

Conditional Product Next Event Estimation for Gaussian Process Implicit Surfaces

Song Shi  Kehan Xu  Wojciech Jarosz 

Dartmouth College



Figure 1: Equal-sample comparison of the dragon scene with a heightfield-like GPIS. The previous NEE method [XBdEJ25] struggles under high anisotropy (top-left inset) since the set of admissible paths approaches a lower-dimensional manifold. Our CPNEE technique importance samples the 1D arc of valid scattering directions restricted to the light’s spherical cap, recovering a clean result in the same number of samples (top-right). Combining all three strategies (unidirectional, emitter NEE, and CPNEE) via multiple importance sampling (MIS) yields robust performance across arbitrary degrees of anisotropy (bottom-right).

Abstract

We propose a new next-event estimation (NEE) technique for Gaussian Process Implicit Surfaces (GPISes). We show that the distribution of surface normals available for sampling at a GPIS hit point collapses from 2D to 1D as the GPIS realizations approach heightfields, causing existing NEE to fail entirely in this limit. We analyze this collapse and develop a new strategy that directly importance samples the resulting 1D arc of valid scattering directions restricted to the spherical cap towards a light source. Our technique enables NEE on heightfield GPISes for the first time and provides substantial variance reduction for highly anisotropic near-heightfield configurations. We combine our strategy with existing techniques via multiple importance sampling for robust performance across all degrees of anisotropy, and support arbitrary emitter shapes via sphere proxies.

CCS Concepts

• **Computing methodologies** → **Rendering**; Ray tracing; • **Mathematics of computing** → Stochastic processes;

1. Introduction

Recent advances in light transport theory have introduced Gaussian Process Implicit Surfaces (GPISes) [SdEJ24] as a powerful, unified representation for rendering surfaces, volumes, and a rich continuum in between. This approach fundamentally differs from classical participating media or microfacet models, which efficiently approximate light transport by sampling statistical distributions without ever explicitly constructing the underlying geometric microstructure. Instead, the GPIS framework operates by explicitly instantiating and ray tracing through random realizations of the stochastic geometry.

By statistically averaging light transport over these explicit intersections, GPISes naturally capture complex multiple scattering effects and spatial correlations along the path – phenomena that classical models often fail to reproduce. Furthermore, this approach unlocks a new appearance continuum that smoothly interpolates between surfaces and volumes. However, this physical fidelity comes at a high computational cost. Despite recent progress in making GPIS rendering more practical [XBdEJ25], explicitly intersecting random realizations remains significantly more expensive than evaluating specialized classical models, limiting its broader adoption.




Figure 2: As anisotropy increases from left to right, GPIS realizations approach Beckmann heightfields (with fewer overhangs). The conditional visible gradient distribution (cvGDF) at a ray marching intersection p_s^f gets narrower in the XZ plane (Y pointing out of the page). We visualize its top-down hemispherical projection and the corresponding distribution of scattering directions $p(\omega_s | \omega_o, g_w)$ in each inset (red: high density). Orange dots show directions sampled towards the emitter; with increasing anisotropy, these increasingly miss the high-contribution regions. In the near-heightfield limit (right), the cvGDF collapses to a 1D arc and emitter NEE fails entirely.

A central challenge is the efficient estimation of light transport under dense stochastic geometry. In particular, next-event estimation (NEE), a cornerstone of physically based rendering, requires careful adaptation in the GPIS setting. Xu et al. [XBdEJ25] introduced the first NEE method for specular GPISes, sampling from the distribution of surface gradients available at each hit point – the *conditional visible gradient distribution function* (cvGDF), which we define formally in Sec. 3.3. As illustrated in Fig. 2, however, this approach degrades under strong anisotropy: as the GPIS realizations approach heightfields, the cvGDF collapses from 2D to 1D and NEE fails entirely. Heightfield-like GPISes – which correspond to rough surface-like appearance handled traditionally by microfacet models – thus remain poorly supported by existing NEE strategies.

