middle) only remembers the value of the process at path vertices. This is reminiscent of the distinction between "free-space" and "particle" path vertices in non-exponential transport [Bitterli et al. 2018] and discards any knowledge about the incoming segment. The Renewal+ model (right) additionally remembers the *gradient* at path vertices, producing results much closer to the ground truth with very little additional cost over the Renewal model.

distance, normal, and exit direction allows us to rewrite Eq. (8) as

$$
\langle L^f(\mathbf{x}, \boldsymbol{\omega}) \rangle_\zeta =
$$
$$
\int_0^\infty \iint_{\mathcal{S}^2} \rho(\mathbf{x}_t) \left\langle \delta^f(\mathbf{x}_t, \mathbf{n}) I^f(0, t)\, L^f(\mathbf{x}_t, \boldsymbol{\omega}_t) \right\rangle_\zeta d\boldsymbol{\omega}_t\, d\mathbf{n}\, dt. \quad (12)
$$

So far, we have gained little — the ensemble average on the right-hand side is still over full 3D realizations f.

From delta functions to conditioned ensembles. Using a relation we derive in Section 1.1.2 of the supplemental, we can pull the delta function out of the ensemble average — i.e. instead of "filtering" realizations *after* sampling, we only average over realizations having the required normal and zero-crossing in the first place, giving us

$$
\langle L^f(\mathbf{x}, \boldsymbol{\omega}) \rangle_\zeta =
$$
$$
\int_0^\infty \iint_{\mathcal{S}^2} \rho(\mathbf{x}_t) \gamma_{\mathbf{x}_t}(0, \mathbf{n} \mid \zeta) \left\langle I^f(0, t)\, L^f(\mathbf{x}_t, \boldsymbol{\omega}_t) \right\rangle_{\zeta \wedge \zeta_\delta} d\boldsymbol{\omega}_t\, d\mathbf{n}\, dt,
$$
$$
(13)
$$

where the condition $\zeta_\delta := (f(\mathbf{x}_t) = 0 \wedge \overline{\nabla} f(\mathbf{x}_t) = \mathbf{n})$ restricts the realizations inside the ensemble average to the ones where the delta function is non-zero and $\gamma_{\mathbf{x}_t}(0, \mathbf{n} \mid \zeta)$ accounts for the density of sampling such realizations. This already constricts the space of realizations we need to consider in the ensemble average on the right-hand side as we move along the path.

Splitting realizations and progressive sampling. Finally, we note that $I^f(0, t)$ only depends on the *1D restriction* of f to the ray segment $(\mathbf{x}, \mathbf{x}_t)$. Gaussian processes, being closed under conditioning, allow us to decompose sampling f into two parts: Sampling the values $f_{\mathbf{x}, \mathbf{x}_t}$ along the ray segment and then sampling f over the remainder of the domain. As long as we condition the second step on $\zeta_{(\mathbf{x}, \mathbf{x}_t)}$, i.e. the results of the first one, the final distribution is the same as if we had sampled all parts of the domain at the same time. Again, using a relation shown in Section 1.1.3 of the supplemental we split

Fig. 7. We show a scene rendered with our range of memory models. Empirically, we found that the Renewal+ model (right) preserves ground truth appearance (left) much better than the Renewal model (middle).

the ensemble average on the right-hand side in two, giving us

$$
\langle L^f(\mathbf{x}, \boldsymbol{\omega}) \rangle_\zeta = \int_0^\infty \iint_{\mathcal{S}^2} \rho(\mathbf{x}_t) \gamma_{\mathbf{x}_t}(0, \mathbf{n} \mid \zeta)
$$
$$
\left\langle I^f(0, t) \left\langle L^f(\mathbf{x}_t, \boldsymbol{\omega}_t) \right\rangle_{\zeta \wedge \zeta_\delta \wedge \zeta_{(\mathbf{x}, \mathbf{x}_t)}} \right\rangle_{\zeta \wedge \zeta_\delta}^{(\mathbf{x}, \mathbf{x}_t)} d\boldsymbol{\omega}_t\, d\mathbf{n}\, dt, \quad (14)
$$

where $\langle \cdot \rangle_\zeta^{(\mathbf{x}, \mathbf{x}_t)}$ represents the conditioned average over realization restricted to a path segment. Eq. (14) is vastly more efficient to estimate, as at each level of recursion, realizations only need to be evaluated on path segments $(\mathbf{x}, \mathbf{x}_t)$ and not the entire scene, reducing the cost to $O(m^3)$, where m is now the number of steps taken along a ray segment (see Fig. 5). Ensemble averaging is done for each path segment individually; however, Eq. (14) is exact, and no approximations are made. This is made possible because each subsequent segment on the path is conditioned on the observed values of the realization seen on all prior segments, ensuring a globally consistent view. However, the cost of sampling realizations on later segments increases as more and more conditions are added.

4.1 Memory Models

While sampling from a 1D process (14) is wildly more practical than sampling global 3D realizations (8), conditioning subsequent path segments on *all* values sampled along prior segments is still computationally intensive. To enable a more practical implementation, we introduce the concept of *memory models*, which "forget" some of the conditioning on prior segments in Eq. (14). This trades off computational cost for approximation error. Effectively, we replace the list of conditions $\zeta' := \zeta \wedge \zeta_\delta \wedge \zeta_{(\mathbf{x}, \mathbf{x}_t)}$ on the recursive ensemble average in Eq. (14) by a simpler form. Figure 6 shows a didactic visualization of the different models we present in this section. Many more are possible, but we restrict ourselves to the most interesting subset.

4.1.1 The Global-n memory model. The Global-n model conditions each segment on only the n prior segments $\zeta'_{G^n} := \zeta^{k-n} \wedge \ldots \wedge \zeta^k$, where each ζ^k is $\zeta_\delta \wedge \zeta_{(\mathbf{x}, \mathbf{x}_t)}$ of Eq. (14) for the k-th path segment. In the limit, for $n = \infty$, the ground truth model is recovered. While still expensive, this model allows setting an upper bound on the evaluation cost of Eq. (14) that does not grow arbitrarily with the length of the light path.

4.1.2 The Renewal memory model. On the other end of the spectrum, the Renewal model retains only minimal state, i.e. that a surface must be present at $\mathbf{x}_k$ on adjacent segments, meaning $\zeta'_R(f) :=$

$(f(\mathbf{x}_k) = 0)$. This is analogous to models of non-exponential media [Bitterli et al. 2018; d'Eon 2018; Jarabo et al. 2018], which remember correlations only up to the preceding path vertex. This semi-Markov model "jumps" to a different realization at each path vertex.

4.1.3 The Renewal+ memory model. D'Eon [2023] showed that a renewal approximation in volumes with stochastic density can introduce significant errors. Our Renewal+ model aims to reduce this error by additionally remembering the previous normal $\mathbf{n}_k$ at $\mathbf{x}_k$, i.e. we condition on $\zeta'_{R+}(f) := \zeta'_R(f) \wedge \left(\overline{\nabla} f(\mathbf{x}_k) = \mathbf{n}_k \right)$. This enforces *gradient consistency* across segment realizations, which reduces approximation error compared to the Renewal model significantly (see Fig. 7) while being much more tractable than the Global model. Conditioning on higher order derivatives may be interesting to explore in future work.

Note that both the Renewal and Renewal+ memory models only retain information about the last collision. We leave the question of whether keeping additional vertices, similar to the Global-n model, would improve results for future work, since these and other memory models are trivial to experiment with.

4.2 Progressive sampling via function-space GPs

To round out this section, we present a method for rendering GPIS light transport that we will use to produce subsequent images and other results. This method supports fully non-stationary covariance kernels (see Section 6) and any choice of memory model. For an alternative weight-space implementation that supports only stationary kernels but can serve as a "ground truth", see Section 3 of the supplemental.

We can write a one-sample Monte Carlo estimator of Eq. (14) as

$$\langle \widehat{L_i(\mathbf{x}^u, \boldsymbol{\omega})} \rangle_\zeta = \frac{\rho(\mathbf{x}_t)\Gamma(t, \mathbf{n} \mid \zeta)}{p(t, \mathbf{n}, \boldsymbol{\omega}_t, f_{(\mathbf{x}, \mathbf{x}_t)})} \langle \widehat{L_i(\mathbf{x}_t, \boldsymbol{\omega}_t)} \rangle_{\zeta'}, \quad (15)$$

$$t, \mathbf{n}, \boldsymbol{\omega}_t, f_{(\mathbf{x}, \mathbf{x}_t)} \sim p(t, \mathbf{n}, \boldsymbol{\omega}_t, f_{(\mathbf{x}, \mathbf{x}_t)}) \quad (16)$$

ALGORITHM 1: Simple light transport in a GPIS

function $L(\vec{r} \mid \zeta)$: color
  $(t, \mathbf{n}, \zeta') \leftarrow \text{nextHit}(\vec{r} \mid \zeta)$
  return $L_e(\vec{r}, t, \mathbf{n} \mid \zeta') + L_r(\vec{r}, t, \mathbf{n} \mid \zeta')$

function $L_r(\vec{r}, t, \mathbf{n} \mid \zeta)$: color
  $(\boldsymbol{\omega}', p_{\boldsymbol{\omega}'}) \leftarrow \text{sampleDirection}(\vec{r}, t \mid \zeta)$
  $\rho \leftarrow \text{evalBSDF}(\vec{r}(t), \vec{r}.\boldsymbol{\omega}, \boldsymbol{\omega}', \zeta)$
  return $\rho \cdot (\mathbf{n}(\vec{r}, t, \zeta).\boldsymbol{\omega}')/p_{\boldsymbol{\omega}'} \cdot L(\text{Ray}(\vec{r}(t), \boldsymbol{\omega}'), \zeta)$

function $\text{nextHit}(\vec{r} \mid \zeta)$: $(t, \mathbf{n}, \zeta')$
  $f \leftarrow \text{sampleRealization}(\vec{r} \mid \zeta)$
           //Sample a 1D realization along the ray.
  $t \leftarrow \arg\min_s f(s) \leq 0$ //Find first zero crossing.
  $\mathbf{n} \leftarrow \text{sampleNormal}(\vec{r}, t \mid \zeta \wedge f)$
           //Sample normal at intersection point.
  $\zeta' \leftarrow M(\zeta \wedge f \wedge \mathbf{n})$ //Filter state based on memory model.
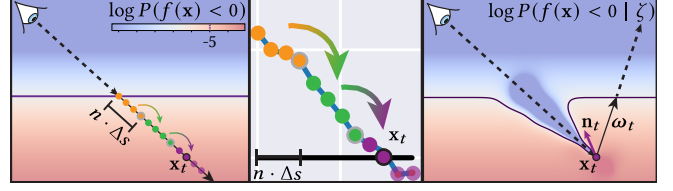 **return** $(t, \mathbf{n}, \zeta')$