In this paper, we introduce a new next-event estimation method that directly addresses this limitation, robustly handling heightfield GPISes while providing substantial variance reduction for near-heightfield configurations. Specifically, our core contributions are:

- We identify and analyze the failure of existing NEE strategies on heightfield-like GPISes, showing how the cvGDF collapses from 2D to 1D in the heightfield limit.
- We propose a novel NEE formulation that directly resolves this singularity by explicitly sampling the 1D scattering arcs of admissible outgoing directions that overlap with a spherical light source (or a bounding sphere proxy).
- We demonstrate how to robustly combine our new estimator with other sampling strategies via multiple importance sampling (MIS) [VG95], providing a unified framework that smoothly bridges the gap between heightfield-like and volume-like GPISes.

2. Related work

Gaussian Process. Gaussian Processes (GPs) [RW06] have recently emerged as a powerful framework for modeling stochastic geometry. Seyb et al. [SdEBJ24] established the theoretical foundations for light transport on Gaussian Process Implicit Surfaces (GPISes) and made their rendering feasible, though the resulting formulation remained far from practical due to its prohibitive $O(N^3)$ cost, where N is the number of ray march steps. Among other improvements, Xu et al. [XBdEJ25] reduced ray-marching complexity from cubic to linear. We defer the technical details of these foundational models to Sec. 3.

Next Event Estimation. NEE [AK90; SWZ96] is a fundamental variance reduction technique in Monte Carlo path tracing. Xu et al. [XBdEJ25] further introduced the first NEE technique for specular GPISes – sampling emitters and weighting the contribution by the cvGDF at each hit point. Unfortunately, as the GPIS covariance becomes anisotropic, the cvGDF narrows (Fig. 2, middle), degrading the benefits of their NEE strategy until it fails entirely in the limit where the cvGDF become a manifold arc (Fig. 2, right). Manifold NEE (MNEE) [HDF15] addresses singularities arising from deterministic specular constraints. Our technique samples the arc-like manifold cvGDFs that arise in GPISes, further restricted to the spherical cap towards an emitter, meaning it samples the *product* of the cvGDF and the emitter. Sampling the arc-like manifolds is conceptually related to line-sampling NEE techniques [BD16; SJ19].

Microfacet Theory. The strongly anisotropic covariance structures that cause NEE to fail include GPIS realizations approaching random heightfields. Such heightfields are the geometric foundation of classical microfacet models [BS63; CT81; dEBWZ23; ON94; WMLT07], the standard tool for rough-surface appearance in both real-time and offline rendering. These models are far more efficient than the GPIS framework because they operate on statistical distributions and avoid intersecting explicit microgeometry. This efficiency, however, comes at the cost of approximations that each require explicit work to correct – for instance, energy loss from the single-scattering assumption [HHdED16]. Concurrent unpublished work [HCWY26] has considered elevating microfacet theory to a volumetric representation [JAM*10] for improved GPIS rendering speed, though again at the expense of introducing further approximations.

GPISes make fewer geometric assumptions and naturally capture multiple scattering, path-level correlations, and non-heightfield microstructures such as porous materials [LRPB23], but at significant efficiency cost. Our work dramatically reduces noise (and hence cost) for the strongly anisotropic covariance regime including heightfields, while gracefully generalizing to arbitrary GPISes via MIS [VG95].

3. Background

3.1. Gaussian Process Implicit Surfaces

A Gaussian process $f(\mathbf{p}) \sim \text{GP}(\mu(\mathbf{p}), \kappa(\mathbf{p}, \mathbf{p}'))$ is a distribution over functions f characterized by the mean $\mu(\mathbf{p}) = \mathbb{E}_f[f(\mathbf{p})]$

and the covariance function $\kappa(\mathbf{p}, \mathbf{p}')$, which encodes how $f(\mathbf{p})$ and $f(\mathbf{p}')$ co-vary for any pair of locations $\mathbf{p}, \mathbf{p}'$. A Gaussian Process Implicit Surface (GPIS) is simply the zero level set of a Gaussian process. We can view a realization f as the sum of a deterministic mean implicit function $\mu(\mathbf{p})$ and a stochastic function $\psi(\mathbf{p}) \sim \mathbf{GP}(0, \kappa(\mathbf{p}, \mathbf{p}'))$ drawn from the GP with the mean removed: $f(\mathbf{p}) = \mu(\mathbf{p}) + \psi(\mathbf{p})$. As proposed by Xu et al. [XBdEJ25], we efficiently evaluate the stochastic component $\psi(\mathbf{p})$ at any point via sparse convolution noise [Lew84; LLDD09], and can obtain different random realizations by changing the random seed used by the noise. With this perspective, tracing a GPIS realization can be simply viewed as tracing a surface displaced by procedural noise.

GPISes can represent surface-type, volume-type, and in-between geometries. This is achieved through careful design of the mean function μ and the covariance kernel κ . For simplicity, we focus on the squared exponential covariance kernel, $\kappa(\mathbf{p}, \mathbf{p}') = \sigma^2 \exp(-\frac{1}{2}(\mathbf{p} - \mathbf{p}')^\top \mathbf{M}^2 (\mathbf{p} - \mathbf{p}'))$ where σ controls the noise amplitude and the 3×3 symmetric positive definite matrix $\mathbf{M}^2 := \mathbf{M}^\top \mathbf{M}$ determines the potentially anisotropic length scales of the kernel. In practice, we can achieve such anisotropic covariance by simply stretching or squishing evaluations of a canonical isotropic procedural noise function $\psi(\mathbf{p}) := \psi_{\text{iso}}(\mathbf{M}\mathbf{p})$, where ψ_{iso} has isotropic unit length scale: $\kappa_{\text{iso}}(\mathbf{p}, \mathbf{p}') = \sigma^2 \exp(-\frac{1}{2}\|\mathbf{p} - \mathbf{p}'\|^2)$. Such stretching shifts GPIS realizations from volumetric shapes with overhangs toward heightfields. For an axis-aligned kernel, $\mathbf{M} = \text{diag}(\ell_x^{-1}, \ell_y^{-1}, \ell_z^{-1})$, where $\ell_x, \ell_y, \ell_z > 0$ are the axis-wise correlation lengths. As $\ell_z \rightarrow \infty$, the kernel becomes independent of z , reducing the covariance to a 2D form $\kappa(\mathbf{p}, \mathbf{p}') = \kappa(\mathbf{p}_{xy}, \mathbf{p}'_{xy})$. This creates perfect vertical correlation and, with a mean gradient in z , forces every GPIS realization to be a heightfield in z (Fig. 2).

3.2. Light transport on GPISes

Within each realization, the light transport is given by the classical rendering equation [ICG86; Kaj86]

$$L^f(\mathbf{p}, \omega) = \int_{S^2} \rho(\mathbf{p}_s^f) L^f(\mathbf{p}_s^f, \omega_s) d\omega_s, \quad (1)$$

where $L^f(\mathbf{p}, \omega)$ is the radiance arriving at point $\mathbf{p}$ from direction ω , given in terms of the radiance leaving the closest intersection $\mathbf{p}_s^f := \mathbf{p} + s\omega$ of the ray with the surface induced by f . This is computed as the spherical integral of light arriving at $\mathbf{p}_s^f$ and being reflected into direction ω , the latter being given by the cosine-weighted BRDF $\rho(\mathbf{p}_s^f) := \rho(\mathbf{p}_s^f, \omega, \omega_s, \mathbf{n}_s^f) |\mathbf{n}_s^f \cdot \omega_s|$. We use a superscript of f to denote quantities that depend on the realization (i.e. the random seed of the procedural noise). To simplify the exposition, we omit the emitted radiance L_e ; its inclusion is straightforward.

While L^f denotes transport in a specific realization f , the total received radiance is the ensemble average over all realizations:

$$\langle L^f(\mathbf{p}, \omega) \rangle = \int_{\text{GP}} L^f(\mathbf{p}, \omega) d\gamma(f), \quad (2)$$

where $\gamma(f)$ denotes the classic Wiener measure – the probability density of sampling $f \sim \text{GP}$. In practice, we can approximate Eq. (2) by ray marching the mean function μ perturbed by the procedural noise ψ and average the light transport over all possible random seeds for ψ . For specular materials, fixing the full 3D stochastic field with a

single random seed locks down the surface entirely: the normal at any hit point is fully determined, leaving nothing to importance sample. Xu et al. [XBdEJ25] resolve this by evaluating only a 1D slice of ψ along each ray. Ray marching this slice yields the intersection and the derivative of ψ along the ray, while leaving gradient components perpendicular to the ray unobserved and available for sampling. These remaining degrees of freedom are what enable importance sampling like NEE. We adopt the Renewal Half+ memory model of Xu et al. [XBdEJ25], which ensures these free components at each hit point are independent of the preceding path segment. As we show in Sec. 3.3, isotropic GPISes retain two such degrees of freedom; under strong anisotropy the distribution compresses, and in the heightfield limit one degree of freedom vanishes entirely. This is the failure mode our method addresses.