Fig. 8. We first determine a conservative envelope around the GPIS based on the point-wise probability of being inside the surface. We then place n (here $n = 4$) sample locations along the ray segment at fixed distances Δs based on the covariance kernel. If the segment is longer than $n \cdot \Delta s$, we stop short, effectively inserting a "null-collision" (grey outline), and continue using a new set of conditioned samples. Once an interpolation of the discrete samples crosses 0 we have found the intersection point $\mathbf{x}_t$ (black outline). We then sample a normal $\mathbf{n}_t$ and a direction from the micro BSDF and continue in the scattered direction $\boldsymbol{\omega}_t$.

where, $p(t, \mathbf{n}, \boldsymbol{\omega}_t, f_{(\mathbf{x}, \mathbf{x}_t)})$ is a joint probability density of sampling intersection distance t, normal at the intersection point $\mathbf{n}$, outgoing direction $\boldsymbol{\omega}_t$ and a 1D realization $f_{(\mathbf{x}, \mathbf{x}_t)}$ that has positive values along the line segment $(\mathbf{x}, \mathbf{x}_t)$ and a zero crossing at t. The question now is whether we can choose $p \sim \rho(\mathbf{x}_t)\Gamma(t, \mathbf{n} \mid \zeta)$. This would both reduce variance, due to importance sampling part of the integrand, and also save us from having to compute $\Gamma(t, \mathbf{n} \mid \zeta)$, for which we don't have an analytic expression. To achieve this, we carefully consider the order of sampling the quantities we need as follows

(1) Sample a 1D realization along the ray $f_{(\mathbf{x}, \mathbf{x}_\infty)} \sim \text{GP}_{(\mathbf{x}, \mathbf{x}_\infty)|\zeta}$.
(2) Find the intersection distance $t = \arg\min_{s \in (0, \infty)} f_{(\mathbf{x}, \mathbf{x}_s)} = 0$. (now $f_{(\mathbf{x}, \mathbf{x}_t)} \sim \text{GP}^+_{(\mathbf{x}, \mathbf{x}_\infty)|\zeta}$ as required).
(3) Sample a normal $\mathbf{n} \sim \text{GP}^{\overline{\nabla}}_{\mathbf{x}_t|\zeta \wedge f_{(\mathbf{x}, \mathbf{x}_t)}}$
(4) Sample the BSDF $\boldsymbol{\omega}_t \sim \rho(\mathbf{x}_t)$

This perfectly importance samples the factor in front of the recursive term in Eq. (15). Afterward, we can recursively compute $\langle \widehat{L_i(\mathbf{x}_t, \boldsymbol{\omega}_t)} \rangle_{\zeta'}$ in exactly the same way, and we continue until we hit a light source or exit the scene. These steps are reflected in Alg. 1.

4.2.1 Practical considerations. Still, the list of steps does not tell the whole story. The main issue lies in sampling $f_{(\mathbf{x}, \mathbf{x}_\infty)}$. Recall that this is a *function*, i.e. an uncountably infinite set of values. Using function space sampling, we can only sample a *finite* set of m values, and the cost grows with $O(m^3)$. Hence, we have to approximate $f_{(\mathbf{x}, \mathbf{x}_\infty)}$ somehow using a finite (and hopefully small) number of samples. Luckily, the processes we consider are smooth, after all we require continuity and second-order differentiability. This means that sample functions should be well approximated by interpolation of a relatively small number of samples. The second issue is that $f_{(\mathbf{x}, \mathbf{x}_\infty)}$ is defined over the positive real number line, i.e. it contains all points along the ray to infinity. Hence even when discretizing the number line, we would still need to sample an infinite number of values. We solve this with a two-layer strategy: First, we derive appropriate scene bounds. Then, we split the ray into segments with a fixed number of sample points, spaced appropriately for the given covariance function, and progressively sample the realization for each segment, much like we progressively sample realizations between path vertices. This process is visualized in Fig. 8 and we
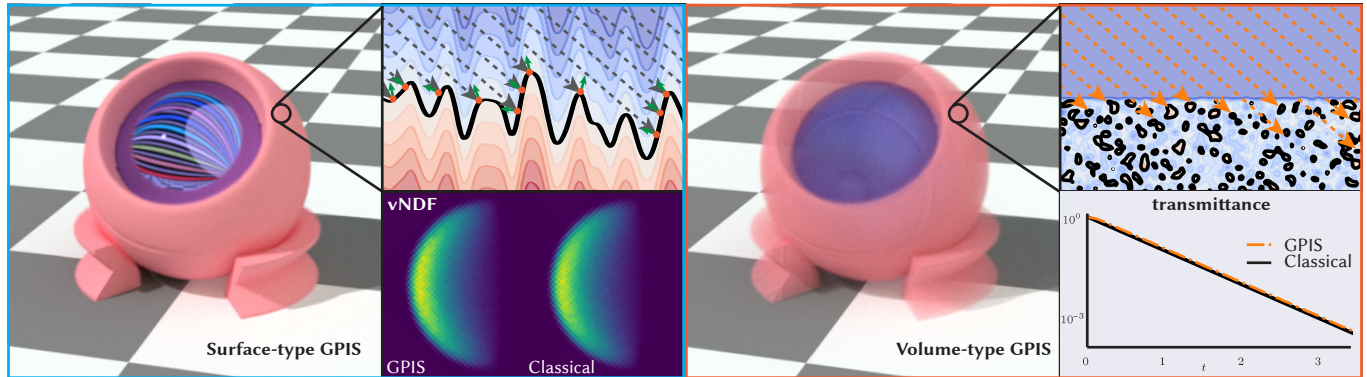
Fig. 9. Here, we explore two extremes of the appearance spectrum that our method supports: microfacet surfaces (left) and participating media (right). We can match classical results by carefully controlling the statistics of the realizations we generate (insets top), and by deriving the central quantities for each in our model (insets bottom, vNDF for microfacets and transmittance for participating media). Our method easily incorporates multiple scattering and allows freely choosing the BSDF of the microgeometry. The micro-bsdf of the sphere above is a mirror, while the pink mesh surrounding it is diffuse at the microscale. This does not change the vNDF/transmittance, but has a large impact on the resulting appearance.

provide details on each individual substep in Section 3.2 of the supplemental.

Note that despite the seemingly intricate procedure described in this section, our method is not difficult to implement, assuming one has access to a path tracer and methods to sample conditioned values from a Gaussian process—which are available in many libraries and easy to implement from scratch. Notably, we did not have to change any of the core integrator code—a GPIS is implemented as a medium with a special phase function and free-flight distance sampling routine; the data for mean and covariance is accessed through the same volume data API that is used for heterogeneous density in classical media; we expand our renderer's existing architecture for non-exponential transport to keep track of the path history.

5 THE APPEARANCE SPACE OF GPISES

Now that we have a theory that allows for representing and rendering GPIS light transport, the remaining question is how to obtain such GPISes in the first place.
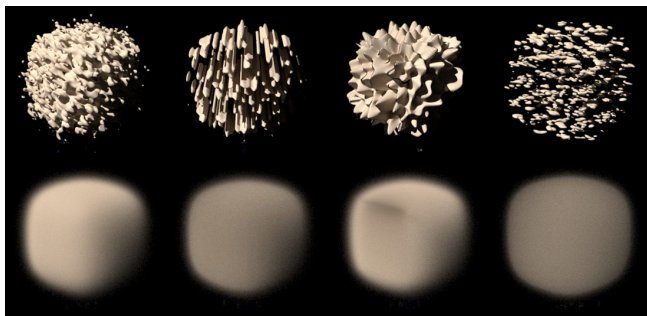

Fig. 10. At the top we show example realizations of Gaussian processes with different covariance kernels, all having the same microfacet roughness α. On the bottom are renders of the ensemble average light transport. Our method makes it trivial to explore light transport in a range of microgeometry without having to generate, store, and process large amounts of data.

Much like traditional scene representations, a GPIS can be manually authored by an artist, either by explicitly specifying its parameters (Section 6.2.1) or implicitly by constraining the GPIS at select points (Section 6.2.2). GPISes may also be combined via CSG operations (Section 6.2.3). The novel degrees of freedom in a GPIS allow to simultaneously specify a surface as well as its "fuzziness"/"certainty", unlocking a new spectrum of authoring workflows.

Beyond artistic control, we show how to render GPISes with parameters derived from data, for example via stochastic Poisson surface reconstruction (Section 6.3). Because GPISes can express correlations at the micro- and mesoscale, they are also a useful tool for level-of-detail, and we show how to obtain GPISes that represent prefiltered implicit surfaces at any desired resolution (Section 6.4).

Inherent to GPISes is the averaging over stochastic microgeometry, and as such they naturally subsume many classical stochastic appearance models. In the following subsections, we focus on two in particular that have seen much popularity: Microfacet surfaces, and participating media. The microgeometry assumed by either of these models can be naturally represented by a GPIS (see Fig. 9). However, because GPISes make fewer approximating assumptions, we can actually achieve more accurate renderings of the appearance described by these classical models, by incorporating correlations along paths through the microgeometry. Beyond making connections to existing appearance models, our method allows us to easily experiment with new microgeometry, and it provides a ground truth against which analytic approximations may be compared (see Fig. 10).