3.3. Existing sampling strategies for GPISes

In the **isotropic case**, the stochastic component of the GPIS is ψ_{iso} . The distribution of gradients $\nabla\psi_{\text{iso}}$ is an isotropic 3D Gaussian. Note that an isotropic Gaussian is independent in any orthonormal coordinate system. We choose an orthonormal coordinate system $[\mathbf{u} \ \mathbf{v} \ \mathbf{w}]$ where $\mathbf{w} := \omega_o$ and $\mathbf{u}, \mathbf{v}$ are two additional vectors that complete the basis ($\mathbf{u}$ points out of the screen). A 3D world-space gradient is

$$\mathbf{g} = \begin{bmatrix} | & | & | \\ \mathbf{u} & \mathbf{v} & \mathbf{w} \\ | & | & | \end{bmatrix} \mathbf{g}_{uvw} := \mathbf{u}g_u + \mathbf{v}g_v + \mathbf{w}g_w, \quad (3)$$

where each of g_u, g_v, g_w are distributed *independently* according to 1D Gaussian distributions.

Marching along the ray gives us a hit point $\mathbf{p}_s^f$ along with $D_{\mathbf{w}}\psi_{\text{iso}}(\mathbf{p}_s^f)$, where $D_{\mathbf{w}}$ denotes the directional derivative operator along the direction $\mathbf{w}$. Evaluating the derivative of a 1D noise realization can be interpreted as sampling one component $g_w = D_{\mathbf{w}}\psi_{\text{iso}}$ of the 3D gradient in the $\mathbf{uvw}$ coordinate system. This leaves two additional components $\mathbf{g}_{uv} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, which follow an isotropic 2D Gaussian (g_v pointing along $\mathbf{v}$ and g_u pointing out of the screen along $\mathbf{u}$). Together, these three components form a full 3D gradient $\mathbf{g}$ via Eq. (3). Adding the mean gradient $\nabla\mu$ gives:

$$\nabla f(\mathbf{g}_{uv} | \omega_o, g_w) = \underbrace{\mathbf{w}g_w + \nabla\mu}_{\text{fixed by ray intersection}} + \underbrace{\mathbf{u}g_u + \mathbf{v}g_v}_{\text{free}}. \quad (4)$$

Geometrically, the distribution of these gradients, $p_{\nabla f}(\mathbf{g}_{uv} | \omega_o, g_w)$, therefore consists of the vectors from the shade point $\mathbf{p}_s^f$ to an isotropic 2D Gaussian $\mathbf{g}_{uv}$ on a plane (dashed black line at the end of $\mathbf{w}g_w$) perpendicular to the ray, shifted by the mean gradient $\nabla\mu$. We call $p_{\nabla f}(\mathbf{g}_{uv} | \omega_o, g_w)$ the cvGDF and visualized its hemispherical projection in Fig. 2. Formally, this is the 2D family of gradients visible from ω_o *conditioned* on the specific g_w observed during ray marching. It is the distribution Xu et al. [XBdEJ25] implicitly rely on for NEE. Notably, this differs from the *visible NDF* (vNDF) of microfacet theory [HdE14] and visualized by Seyb et al. [SdEJ24]




Figure 3: The cvGDF's projections onto the upper hemisphere conditioned on a fixed ray direction ω_o and different g_w values (left 3 columns). The classical vNDF (right column) corresponds to the weighted average across all $g_w \sim p_{rm}$ generated by ray marching. In the heightfield limit, the vNDF of GPIS (upper half) matches the Beckmann vNDF with Smith approximation (lower half).

for GPISeS, which marginalizes over all values of g_w . Fig. 3 illustrates this relationship.

The ensemble-averaged light transport integrates over all possible visible gradient directions, so the outgoing radiance at a hit point $\mathbf{p}_s^f$ back along the ray-marching direction ω_o is


$$\langle L_o^f(\mathbf{p}_s^f, \omega_o) \rangle = \int_{\mathbb{R}} p(g_w) \left(\int_{\mathbb{R}^2} \rho(\mathbf{p}_s^f) \langle L^f(\mathbf{p}_s^f, \omega_s) \rangle p(\mathbf{g}_{uv} | \omega_o, g_w) d\mathbf{g}_{uv} \right) dg_w, \quad (5)$$

where the surface normal for the BRDF ρ at the shade point is the normalized gradient $\nabla f / \|\nabla f\|$, and $p(\mathbf{g}_{uv} | \omega_o, g_w) = \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ is the cvGDF with the mean removed. Ray marching importance samples g_w so $p_{rm}(g_w) = p(g_w)$ cancel, leaving the Monte Carlo estimator:

$$\widehat{L_o^f}(\mathbf{p}_s^f, \omega_o) = \frac{p(g_w) \rho(\mathbf{p}_s^f) \langle L^f(\mathbf{p}_s^f, \omega_s) \rangle p(\mathbf{g}_{uv} | \omega_o, g_w)}{p_{rm}(g_w) p_{mc}(\mathbf{g}_{uv})}. \quad (6)$$

Prior work [XBdEJ25] proposed two strategies for sampling the remaining terms: unidirectional (which sets p_{mc} to cancel $p(\mathbf{g}_{uv} | \omega_o, g_w)$) and NEE sampling (which samples directions towards the light sources and evaluates $p(\mathbf{g}_{uv} | \omega_o, g_w)$ in the numerator). We refer to these as the UNI and emitter NEE strategies.

If the stochastic component is anisotropic (e.g., stretched along z as in the inset figure), we can write it as $\psi(\mathbf{p}) = \psi_{iso}(\mathbf{M}\mathbf{p})$, where $\mathbf{M}$ is a 3×3 matrix that transforms from world space to isotropic space (equivalently, we can think of ψ as ψ_{iso} trans-



formed by $\mathbf{M}^{-1}$). Due to the covector transformation rule, isotropic-space gradients get transformed by $\mathbf{M}^T$ (the inverse transpose of $\mathbf{M}^{-1}$), hence $\nabla\psi = \mathbf{M}^T \nabla\psi_{iso}$.

To reason about the gradient $\mathbf{g}$, we will again consider a ray-aligned orthonormal basis $\mathbf{uvw}$ in isotropic space, where now $\mathbf{w} = \mathbf{M}\omega_o / \|\mathbf{M}\omega_o\|$. Gradients $\mathbf{g}_{uvw}$ in this space remain distributed according to an isotropic 3D Gaussian, however, since $\mathbf{M}$ is not orthonormal, the world-space gradient becomes:

$$\mathbf{g} = \begin{bmatrix} \hat{\mathbf{u}} & \hat{\mathbf{v}} & \hat{\mathbf{w}} \end{bmatrix} \mathbf{g}_{uvw} := \hat{\mathbf{u}}g_u + \hat{\mathbf{v}}g_v + \hat{\mathbf{w}}g_w, \quad (7)$$

where $\hat{\mathbf{u}}\hat{\mathbf{v}}\hat{\mathbf{w}}$ are the (skewed) dual basis vectors formed by transforming $\mathbf{uvw}$ into world space: $\begin{bmatrix} \hat{\mathbf{u}} & \hat{\mathbf{v}} & \hat{\mathbf{w}} \end{bmatrix} = \mathbf{M}^T \begin{bmatrix} \mathbf{u} & \mathbf{v} & \mathbf{w} \end{bmatrix}$.

As before, at the hit point of a realization, the directional derivative $g_w = D_{\mathbf{w}}\psi_{iso}$ samples one component of the gradient, and the remaining two components $\mathbf{g}_{uv} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ are distributed according to a zero-mean isotropic 2D Gaussian. However, when this isotropic Gaussian is transformed into world space via Eq. (7), we get a squished 2D distribution along the $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ axes. Note that $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ remain perpendicular to the ray direction ω_o , but $\hat{\mathbf{w}}$ is no longer aligned with the ray direction.

As the anisotropy along the z -axis increases toward infinity

(see inset), all realizations of the GPIS approach standard Beckmann heightfields. The 2D Gaussian distribution $\mathbf{g}_{uv}$ becomes increasingly compressed, eventually degenerating into a 1D Gaussian. Its hemispherical projection therefore collapses to an arc on the hemisphere, as illustrated in Fig. 2. Under this degenerate configuration, NEE fails completely.