5.1 Joint distribution of free-flight distances and normals

The appearance of a microfacet surface is characterized by a normal distribution, a microfacet BSDF, and a shadowing-and-masking term. Similarly, a participating medium is described by its phase function and the transmittance function (or equivalently, its free-flight distribution). Though obscured by the rigorous treatment of correlations, these quantities also appear in our recursive GPIS light transport formulation, which we make explicit in the following.
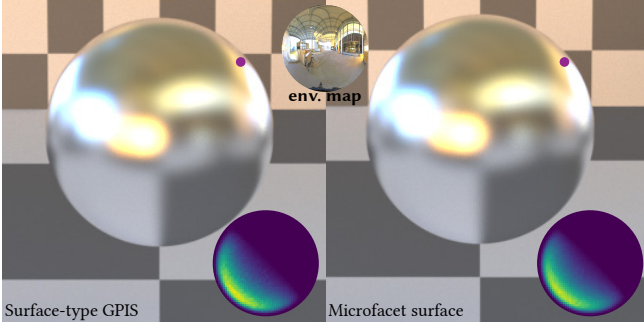
Fig. 11. An appropriately chosen surface-type GPIS can match microfacet models. We show renderings of a scene lit by an environment map, containing a surface-type GPIS with mirror microfacets (left) and a Beckmann microfacet model (right) with roughness matched via Eq. (20). The insets show the vNDF for a ray hitting the surface at a grazing angle, as computed by our method and derived from microfacet theory. For the relatively smooth surface shown here, the Smith approximation is close to the ground truth computed with our method as expected [Bourlier et al. 2000].

Under the Renewal+ and Renewal models, Eq. (14) can be simplified further, as $\langle L(\mathbf{x}_t, \boldsymbol{\omega}_t) \rangle_{\zeta, \zeta'}$ no longer depends on a specific realization $f_{(\mathbf{x}, \mathbf{x}_t)}$ along the incoming ray segment. This yields

$$\langle L(\mathbf{x}, \boldsymbol{\omega}) \rangle_\zeta \approx$$
$$\int_0^\infty \iint \rho(\mathbf{x}_t) \underbrace{\Gamma(t, \mathbf{n} \mid \zeta)}_{\substack{:=\gamma_{\mathbf{x}_t}(0, \mathbf{n} \mid \zeta) T(\mathbf{x}_t \mid \zeta) \\ \text{Joint probability density of} \\ \text{normal } \mathbf{n} \text{ and free-flight distance } t}} \langle L(\mathbf{x}_t, \boldsymbol{\omega}_t) \rangle_{\zeta \wedge \zeta'} \, d\boldsymbol{\omega}_t \, d\mathbf{n} \, dt, \quad (17)$$

with

$$T(\mathbf{x}_t \mid \zeta) = \int_{\text{GP}_{(\mathbf{x}, \mathbf{x}_t)} \mid \zeta \wedge \zeta_\delta} I^f(0, t) \, d\gamma(f_{(\mathbf{x}, \mathbf{x}_t)} \mid \zeta \wedge \zeta_\delta)$$
$$= P(f_{(\mathbf{x}, \mathbf{x}_t)} > 0 \mid \zeta \wedge \zeta_\delta) \quad (18)$$

being the probability of sampling *any* realization $f_{(\mathbf{x}, \mathbf{x}_t)} > 0$ conditioned on $\zeta \wedge f(\mathbf{x}_t) = 0 \wedge \overline{\nabla} f(\mathbf{x}_t) = \mathbf{n}$. We provide a step-by-step derivation in Section 1.3 of the supplemental.

The "GPIS density" $\Gamma(t, \mathbf{n} \mid \zeta)$ allows us to connect our method to participating media and microfacet surfaces, and derive parameters for our model to match their appearance.

5.2 Surface-type GPISes

Microfacet surfaces, much like GPISes, describe surfaces as realizations of a stochastic process—in fact, the popular Beckmann model [Beckmann 1965] even assumes Gaussian process height fields. A crucial microsurface property is the *distribution of visible normals* (vNDF) $D_v(\mathbf{n} \mid \boldsymbol{\omega})$, which describes the distribution of surface normals visible from direction $\boldsymbol{\omega}$. We can derive a vNDF directly from the GPIS density as

$$D_v(\mathbf{n} \mid \boldsymbol{\omega}) = \int_0^\infty \Gamma(t, \mathbf{n} \mid \boldsymbol{\omega}, \zeta) \, dt. \quad (19)$$

This gives the vNDF for any GPIS. To match microfacet theory, we can explicitly induce a heightfield surface via a linearly varying
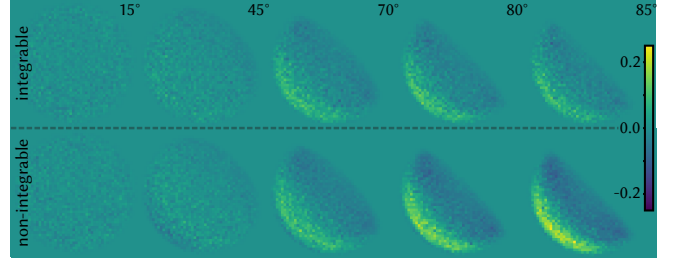

Fig. 12. We show the difference between the vNDF computed by our method and the vNDF predicted for a Beckmann surface with Smith shadowing and masking at various incident ray angles (left to right) and for different covariance kernels (top vs. bottom) for a relatively rough surface. As expected [Bourlier et al. 2000], at non-grazing angles, the Smith approximation is excellent, independent of the covariance kernel, whereas at grazing angles non-integrable covariance kernels (such as the rational quadratic kernel shown here) induce significantly different vNDFs.

mean along the z-axis $\mu(\mathbf{p}) = \mu(\mathbf{p}_z)$ and a strongly anisotropic kernel $k(\mathbf{p}, \mathbf{q}) = k(\mathbf{p}_{x,y}, \mathbf{q}_{x,y})$ that induces complete correlation along the z-axis. What remains is how to determine the parameters of the covariance kernel to match a given microfacet model.

Classical microfacet theory disregards certain correlations [Smith 1967], making the vNDF only depend on α, the "roughness" of the surface, which we can derive for any kernel [Pharr et al. 2016] as

$$\alpha = \sqrt{-2 \, k''(0)}. \quad (20)$$

As long as the kernel is twice differentiable, we can easily find a GPIS that matches a given microfacet surface by (1) aligning the gradient of the mean with the macro surface normal, and (2) picking a kernel with perfect correlation along the macro surface normal and $k''(0) = -\alpha^2/2$. In Fig. 11, we verify that our method can produce results consistent with microfacet theory by rendering a heightfield GPIS matched to a classical microfacet surface via Eq. (20). Note that without the Smith approximation to shadowing and masking, the vNDF depends on the full shape of the covariance kernel. In our method, we do not apply this approximation and we can show (Fig. 12) that, while the Smith approximation is excellent in some scenarios, for grazing angles and certain kernels, taking correlations into account does matter.
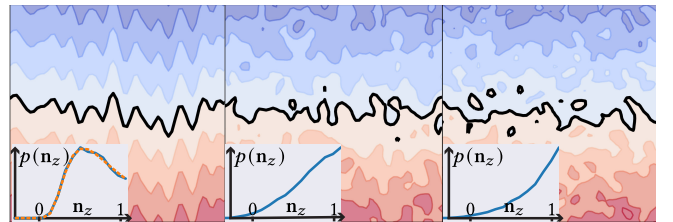

Fig. 13. As we move from the heightfields produced by strongly anisotropic covariance kernels (left) towards isotropic ones (right), the NDF (inset, blue) goes from the one predicted for Beckmann surfaces (dashed orange) towards a von-Mises-Fisher distribution, as discussed by d'Eon [2021].
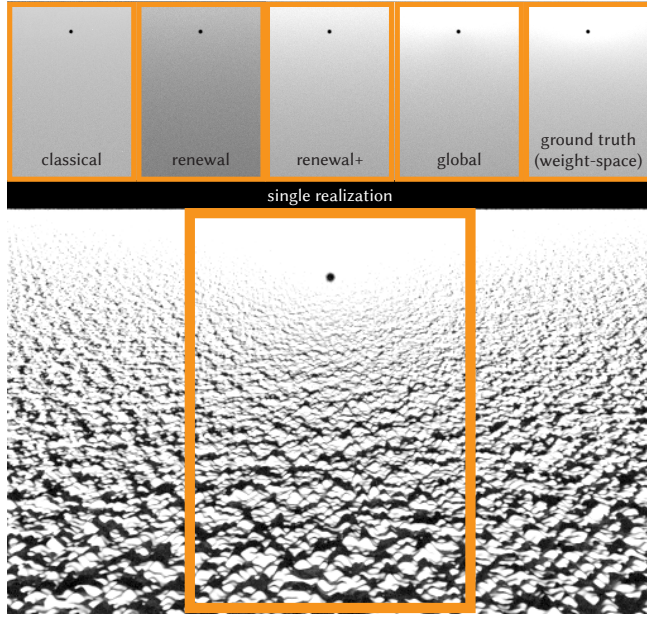
Fig. 14. We demonstrate the opposition effect with a single realization drawn from a Gaussian Process (bottom). The "shadow hiding" phenomenon is very visible in the image center. The top row shows different methods of computing the ensemble average light transport: Classical microfacet theory with uncorrelated shadowing and masking (top left) does not produce the opposition effect at all. Our memory models produce results closer to the ground truth as we increase the types of correlation we keep track of: Renewal model loses energy due to the lack of gradient continuity between "bounce" realizations; Renewal+ captures some backward scattering, but its conditioning is not enough to ensure perfect correlation along backward scattering paths; and the Global model near perfectly reproduces the bright halo around the observer's shadow

NDFs for non-heightfield surfaces. Our theory is not restricted to heightfields, and we can model non-heightfield surfaces and their corresponding (v)NDFs by relaxing the correlation constraint along the z-axis. In Fig. 13, we show empirically that these start to approach prior results for aggregate NDFs [d'Eon 2021]. Intuitively, this is explained by examining individual realizations, which start to resemble porous surfaces rather than bumpy heightfields.

Opposition effect. Rough surfaces lit from the viewing direction exhibit the "opposition effect", a noticeable brightening of the surface around the observer's shadow (Fig. 14). The main contributing factor to this effect—"shadow hiding"—is not captured in shadowing/masking terms derived from Smith theory, but can be reproduced in varying degrees of accuracy using our theory, depending on the memory model. The Renewal+ model provides a reasonable performance/accuracy trade-off.