4. Method

We address the failure of the existing NEE technique with *conditional product next-event estimation (CPNEE)*, which explicitly samples the 1D arc of valid scattering directions that intersect the light source (Secs. 4.1 and 4.2). We derive CPNEE for the degenerate anisotropic case before extend it to general 3D GPISeS (Sec. 4.3). We then discuss how to combine it with existing strategies via MIS (Sec. 4.4), handle non-spherical emitters via bounding-sphere proxies (Sec. 4.5), and consider non-stationary cases in Sec. 4.6.

4.1. Double-circle Geometry Model

In the heightfield GPIS case, after intersecting the GPIS at $\mathbf{p}_s^f$ and obtaining the directional derivative g_w along the ray, the gradient $\nabla f(\mathbf{p}_s^f)$ has one remaining degree of freedom, $g_u \sim \mathcal{N}(0, \sigma)$, along the $\hat{\mathbf{u}}$ axis (see Fig. 4 left). The full gradient is centered at the mean gradient plus the ray directional component, i.e., $\nabla\mu(\mathbf{p}_s^f) + g_w\hat{\mathbf{w}}$. Projecting this gradient distribution onto the sphere centered at $\mathbf{p}_s^f$ yields a semicircular arc of admissible normals $\mathbf{n}$. Reflecting ω_o about these normals produces a full circle of possible scattering directions ω_s (the yellow-purple circle in Fig. 4 right). We refer to this circle as the *scattering circle*.





Figure 4: Geometric setup of the double-circle model for height-field GPISes. After normalizing the gradient, the set of admissible normals forms a semicircular arc on the upper hemisphere (left), and the associated scattered directions constitute a circle – the scattering circle (right). The spherical light source also projects onto the sphere as a circle – the light cap (right). Note that both circles may span the entire sphere.

A spherical light source projects onto the same sphere centered at p_s^f as another circle (the orange circle and cone in Fig. 4 right). The spherical *light cap* within the circle represents the range of outgoing directions that intersect the light. To obtain the set of scattering directions that can both be generated by the GPIS and hit the light, we first determine the portion of the scattering circle that lies inside the light cap.



This is a purely geometric problem. Once this range is known, we describe how to importance sample the scattering directions in Sec. 4.2. To provide intuition, we show in the left inset the distribution of scattering directions at a GPIS intersection in a rendered scene produced by our method.

There are three distinct cases: (1) the scattering circle lies entirely outside the light cap; (2) the scattering circle is completely contained within the light cap (Fig. 5a); and (3) the two circles intersect at two points ω_{s1} and ω_{s2} , which form the bounds of a scattering arc toward the light (Fig. 5b). We need to distinguish among these three situations, and in the third case determine the arc bounds ω_{s1}, ω_{s2} .

Our method proceeds by expressing each circle as the intersection of a plane with a unit sphere. We denote by $\mathbf{c}_l$ and $\mathbf{c}_s$ (Fig. 5) the *Euclidean* centers of the light and scattering circles. This allows us to write implicit equations for two planes as $\mathcal{P}_l : \mathbf{c}_l \cdot \mathbf{o} = \mathbf{c}_l \cdot \mathbf{c}_l$ and $\mathcal{P}_s : \mathbf{c}_s \cdot \mathbf{o} = \mathbf{c}_s \cdot \mathbf{c}_s$, where $\mathbf{c}_l$ and $\mathbf{c}_s$ are used for both the offset and normals of their respective planes. Any point $\mathbf{o}$ satisfying an equation lies on the associated plane.

Now consider the line defined by the intersection of the two planes $\mathcal{P}_l$ and $\mathcal{P}_s$. Any point $\mathbf{o}$ on this line satisfies both plane equations. This line is orthogonal to both $\mathbf{c}_l$ and $\mathbf{c}_s$, and its direction is therefore given by $\mathbf{d} = \mathbf{c}_l \times \mathbf{c}_s$. We focus on the minimum-norm point $\mathbf{o}^*$, i.e., the point on the line closest to the origin. To solve for $\mathbf{o}^*$, we construct




Figure 5: (a): The case where the scattering circle is fully contained within the light cap. (b): the case where the admissible scattered directions form a minor arc (blue) on the scattering circle $\mathbf{c}_s$. (c): the major arc case. The red vector determines which arc is contained in the light cap.

a 3×3 linear system, where the first two equations enforce the plane constraints, and the third enforces that $\mathbf{o}^*$ is orthogonal to $\mathbf{d}$:

$$\begin{bmatrix} - & \mathbf{c}_s^T & - \\ - & \mathbf{c}_l^T & - \\ - & \mathbf{d}^T & - \end{bmatrix} \mathbf{o}^* = \begin{bmatrix} \mathbf{c}_s \cdot \mathbf{c}_s \\ \mathbf{c}_l \cdot \mathbf{c}_l \\ 0 \end{bmatrix} \quad (8)$$

We can write out the analytic solution to this linear system:

$$\mathbf{o}^* = \frac{\|\mathbf{c}_s\|^2 (\mathbf{d} \times \mathbf{c}_l) - \|\mathbf{c}_l\|^2 (\mathbf{d} \times \mathbf{c}_s)}{\|\mathbf{d}\|^2}. \quad (9)$$


If $\|\mathbf{o}^*\|^2 > 1$, the two circles do not intersect, and the scattering circle is either (1) completely outside or (2) completely inside the light cap. For (2), two conditions must hold (Fig. 5a): first, the radius of the light cap must exceed the radius of the scattering circle, i.e. $\|\mathbf{c}_l\|^2 > \|\mathbf{c}_s\|^2$; second, the angle β between $\mathbf{c}_l$ and $\mathbf{c}_s$ must be smaller than the difference of the half-angles subtended by the light and scattering circles, θ_l and θ_s .

If $\|\mathbf{o}^*\|^2 < 1$, the two circles intersect at two points, ω_{s1} and ω_{s2} (Fig. 5b). These intersections are obtained by intersecting the unit sphere with the line $\mathbf{o}^* + t\mathbf{d}$. The two points define the start and end of the scattering arc. In this case we must decide whether the arc between ω_{s1} and ω_{s2} is the minor or the major one. This can be determined by checking the sign of the dot product $((\omega_{s1} - \mathbf{c}_s) + (\omega_{s2} - \mathbf{c}_s)) \cdot \mathbf{c}_l$, which measures the angle between $\mathbf{c}_l$ and the midpoint of the two arc endpoints. If this dot product is positive, the arc is the minor one (Fig. 5b); otherwise it is the major arc (Fig. 5c).

Given ω_{s1} and ω_{s2} , we recover the corresponding range of the gradient component along $\hat{\mathbf{u}}$, denoted g_{u1} and g_{u2} . Both ω_{s1} and ω_{s2} map back to their associated normals via $\mathbf{n} = \frac{\omega_o + \omega_s}{\|\omega_o + \omega_s\|}$. Intersecting each $\mathbf{n}$ with the line representing the one-dimensional Gaussian distribution (Fig. 4, top left) yields the value of g_u and the magnitude $\|\nabla f(\mathbf{p}_s^f)\|$. An intersection is guaranteed because each normal originates from an outgoing direction on the scattering circle.

4.2. Sample the 1D Gradient

If the scattering circle is fully contained by the light cap (Fig. 5a), we sample g_u according to its 1D distribution, achieving perfect importance sampling (except for visibility).



In the case where the two circles intersect (Fig. 5b), we obtain the 1-D gradient values g_{u_1}, g_{u_2} that correspond to the two boundary points $\omega_{s_1}, \omega_{s_2}$ described in Sec. 4.1. The next step is to sample g_u from a truncated Gaussian distribution within the interval defined by g_{u_1} and g_{u_2} . There are two possibilities: the admissible range may correspond to the interior of the Gaussian, or to the two tails on either side. Returning to the geometric setup in Fig. 4, we observe that the two infinity points map to one same location on the scattering circle. We have already determined in Sec. 4.1 which arc (major or minor) lies inside the light cap. If this infinity point lies within that arc (inset, blue segment), we sample the two tails of the Gaussian; otherwise we sample the interior (inset, red segment).