5.3 Volume-type GPISes

The central quantity in volumetric transport is the free-flight distribution. We can derive it for our model by integrating the GPIS density over the sphere of all possible normals at the intersection

point x_t:

$$\Gamma(t \mid \zeta) = \int_{S^2} \Gamma(t, \mathbf{n} \mid \zeta)\, d\mathbf{n}. \tag{21}$$

$\Gamma(t \mid \zeta)$ represents the probability density of finding the first zero crossing of the GPIS at distance t. Assuming the Renewal model, it is sufficient to examine the case when $\zeta_c \coloneqq f(\mathbf{x}) = 0$ (i.e. the current ray starts on a zero crossing, e.g. after leaving a scattering event) and $\zeta_u \coloneqq f(\mathbf{x}) > 0$ (i.e. the ray origin is uncorrelated with the GPIS e.g. starting at the camera or leaving a non-GPIS surface). This is analogous to recent work in non-exponential transport [Bitterli et al. 2018], and the same properties apply: $\Gamma(t \mid \zeta_c)$ can be derived from $\Gamma(t \mid \zeta_u)$ and vice-versa (see Fig. 15).

5.3.1 Matching participating media. Just as in microfacet theory, disregarding correlations along light paths allows us to describe a medium by a single number, its *density* σ_t, i.e. a value indicating how many particles occupy a certain region of space. As before, we want to derive the parameters, i.e. σ_t, of this simplified classical model from the parameters of a GPIS and vice versa. Since volume-type GPISes have not been studied much in graphics literature, this requires us to pull in knowledge from a wider range of literature on stochastic processes. There, the free-flight distributions $\Gamma(t \mid \zeta_{u/c})$ are known as the *first-passage-time* (for ζ_u) and the *excursion-time density* (for ζ_c). The first-passage-time is the "time" (in our case distance) at which a realization of a stochastic process first crosses a particular level. The excursion time is then the time that a realization of a stochastic process takes to *return* to the given level for the first time. This corresponds to the distance from one intersection point to another. In the Gaussian process case, it is known that excursion times are well-defined only for *smooth processes* [Bray et al. 2013; Torquato and Lu 1993]. A process is smooth iff the Taylor expansion of $k(\mathbf{x}, \mathbf{y})$ has the form $k(\mathbf{x}, \mathbf{y}) = 1 - a\|\mathbf{x} - \mathbf{y}\|^\alpha + O(\|\mathbf{x} - \mathbf{y}\|^\alpha)$, with $\alpha = 2$ [Bray et al. 2013]. For non-smooth processes—like the Ornstein-Uhlenbeck (OU) process and Brownian motion—the excursion time density is simply a delta function at 0. For us this means that any photon leaving a scatterer would scatter again immediately, no progress is made and no light transport is possible. Hence we have to limit ourselves to smooth processes. Unfortunately, it is also true that only Markov processes have known analytic solutions for first-passage times [Buonocore et al. 1987]. Since the only Gaussian Markov process is the OU process, which is not smooth, there are
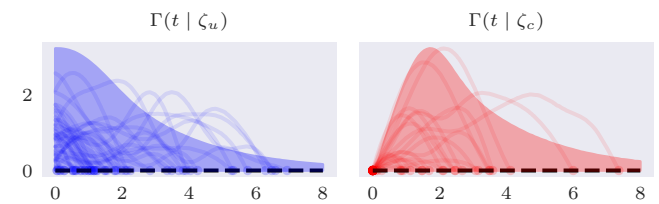

Fig. 15. Left: Realizations sampled along rays starting in free-space. The resulting distribution of zero crossing distances is the *first-passage time density* and corresponds to $\Gamma(t \mid \zeta_c)$. Right: Realizations sampled along rays starting at a GPIS intersection. The distribution of zero crossing distances here is the *excursion time density* and correspond to $\Gamma(t \mid \zeta_c)$. Realizations are sampled from a GP with squared exponential covariance and zero mean.
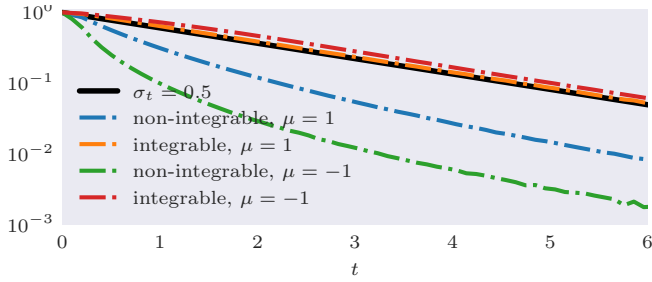
Fig. 16. We numerically compute transmittance using RIND (dashed) for a range of homogeneous GPISes, allowing us to derive the extinction coefficient $\sigma_t = b_k$ of the closest exponential medium (solid black). Based on the differences between the dashed and solid lines, we can confirm that processes with a higher mean μ and integrable kernels produce more exponential first-passage times.

no Gaussian processes for which we (1) can hope to derive a non-zero particle/particle free-flight density and (2) for which such a free-flight density has a known analytic form. That said, it is well known that for many covariance kernels, the *tail behaviour* of the first-passage time is exponential. In particular, we can write the *persistence exponent* b_k [Dembo and Mukherjee 2017] for a given kernel k as

$$b_k := -\lim_{t \to \infty} \frac{1}{t} \log \Gamma(t \mid \zeta_u), \qquad (22)$$

with $b_k \in [0, \infty]$ as long as k is non-negative and stationary. Note that if $b_k = 0$, we have longer-than-exponential tails for the first-passage time. This is the case iff $\int_0^\infty k(t) \, dt = \infty$, i.e. when the covariance kernel is non-integrable. In these cases, the decay rate is often of the form $t^{-\alpha}$, with $\alpha \in (0, 1]$. There are strong similarities to the Davis model of non-exponential transmittance [Davis and Mineev-Weinstein 2011], which Bitterli et al. [2018] previously imported to graphics, where the authors derive longer-than-exponential tails in cases where the random fluctuations of the medium are "self-similar". Now, matching σ_t to the persistence exponent b_k produces
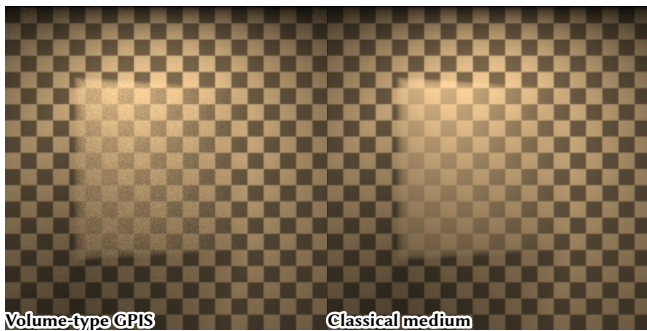

Fig. 17. Here we show results for an index-matched homogeneous medium in a "wedge" (thicker on the left, thinner on the right), set in front of a checkerboard, and lit from above. On the left we render a volume-type GPIS with parameters based on the persistence exponent match discussed in this section and a mirror BSDF. On the right, we render the classical medium with the σ_t we used to match the GPIS and an isotropic phase function.
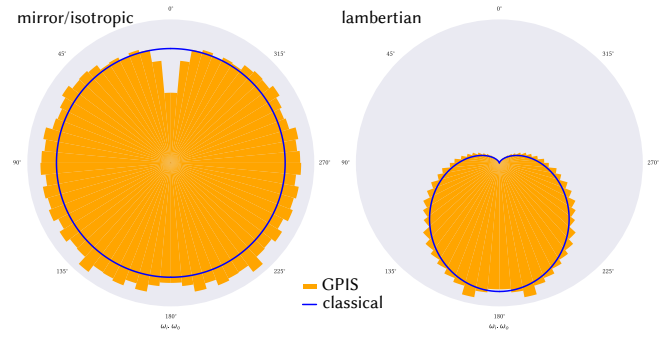

Fig. 18. Isotropic volume-type GPISes induce phase functions that approach ones predicted by prior work. For a GPIS with a mirror BSDF, we recover a close-to-isotropic phase function (left), whereas for a GPIS with a diffuse BSDF we approach the phase function for Lambertian spherical scatterers described by [d'Eon 2021] (right).

the classical medium "closest" to a given GPIS, which is what we have been looking for. While there are no known analytical results to compute b_k, Lindgren et al. [2020] suggest to derive the persistence exponent from numerical first-passage time results computed via their method RIND. RIND is efficient and accurate, and we can easily fit an exponential to the tail of the numerical results as shown in Fig. 16. In Fig. 17 we show renders using parameters computed with this method to verify its use in practice.

Phase function reproduction. Additionally, we can show that the phase functions in our volume-type GPISes behave as expected. In particular, as we increase the mean (making the medium more dilute) and decrease the length scale (reducing the size and correlation of particles), we expect the phase function to be increasingly isotropic for a GPIS with a mirror micro BSDF. Applying a diffuse micro-BSDF should reproduce the phase function for Lambertian scatterers. Figure 18 shows that this is the case. We pose that the dip in the forward scattering direction for the mirror BSDF (left) is due to our ray marching based intersection routine being more likely to miss intersections at grazing angles.

5.4 The in-betweens & mesoscale geometry

We have discussed two extreme ends of the appearance spectrum that we can cover with GPIS, but for both of them, specialized methods exist which perform extremely well compared to our more general method. The value that we see in our method is that it not only does not require an *a priori* choice between either end of the spectrum, but that it can also represent cases that don't fit neatly into either category. Particularly interesting is that we can represent *mesoscale* detail easily, which is something that other geometry models struggle with. We define mesoscale detail as geometry that is fine enough to require relatively large amounts of storage to represent explicitly, while being large scale enough that the assumptions necessary for efficient microscale models break down. This covers, for example, tree bark and leaves, grains, and fibers or fur, and we show some didactic examples in Fig. 10. We leave the further investigation to future work, but we believe that at mesoscale level uncertainty, our representation behaves much like certain classes
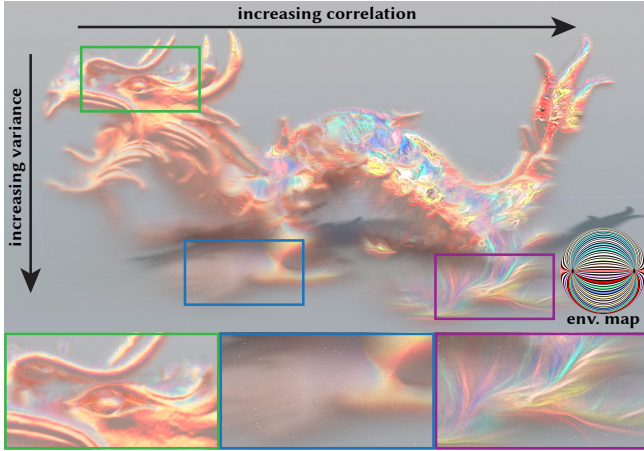
Fig. 19. Using the model we present in Section 6.1 we can spatially vary all of the parameters we have identified as important to the appearance in Section 5. As shown here, this allows us to represent these different types of appearance not only in the same scene, but in the same *geometry primitive*, with smooth transitions between them.

of anisotropic media [Jakob et al. 2010]. In Gaussian process terms, mesoscale detail appears when the *variance* of the process is on the same scale as variations in the mean. For example, in Fig. 10, the mean shape has an extent of 50 units in each direction, while the variance is 25 for all examples.

6 A NON-STATIONARY GPIS MODEL FOR ACQUISITION & CREATION

In this section, we present a concrete, flexible model that covers a wide range of practical use cases. In particular this model supports the whole range of volume- to surface-type GPIS discussed in Section 5 as well as stochastic mesoscale detail. We restrict ourselves to *non-stationary Gaussian process implicit surfaces* with a parameterized prior mean function and covariance kernel with parameters Φ, as well as a set of conditioning points C, where each point $\mathbf{c} \in C$ is associated with a location $\mathbf{c}_\mathbf{x}$, a value $\mathbf{c}_v$ and a normal derivative direction c_∇. Points that condition the *value* of the process are encoded as $\mathbf{c}_\nabla = 0$ and points that condition the *derivative* of the process as $\mathbf{c}_\nabla \neq 0$. The resulting process, $\mathrm{GP}(\mu_{\Phi|C}(\mathbf{x}), k_{\Phi|C}(\mathbf{x}, \mathbf{y}))$, has mean and covariance given by

$$\mu_{\Phi|C}(\mathbf{x}) = \mu_\Phi(\mathbf{x}) + k_\Phi(\mathbf{x}, C_X) k_\Phi(C_\mathbf{x}, C_\mathbf{x})^{-1}(C_v - \mu_\Phi(C_\mathbf{x})) \quad (23)$$

$$k_{\Phi|C}(\mathbf{x}, \mathbf{y}) = k_\phi(\mathbf{x}, \mathbf{x}) - k_\Phi(\mathbf{x}, C_\mathbf{x}) k_\Phi(C_\mathbf{x}, C_\mathbf{x})^{-1} k_\Phi(C_\mathbf{x}, \mathbf{y}), \quad (24)$$

which are simply the standard posterior mean and covariance kernel we introduced in Section 3. We are left with having to choose a prior covariance kernel k_Φ.

6.1 A non-stationary kernel

Since appearance depends strongly on the parameters of the kernel we would like to be able to vary them spatially, e.g. to allow for spatially varying roughness in a surface-type GPIS. Building non-stationary covariance kernels with control over local behavior is difficult in general since it can be hard to prove that the global
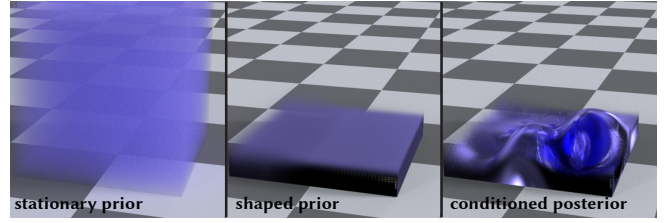
Fig. 20. We show the effect of the different editing methodologies we discuss in Section 6.2. For example, we can start with a volume-type GP (i.e. homogeneous mean and isotropic covariance, left) and use prior shaping to insert a solid object (e.g. a plane with varying uncertainty, middle) and then deform that object using posterior conditioning (9 handles, right).

function is positive semi-definite. For example, we can't simply average between different locally stationary kernels. In the following, we choose to use the non-stationary covariance kernel derived by Paciorek and Schervish [2006]:

$$k_\Phi^{NS}(\mathbf{x}, \mathbf{y}) = \sigma_\Phi(\mathbf{x}) \sigma_\Phi(\mathbf{y}) \frac{|\Sigma_\Phi(\mathbf{x})|^{\frac{1}{4}} |\Sigma_\Phi(\mathbf{y})|^{\frac{1}{4}}}{|\frac{\Sigma_\Phi(\mathbf{x}) + \Sigma_\Phi(\mathbf{y})}{2}|^{\frac{1}{2}}} k_\Phi^S(\sqrt{Q_\Phi(\mathbf{x}, \mathbf{y})}), \quad (25)$$

where $k_\Phi^S(t)$ is any stationary covariance kernel and

$$Q_\Phi(\mathbf{x}, \mathbf{y}) = (\mathbf{x} - \mathbf{y})^\top \left(\frac{\Sigma_\Phi(\mathbf{x}) + \Sigma_\Phi(\mathbf{y})}{2}\right)^{-1} (\mathbf{x} - \mathbf{y}). \quad (26)$$

This kernel is characterized by the parameters of the global stationary kernel k_Φ^S, as well as two spatially varying fields: the "local variance" $\sigma_\Phi : \mathbb{R}^3 \to \mathbb{R}$ and the "local anisotropy" $\Sigma_\Phi : \mathbb{R}^3 \to \mathbf{S}_+^3$ where $\mathbf{S}_+^3$ is the set of positive semi-definite 3×3 matrices. Note that Σ_Φ is in concept very close to how anisotropy is expressed in the *SGGX* [Heitz et al. 2015] microflake phase function. The kernel k_Φ^{NS} is very flexible, has intuitive parameters, and has been used successfully to model complex priors [Dexheimer and Davison 2023]. We show an example of the results that we can generate by varying these parameters spatially in Fig. 19.

6.1.1 Model Storage. We store the mean, variance, and anisotropy fields, $\mu_\Phi(\mathbf{x})$, $\sigma_\Phi(\mathbf{x})$, and $\Sigma_\Phi(\mathbf{x})$, on a voxel grid and retrieve values at arbitrary evaluation locations via interpolation. For $\Sigma_\Phi(\mathbf{x})$, we use the encoding and linear interpolation method proposed by Heitz et al. [2015]. This allows for smooth interpolation while guaranteeing that the interpolated matrices are positive semi-definite.

6.2 Manual editing of GPIS

The most basic way to create GPIS scenes is to specify the underlying fields directly. This is how we created most of the didactic examples so far since it allows for a high degree of control. We propose three intuitive editing approaches, which we show in Fig. 20.

6.2.1 Implicit control via prior shaping. The mean can be edited like any other classical implicit surface. That is, we can easily apply CSG operations and warping functions, offset the isosurface, and more. Similarly, the variance field can be edited like any other scalar field, for example, via 3D "brushes" that modify the stored values in their influence region *à la* the looseness control in DREAMS [Media Molecule 2020]. For the anisotropy field, we can similarly provide

"anisotropy brushes" that paint the appropriate vector fields to parameterize the anisotropy. We created most of the simple scenes in this paper using this type of editing, as it provides very direct control over all aspects of the GPIS.

6.2.2 Explicit control via posterior conditioning. Additionally, we can allow the user to exploit the *conditioning* capabilities of the Gaussian process representation. To do so, we enable the user to specify points at which they can place value and derivative constraints. This is, in principle, an extension to Witkin and Heckbert's control particle approach [Hart et al. 2002; Witkin and Heckbert 1994] for controlling implicit surfaces. The GPIS is then conditioned on taking on the specified values using classical Gaussian process regression techniques as discussed in Section 3.1.3. When rendering, we can simply add the "global" conditioning to our path conditioning variable ζ, since they behave exactly like any other conditioning we introduce along the path (except that they are not "forgotten" by any memory model).

6.2.3 CSG on GPISes. Ideally, one would like to perform CSG operations on GPISes directly. Unfortunately, GPISes are only closed under linear transforms (e.g. the sum of two GPISes is a GPIS, but the product is not). Many CSG operations contain non-linear operations, such as the min or the max. This means that we must support a certain set of non-Gaussian stochastic implicit surfaces to support the full range of CSG operations. Luckily, this is trivial to do in our current implementation. We can simply generate a realization for each leaf node in the CSG tree and then apply the CSG operations to the realizations to get the final sample. Our implementation hence supports the non-Gaussian SISes produced by the non-linear combination of "leaf GPs". This flexibility does, however, mean that the cost of generating realizations depends on the number of leaves in the CSG tree, i.e. we can not "bake down" to a more compact representation. Additionally, we must keep track of path conditioning variables for *each* leaf GPIS. A lossy alternative is to fit a GPIS to the non-Gaussian SIS. We show the results of this when applied to the intersection of two sphere GPISes with different variances in Fig. 21. CSG GPISes are not only useful for artistic control but have also proven themselves as a rich model for real-world microstructure in other fields [Torquato and Lu 1993].
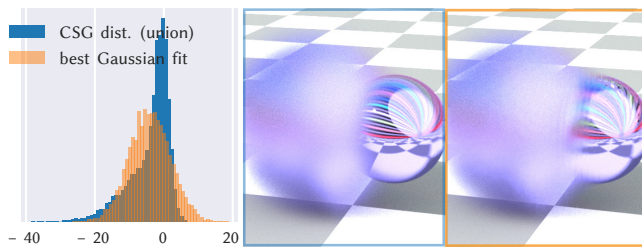

Fig. 21. We can use CSG operations to combine multiple GPISes (left, blue: point distribution). The resulting process is not Gaussian anymore, but our rendering algorithm is still able to visualize it by sampling from each GPIS individually and applying the CSG operations on each set of samples (middle). If we fit the closest Gaussian to the CSG distribution (left, orange), we can represent the result approximately with a single GPIS (right).
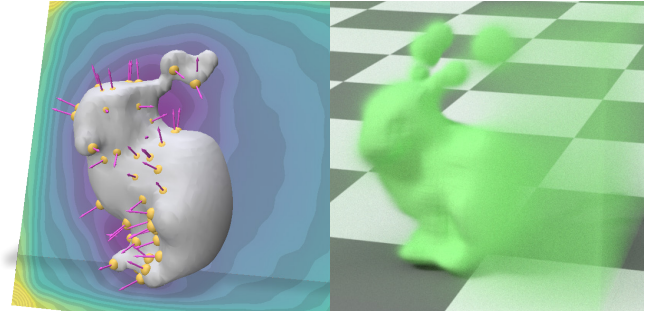

Fig. 22. We compute the light transport through the GPIS defined by the mean and variance fields that *stochastic Poisson surface reconstruction* (SPSR) [Sellán and Jacobson 2022] produces (left) with our rendering method (right). Note that the surface is well-defined in the densely sampled area on the left, while regions with fewer samples produce a fuzzy volume-like appearance.

6.3 Stochastic Poisson surface reconstruction

Recently, Sellán and Jacobson [2022] presented a surface reconstruction method, based on the widely used Poisson surface reconstruction algorithm, that recovers both the mean implicit surface as well as a spatially varying variance field. We can visualize the output of their method using our rendering algorithm and show a simple example in Fig. 22. While Sellán and Jacobson [2022] do show a rendering-related example and even take the full correlations along single rays into account [Sellán and Jacobson 2023], they are mainly interested in the uncertainty of primary hit locations and do not consider global illumination.

6.4 Filtering implicit surfaces

Inspired by the connections between microfacet surfaces and GPISes that we discussed in Section 5.2, we present a method to use GPISes for implicit surface downsampling by representing the sub-resolution surface details via their *statistics*. In the following, we show how to leverage this intuition for a principled way to construct low-capacity GPISes that approximate high-capacity surface models well. This is, in spirit, similar to methods that learn the statistics of textures from examples [Galerne et al. 2012; Guehl et al. 2020] or downsample normal maps by fitting distributions over the filtered area [Dupuy et al. 2013; Olano and Baker 2010].
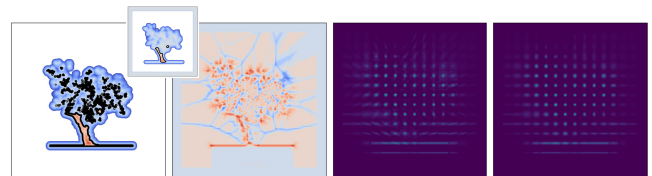

Fig. 23. We downsample an implicit surface with non-stationary statistics (left) by first separating it into a mean field (inset) and a residual (middle-left). We then use a short-time Fourier transform based method to recover an approximation to the non-stationary autocorrelation function (ACF) of the residual (middle-right). Finally, we fit a stationary kernel to each of the ACF blocks (right). We can then interpolate the parameters of this kernel to compute the global non-stationary covariance kernel in Eq. (25).
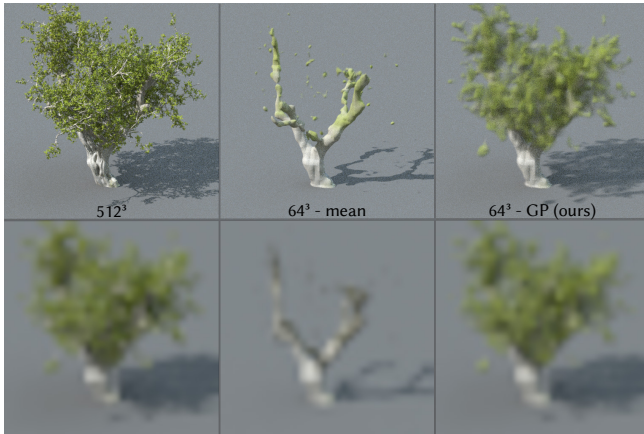
Fig. 24. We show full-resolution renders on the top and 32x32 pixel images at the bottom. If we simply downsample the high-resolution mean surface (left), we get a hard surface, which does not represent any of the higher frequency detail that was filtered out, leading to drastically different light transport (middle). With our representation, we can create a GP that approximates the statistics of the filtered data, preserving light transport better (right).

Given some high-resolution representation of an implicit surface $f(\mathbf{x})$, we want to *downsample* it, for example, to use less storage, while preserving the geometry information and any derived quantities, such as light transport, as well as possible. In general, we parameterize the downsampled approximation $\hat{f}_\Phi$ by some finite set of values Φ, chosen such that the *residual* $R_\Phi(\mathbf{x}) = \hat{f}_\Phi(\mathbf{x}) - f(\mathbf{x})$ is minimal over the domain, i.e. Φ is chosen to minimize $\int_\Omega R_\Phi(\mathbf{x})^2 \, d\mathbf{x}$. Traditionally, this is done using maximum likelihood estimation, but this discards any information other than the *mean*. This results in a drastic difference in the light transport as we show in Fig. 24. We can go one step further and instead find the Gaussian process with mean $\mu_\Phi(\mathbf{x}) = \hat{f}_\Phi(\mathbf{x})$ that is most likely to sample $f(\mathbf{x})$. That is, we effectively capture the original function as well as possible with the limited capacity we have and then estimate the *statistics* of the residual. The resulting process has the form $\text{GP}(\mu_\Phi(\mathbf{x}), k_\Theta(\mathbf{x}, \mathbf{y}))$. We now need to find the parameters Θ of the kernel k_Θ such that the prior process has statistics that are as close to possible to the residual. This is a common task in Gaussian process regression and the parameters can be found by minimizing the negative log marginal likelihood, but this process is expensive and can be unstable when the number of parameters is large. Instead, we exploit the local stationarity of our model and compute kernel parameters assuming stationarity over some finite region. We do so by numerically computing the autocorrelation function (ACF) of the residual using a Fourier transform-based method over subregions of the domain. We then find kernel parameters by fitting our model to the numerically retrieved ACF in each region. This gives us a grid of kernel parameters, which we can interpolate to approximate the full nonstationary kernel. We illustrate this process in Fig. 23. Note that here we only discuss fitting based on geometric error. We could alternatively define other loss functions, such as an appearance-based one, or one that aims to match transmittance between points in space

[Vicini et al. 2021]. We leave this, and exploring how our method relates to recent work that exploits correlation-aware downsampling for level of detail [Weier et al. 2023], to future work.

7 LIMITATIONS OF GPISES

Gaussian process implicit surfaces inherently share the characteristics of implicit surfaces, enabling robust representation capabilities. However, this also implies certain challenges inherent to implicit surfaces. For instance, while editing tools are generally more developed for explicit representations like meshes, GPISes face challenges in applying techniques like texture mapping due to complexities in parameterizing the 3D surface. Further, the lack of analytic expressions for transmittance and normal distributions prevents our current formalism from being applied to techniques such as differentiable rendering. Future progress in this respect would be of great interest in light of the recent focus on neural radiance field techniques and related representations.

Performance. In terms of rendering, implicit surfaces, including those not based on Gaussian processes, are typically slower than triangle mesh rendering, which may limit their use in real-time applications like video games. Nonetheless, for problems where strong spatial correlations in stochastic geometry significantly influence the light transport, our proposed rendering framework is orders of magnitude more efficient than a brute force global-sampling approach and enables novel appearances that were not previously practical. The performance of our method is mainly determined by the amount of correlations one keeps track of—both across ray intersections and within one ray segment, see Section 3 of the supplemental. We have found that the acceleration techniques we provide make it possible to produce images of simple scenes (e.g. Figs. 1, 9 and 19) in a "reasonable" amount of time on a standard desktop PC, i.e. around 30 seconds for a 512×512 pixel render at 1 sample per pixel when using the Renewal+ memory model. While more than an order of magnitude slower than rendering e.g. microfacet surfaces directly, it is fast enough for us to explore the expressive power of GPISes and efficiently iterate on improvements in future work. In contrast, consider the global-sampling method of sampling many realizations and averaging many rendered images. In Fig. 1, for example, the variations in the surface of the rhinoceros are on the order of 1 mm with the statue being roughly 1.6 m tall. To resolve these details we would need to discretize to a grid of around 3200^3 and to sample a realization at that resolution then requires inverting a $3200^3 \times 3200^3$ matrix, which would need 4,294,967,200,000 GB for storage alone.

Discrete sampling/ray marching. The function-space algorithm we present in Section 4.2 represents realizations along ray segments using a discrete set of samples. Finding the first zero-crossing is then done via ray marching. If the step size is too large, we encounter the usual issues that appear when finding implicit surface ray intersections with ray marching, such as missing small surface features. In practice, we choose the step size relative to the correlation length of the covariance kernel (e.g. we set Δt such that $k(x, x + \Delta t \vec{\omega})/k(x, x) > 0.95$). A more advanced "stochastic root finding" algorithm and smarter sample distribution are interesting

topics for future work that have the potential to improve the performance of GPIS rendering without large structural changes to the method we present here.

Lack of next-event estimation. As mentioned in Section 1, we currently only support next-event estimation (NEE) when using diffuse or glossy micro-BSDFs, but not when using a perfect mirror micro-BSDF. This is because we treat the normal sampled at the intersection point (see Alg. 1) as a deterministic quantity when evaluating the micro-BSDF. While this choice does reduce implementation complexity, it is not a fundamental limitation of our approach—after all, the normal is *sampled* from a distribution, and hence, we should be able to apply NEE. We present some thoughts on how this could be done in Section 3.3 of the supplemental.

8 FUTURE WORK AND CONCLUSION

We see the example use cases that we presented in Section 6 as validating the expressiveness of GPISes, but concede that they fall short of being "applications" immediately useful in practice. Hence we see much room for future work.

8.1 Correlated appearance

For a full rendering method, we need to be able to retrieve additional properties, such as surface albedo, at intersection points. Of course, we can simply store these in a separate volumetric data structure, which we do in this work, but Gaussian processes provide us with an elegant mechanism for a more principled solution. Multi-output or multi-task GPs predict vector-valued functions $f_N : \mathbb{R}^D \to \mathbb{R}^N$ and allow for cross-correlations in their outputs. We should be able to extend the fitting approach presented in Section 6.4 to include BSDF parameters.

8.2 Solving elliptic partial differential equations

A central property of this volumetric *geometry* representation is that it can be used for a wide variety of geometry processing tasks instead of just rendering. In particular, we are optimistic that a large part of the derivations in Section 4 are applicable to other recursive integral equations, such as the ones that appear in walk-on-sphere [Sawhney and Crane 2020] and walk-on-boundary [Sugimoto et al. 2023] methods. This should allow us to solve certain classes of elliptic partial differential equations while taking uncertainty in the domain boundary and boundary conditions into account.

8.3 More general stochastic processes

So far we have limited ourselves to *Gaussian* processes. This is a common choice because they possess many desirable properties, such as being closed under conditioning and taking derivatives. More general stochastic processes would allow for an even more expressive function space, but at the cost of fewer analytical results. There is existing work that aims to use Gaussian processes to approximate more general stochastic processes [Gurley 1997; Hong et al. 2023; Shields et al. 2011], but any Gaussian process will always produce Beckmann normal statistics. This is limiting because prior work has shown that normal distributions in the real-world often have longer tails [Walter et al. 2007]. A logical first choice for a more complex class of stochastic processes would be the student-t process,

which is essentially a weighted sum of infinitely many Gaussian processes. There is recent work on student-t microfacet surfaces [d'Eon 2023] that shows that the resulting vNDFs match real-world results very well.

8.4 Advanced weight-space methods

The global realization sampling method via the weight-space approximation is very attractive because it conceptually simplifies GPIS rendering. Unfortunately, the standard formulation of weight-space sampling via an expansion of basis functions according to the process' spectral density is limited to stationary covariance kernels. We are not aware of any prior work that extends this to general non-stationary kernels other than methods that use the Karhunen-Loève expansion [Huang et al. 2001], which are generally not seen as practical in applications. That said, we think it is worthwhile to explore how to extend weight-space sampling to support a, maybe limited, range of non-stationary kernels. We could imagine that assuming local stationarity and blending basis-function representations spatially would produce acceptable results for certain applications.

8.5 Uncertainty Quantification

A natural question is whether we can propagate the uncertainty in the scene geometry to uncertainty in the rendered results. That is, instead of just computing the ensemble average, we would like to extract higher-order moments, primarily variance. This problem is known as "uncertainty quantification" [Smith 2013] and progress here would be useful in a range of real-world applications. Doing this efficiently when we only have access to Monte Carlo estimates of the ensemble average is non-trivial and we discuss some reasons for this in Section 4 of the supplemental. Intuitively, it requires us to untangle the variance due to uncertainty in the geometry from the variance due to Monte Carlo sampling of light paths. Also known as separating *parametric* uncertainty from *intrinsic* uncertainty, methods based on polynomial chaos [Mueller et al. 2023] seem promising but assume relatively low intrinsic uncertainty. Extending these to the high-variance Monte Carlo models common in graphics would allow us to not only compute variation due to geometry but also other scene parameters.

8.6 Conclusion

In this paper, we have shown that reasoning about the light transport in GPISes allows us to make deep connections between existing stochastic representations in graphics. We have done so by relying on Monte Carlo sampling of the involved integrals without being overly concerned with performance. However, real-world applications rely on the availability of performant methods. While compute power grows with time, and methods that were previously seen as intractable become widely used, we believe that there is much room for advancing the *theory* of GPIS rendering. For example, there is much interest in GPs for large data sets, and we expect that advances there would directly translate to faster rendering. Analytic approximations for first passage times are important in many areas of research and would allow us to decouple rendering times from sampling efficiency. Tabulating transport functions that cannot be

computed analytically ahead of time has been used in rendering before [Müller et al. 2016] and we could use similar techniques to render with tabulated first-passage times. Finally, improving the performance of Monte Carlo methods by orders of magnitude is the bread and butter of the rendering community and with this work we hope to inspire researchers to apply their knowledge to Gaussian process simulations, enriching both fields.

ACKNOWLEDGMENTS

REFERENCES

Robert J Adler and Jonathan E Taylor. 2009. *Random fields and geometry.* Springer Science & Business Media.

Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. 2021. Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields. *Proceedings of the International Conference on Computer Vision (ICCV)* (2021).

Petr Beckmann. 1965. Shadowing of random rough surfaces. *IEEE Transactions on Antennas and Propagation* 13, 3 (1965). https://doi.org/10/c3g9cs

Benedikt Bitterli, Srinath Ravichandran, Thomas Müller, Magnus Wrenninge, Jan Novák, Steve Marschner, and Wojciech Jarosz. 2018. A Radiative Transfer Framework for Non-Exponential Media. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 37, 6 (Nov. 2018), 225:1–225:17. https://doi.org/10/gfz2cm

Patrick Boissé. 1990. Radiative transfer inside clumpy media-The penetration of UV photons inside molecular clouds. *Astronomy and Astrophysics* 228 (1990). http://adsabs.harvard.edu/abs/1990A%26A...228..483B

Christophe Bourlier, Joseph Saillard, and Gerard Berginc. 2000. Effect of correlation between shadowing and shadowed points on the Wagner and Smith monostatic one-dimensional shadowing functions. *IEEE Transactions on Antennas and Propagation* 48, 3 (March 2000). https://doi.org/10.1109/8.841905

Alan J Bray, Satya N Majumdar, and Grégory Schehr. 2013. Persistence and first-passage properties in nonequilibrium systems. *Advances in Physics* 62, 3 (June 2013). https://doi.org/10/gfv7wc

Ken Brodlie, Rodolfo Allendes Osorio, and Adriano Lopes. 2012. A review of uncertainty in data visualization. In *Expanding the Frontiers of Visual Analytics and Visualization.* Springer London, London. https://doi.org/10/gjvc64

Antonio Buonocore, Amelia G. Nobile, and Luigi M. Ricciardi. 1987. A New Integral Equation for the Evaluation of First-Passage-Time Probability Densities. *Adv. Appl. Probab.* 19, 4 (1987). https://doi.org/10/c4q5hb

Anthony B. Davis and Mark B. 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, 4 (March 2011), 632–645. https://doi.org/10/fgxnc4

Amir Dembo and Sumit Mukherjee. 2017. Persistence of Gaussian processes: non-summable correlations. *Probab. Theory Relat. Fields* 169, 3-4 (Dec. 2017). https://doi.org/10/gcj348

Eugene d'Eon. 2021. An Analytic BRDF for Materials with Spherical Lambertian Scatterers. *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 40, 4 (2021), 153–161. https://doi.org/10/gsmm3v

Eugene d'Eon. 2023. Student-T and beyond: Practical tools for multiple-scattering BSDFs with general NDFs. In *ACM SIGGRAPH 2023 Talks.* ACM, New York, NY, USA. https://doi.org/10/mr42

Eric Dexheimer and Andrew J. Davison. 2023. Learning a Depth Covariance Function. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).* IEEE. https://doi.org/10/mr43

Stanimir Dragiev, Marc Toussaint, and Michael Gienger. 2011. Gaussian process implicit surfaces for shape estimation and grasping. In *2011 IEEE International Conference on Robotics and Automation.* ieeexplore.ieee.org. https://doi.org/10/bdfbb5

Jonathan Dupuy, Eric Heitz, and Eugene d'Eon. 2016. Additional Progress towards the Unification of Microfacet and Microflake Theories. In *Proceedings of EGSR (Experimental Ideas & Implementations).* Eurographics Association, 55–63. https://doi.org/10/gfz72z

Jonathan Dupuy, Eric Heitz, Jean-Claude Iehl, Pierre Poulin, Fabrice Neyret, and Victor Ostromoukhov. 2013. Linear Efficient Antialiased Displacement and Reflectance Mapping. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 32, 6, Article 211 (2013), 11 pages. https://doi.org/10/gbd53f

Eugene d'Eon. 2018. A reciprocal formulation of nonexponential radiative transfer. 1: Sketch and motivation. *Journal of Computational and Theoretical Transport* 47, 1-3 (2018). https://doi.org/10/gg5kzj

Eugene d'Eon. 2023. Beyond Renewal Approximations: A 1D Point Process Approach to Linear Transport in Stochastic Media. *ANS M&C 2023 - The international conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering - Niagara Falls, Ontario, Canada, Aug.* (2023).

David S. Ebert, F. Kenton Musgrave, Darwyn Peachey, Kenneth Perlin, and Steven Worley. 2003. *Texturing and modeling: a procedural approach* (3rd ed.). Morgan Kaufmann, San Francisco, CA, USA.

Anne Estrade, Ileana Iribarren, and Marie Kratz. 2012. Chord-length distribution functions and Rice formulae. Application to random media. *Extremes* 15, 3 (Sept. 2012). https://doi.org/10/dkrb8f

Nathaniel Fout and Kwan-Liu Ma. 2012. Fuzzy volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (Dec. 2012). https://doi.org/10/f4fq83

Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. 2022. Plenoxels: Radiance Fields without Neural Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR).* https://doi.org/10/gq34wg

Leonid Fukshansky, A. Martinez von Remisoksky, and John McClendon. 1993. Absorption spectra of leaves corrected for scattering and distributional error: a radiative transfer and absorption statistics treatment. *Photochemistry and Photobiology* 57, 3 (1993), 538–555. https://doi.org/10/dnthw2

Bruno Galerne, Ares Lagae, Sylvain Lefebvre, and George Drettakis. 2012. Gabor noise by example. *ACM Transactions on Graphics* 31, 4 (Aug. 2012). https://doi.org/10/gbbwm7

Manuel Gamito. 2009. *Techniques for Stochastic Implicit Surface Modelling and Rendering.* Ph. D. Dissertation. University of Sheffield. https://etheses.whiterose.ac.uk/116/

Pascal Guehl, Rémi Allègre, J-M Dischler, Bedrich Benes, and Eric Galin. 2020. Semi-procedural textures using point process texture basis functions. *Computer Graphics Forum* 39, 4 (July 2020). https://doi.org/10/hn27

Kurtis Robert Gurley. 1997. *Modelling and simulation of non-Gaussian processes.* Ph. D. Dissertation. University of Notre Dame.

John C. Hart. 1996. Sphere Tracing: A Geometric Method for the Antialiased Ray Tracing of Implicit Surfaces. *The Visual Computer* 12, 10 (Dec. 1996), 527–545. https://doi.org/10/b3q2p6

John C. Hart, Ed Bachta, Wojciech Jarosz, and Terry Fleury. 2002. Using Particles to Sample and Control More Complex Implicit Surfaces. In *Proceedings of Shape Modeling International.* 129–136. https://doi.org/10/dfw2ss

Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. 2015. The SGGX Microflake Distribution. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 34, 4 (July 2015), 48:1–48:11. https://doi.org/10/f7m2n2

Han-Ping Hong, M Y Xiao, Xizhong Cui, and Yongxu Liu. 2023. A framework for conditional simulation of nonstationary non-Gaussian random field and multivariate processes. *Mechanical Systems and Signal Processing* 183, 109646 (Jan. 2023). https://doi.org/10/mr46

Shuping Huang, ST Quek, and K. Phoon. 2001. Convergence study of the truncated Karhunen–Loeve expansion for simulation of stochastic processes. *International journal for numerical methods in engineering* 52, 9 (2001). https://doi.org/10/ffhsdb

Wenzel Jakob, Adam Arbree, Jonathan T. Moon, Kavita Bala, and Steve Marschner. 2010. A Radiative Transfer Framework for Rendering Materials with Anisotropic Structure. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 29, 4 (July 2010), 53:1–53:13. https://doi.org/10/ftr7st

Adrian Jarabo, Carlos Aliaga, and Diego Gutierrez. 2018. A Radiative Transfer Framework for Spatially-Correlated Materials. *ACM Transactions on Graphics (Proceedings*

of SIGGRAPH) 37, 4 (July 2018), 83:1–83:13. https://doi.org/10/gd52qq

Yang Jiao, F H Stillinger, and Salvatore Torquato. 2007. Modeling heterogeneous materials via two-point correlation functions: basic principles. *Physical Review E* 76 (Sept. 2007). https://doi.org/10/d6nrzc

Yang Jiao, F H Stillinger, and Salvatore Torquato. 2008. Modeling heterogeneous materials via two-point correlation functions. II. Algorithmic details and applications. *Physical Review E* 77 (March 2008). https://doi.org/10/bfzw5b

James T. Kajiya. 1986. The Rendering Equation. *Computer Graphics (Proceedings of SIGGRAPH)* 20, 4 (Aug. 1986), 143–150. https://doi.org/10/cvf53j

John T. O. Kirk. 1975. A theoretical analysis of the contribution of algal cells to the attenuation of light within natural waters I. General treatment of suspensions of pigmented cells. *New phytologist* 75, 1 (1975). https://doi.org/10/fdpx4w

Georg Lindgren, Krzysztof Podgorski, and Igor Rychlik. 2020. Effective computations of joint excursion times for stationary Gaussian processes. (July 2020). arXiv:2007.14220 [math.PR] http://arxiv.org/abs/2007.14220

Yu-Tao Liu, Li Wang, Jie Yang, Weikai Chen, Xiaoxu Meng, Bo Yang, and Lin Gao. 2023. NeUDF: Leaning Neural Unsigned Distance Fields with Volume Rendering. In *Computer Vision and Pattern Recognition (CVPR)*.

Binglin Lu and Salvatore Torquato. 1992. Lineal-path function for random heterogeneous materials. *Physical Review A* 45, 2 (Jan. 1992). https://doi.org/10/ddmhmc

Wolfram Martens, Yannick Poffet, Pablo Ramón Soria, Robert Fitch, and Salah Sukkarieh. 2017. Geometric Priors for Gaussian Process Implicit Surfaces. *IEEE Robotics and Automation Letters* 2, 2 (April 2017). https://doi.org/10/ghk78w

Media Molecule. 2020. *Dreams.* https://www.mediamolecule.com/games/dreams

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 Transactions on Graphics (Proceedings of SIGGRAPH)* 34, 4 (July 2015), 49:1–49:13. https://doi.org/10/gfzndr

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *Proceedings of the European Conference on Computer Vision (ECCV)*. https://doi.org/10/gn9vj9

Jonathan T. Moon, Bruce Walter, and Stephen R. Marschner. 2007. Rendering Discrete Random Media Using Precomputed Scattering Solutions. In *Rendering Techniques (Proceedings of the Eurographics Symposium on Rendering)*. Eurographics Association, 231–242. https://doi.org/10/gfzp5n

Joy N Mueller, Khachik Sargsyan, and Habib N Najm. 2023. Polynomial Chaos Surrogate Construction for Stochastic Models with Parametric Uncertainty. In *14th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP14), Dublin, Ireland.*

Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM Transactions on Graphics* 41, 4, Article 102 (July 2022), 15 pages. https://doi.org/10/gqkpt7

Thomas Müller, Marios Papas, Markus Gross, Wojciech Jarosz, and Jan Novák. 2016. Efficient Rendering of Heterogeneous Polydisperse Granular Media. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 35, 6 (Nov. 2016), 168:1–168:14. https://doi.org/10/f9cm65

Marcus M Noack and James A Sethian. 2022. Advanced stationary and nonstationary kernel designs for domain-aware gaussian processes. *Communications in Applied Mathematics and Computational Science* 17, 1 (2022). https://doi.org/10/mr4x

Marc Olano and Dan Baker. 2010. LEAN Mapping. In *Proceedings of the Symposium on Interactive 3D Graphics and Games*. 181–188. https://doi.org/10/fkbvpn

Christopher J Paciorek and Mark J Schervish. 2006. Spatial modelling using a new class of nonstationary covariance functions. *Environmetrics* 17, 5 (2006). https://doi.org/10/bt6vvk

Tobias Pfaffelmoser, Matthias Reitinger, and Rüdiger Westermann. 2011. Visualizing the positional and geometrical variability of isosurfaces in uncertain scalar fields. *Computer Graphics Forum* 30, 3 (June 2011). https://doi.org/10/d3pjgc

Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically Based Rendering: From Theory to Implementation* (3 ed.). Morgan Kaufmann, Cambridge, MA.

Gerald C Pomraning. 1991. *Linear kinetic theory and particle transport in stochastic mixtures.* Vol. 7. World Scientific. https://doi.org/10/mr49

Ali Rahimi and Benjamin Recht. 2007. Random Features for Large-Scale Kernel Machines. In *Advances in Neural Information Processing Systems (NIPS)*. Curran Associates Inc., Red Hook, NY, USA, 1177–1184.

Anthony P. Roberts and Salvatore Torquato. 1999. Chord-distribution functions of three-dimensional random media: approximate first-passage times of Gaussian processes. *Physical Review E* 59, 5 Pt A (May 1999). https://doi.org/10/fhhj8g

Rohan Sawhney and Keenan Crane. 2020. Monte Carlo Geometry Processing: A Grid-Free Approach to PDE-based Methods on Volumetric Domains. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 39, 4 (2020). https://doi.org/10/ghw7t7

Silvia Sellán and Alec Jacobson. 2022. Stochastic Poisson Surface Reconstruction. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 41, 6 (Dec. 2022). https://doi.org/10/mr5b

Silvia Sellán and Alec Jacobson. 2023. Neural Stochastic Screened Poisson Reconstruction. *ACM SIGGRAPH Asia Conference Papers* (2023). https://doi.org/10/mr5c

Michael D. Shields, George Deodatis, and Paolo Bocchini. 2011. A simple and efficient methodology to approximate a general non-Gaussian stationary stochastic process by a translation process. *Probabilistic Engineering Mechanics* 26, 4 (Oct. 2011). https://doi.org/10/cc5vkg

B. Smith. 1967. Geometrical Shadowing of a Random Rough Surface. *IEEE Transactions on Antennas and Propagation* 15, 5 (Sept. 1967), 668–671. https://doi.org/10/b754nw

Ralph C Smith. 2013. *Uncertainty quantification: theory, implementation, and applications.* Vol. 12. Siam.

Jorge Stol and Luiz Henrique De Figueiredo. 1997. Self-validated numerical methods and applications. In *Monograph for 21st Brazilian Mathematics Colloquium, IMPA, Rio de Janeiro. Citeseer*, Vol. 5. Citeseer.

Ryusuke Sugimoto, Terry Chen, Yiti Jiang, Christopher Batty, and Toshiya Hachisuka. 2023. A Practical Walk-on-Boundary Method for Boundary Value Problems. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 42, 4 (July 2023), 81:1–81:16. https://doi.org/10/gsk4f7

Towaki Takikawa, Andrew Glassner, and Morgan McGuire. 2022. A Dataset and Explorer for 3D Signed Distance Functions. *Journal of Computer Graphics Techniques (JCGT)* 11, 2 (27 April 2022), 1–29. http://jcgt.org/published/0011/02/01/

Michael Eugene Taylor. 2006. *Measure theory and integration.* American Mathematical Society.

Salvatore Torquato. 1986. Microstructure Characterization and Bulk Properties of Disordered Two-Phase Media. *Journal of Statistical Physics* 45, 5 (Dec. 1986), 843–873. https://doi.org/10/djpwzw

Salvatore Torquato. 2005. *Random Heterogeneous Materials: Microstructure and Macroscopic Properties.* Springer New York.

S. Torquato and B. Lu. 1993. Chord-Length Distribution Function for Two-Phase Random Media. *Physical Review E* 47, 4 (April 1993), 2950–2953. https://doi.org/10/dwggd5

V. Tuchin. 2000. *Tissue Optics: Light Scattering Methods and Instruments for Medical Diagnosis.*

Delio Vicini, Wenzel Jakob, and Anton Kaplanyan. 2021. A Non-Exponential Transmittance Model for Volumetric Scene Representations. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 40, 4 (July 2021), 136:1–136:16. https://doi.org/10/hshg

Emily H Vu and Aaron J Olson. 2021. Conditional Point Sampling: A Stochastic Media Transport Algorithm with Full Geometric Sampling Memory. *Journal of Quantitative Spectroscopy and Radiative Transfer* (2021). https://doi.org/10/gj9h3v

Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet Models for Refraction through Rough Surfaces. In *Rendering Techniques (Proceedings of the Eurographics Symposium on Rendering)*. Eurographics Association, 195–206. https://doi.org/10/gfz4kg

Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. 2021. NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction. *Advances in Neural Information Processing Systems (NIPS)* 34 (2021).

Philippe Weier, Tobias Zirr, Anton Kaplanyan, Ling-Qi Yan, and Philipp Slusallek. 2023. Neural Prefiltering for Correlation-Aware Levels of Detail. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 42, 4 (July 2023), 78:1–78:16. https://doi.org/10/gsk4f9

Christopher KI Williams and Carl Edward Rasmussen. 2006. *Gaussian processes for machine learning.* Vol. 2. MIT press Cambridge, MA.

Michael M. R. Williams. 1974. *Random processes in nuclear reactors.* Pergamon Press.

Oliver Williams and Andrew Fitzgibbon. 2006. Gaussian process implicit surfaces. In *Gaussian Processes in Practice.*

Andrew P. Witkin and Paul S. Heckbert. 1994. Using Particles to Sample and Control Implicit Surfaces. In *Annual Conference Series (Proceedings of SIGGRAPH)*. ACM, New York, NY, USA, 269–277. https://doi.org/10/bv24kc

George Zimmerman and Marvin L Adams. 1991. Algorithms for Monte Carlo particle transport in binary statistical mixtures. *Transactions of the American Nuclear Society* 64 (1991).