Returning to the geometric setup in Fig. 4, we observe that the two infinity points map to one same location on the scattering circle. We have already determined in Sec. 4.1 which arc (major or minor) lies inside the light cap. If this infinity point lies within that arc (inset, blue segment), we sample the two tails of the Gaussian; otherwise we sample the interior (inset, red segment).

The NEE strategy for a GPIS heightfield yields a Monte Carlo estimator that closely resembles Eq. (6):

$$\widehat{L}_o(\mathbf{p}_s^f, \omega_o) = \frac{p(g_w)p(\mathbf{p}_s^f)\langle L^f(\mathbf{p}_s^f, \omega_s) \rangle p(g_u | \omega_o, g_w)}{p_{\text{rm}}(g_w)p_{\text{mc}}(g_u)}. \quad (10)$$


The only difference is that, instead of a 2D gradient distribution $\mathbf{g}_{uv}$, the degenerate case involves only one degree of freedom in g_u , in addition to the directional derivative g_w that is already determined at the ray intersection. Here, $p_{\text{mc}}(g_u)$ denotes the Monte Carlo sampling density for g_u , which follows a 1D (truncated) Gaussian distribution as discussed above.

Our method samples g_u from the product of the cvGDF $p(g_u | \omega_o, g_w)$ and the light source contribution, hence the name *conditional product next event estimation*.

4.3. Extending to General 3D GPISes

After conditioning on the directional derivative g_w , general 3D GPISes with finite correlation still have two degrees of freedom (g_u and g_v), which together define a 2D gradient plane (dashed rectangle in the inset figure). The plane center is defined by the mean gradient plus the directional derivative along the ray, $\nabla\mu(\mathbf{p}_s^f) + g_w\hat{\mathbf{w}}$. We first sample $g_v \sim \mathcal{N}(0, \sigma)$ from a 1D Gaussian distribution, and $\nabla\mu(\mathbf{p}_s^f) + g_w\hat{\mathbf{w}} + g_v\hat{\mathbf{v}}$ determines the conditional mean of g_u . We then follow the same steps as in Secs. 4.1 and 4.2 to construct the normal arc and scattering circle, and finally intersect it with the light cap.

For general GPISes, our Monte Carlo estimator is almost identical



to the degenerate case in Eq. (10) but with a $p(g_v | \omega_o, g_w)$ term (written as $p(g_v)$ for brevity) that we importance sample via $p_{\text{mc}}(g_v)$:

$$\widehat{L}_o(\mathbf{p}_s^f, \omega_o) = \frac{p(g_w)p(g_v)p(\mathbf{p}_s^f)\langle L^f(\mathbf{p}_s^f, \omega_s) \rangle p(g_u | \omega_o, g_w)}{p_{\text{rm}}(g_w)p_{\text{mc}}(g_v)p_{\text{mc}}(g_u)}. \quad (11)$$

4.3.1. The Choice of the $\mathbf{u}$ and $\mathbf{v}$ Axes

In the degenerate case, $\mathbf{u}$ in the isotropic space is uniquely determined as the cross product between the ray direction $\mathbf{w}$ and the direction of infinite correlation. In contrast, for general 3D GPISes, we sample a 2D distribution g_{uv} in the isotropic space. Together, $\mathbf{u}$ and $\mathbf{v}$ form an orthonormal basis perpendicular to $\mathbf{w}$. This basis is not unique, introducing an additional degree of freedom in the choice of axes. Because g_u and g_v are sampled using different routines, different choices of $(\mathbf{u}, \mathbf{v})$ lead to different estimators and variance characteristics. In our implementation, we choose $\mathbf{u}$ as the cross product between $\mathbf{w}$ and the direction of anisotropy, a natural extension of the heightfield case. The vector $\mathbf{v}$ is then given by $\mathbf{u} \times \mathbf{w}$. We further discuss the choice of axes in Sec. 6.

4.4. Multiple Importance Sampling

For 3D GPISes, a key source of variance in CPNEE arises from the sampling of g_v . Specifically, we first sample g_v , which – together with g_w and the mean gradient $\nabla\mu(\mathbf{p}_s^f)$ – determines the center of the conditional 1D Gaussian for g_u , and consequently the normal arc and scattering circle. Depending on the sampled g_v , the resulting scattering circle may fail to intersect the light cap. In contrast, standard NEE directly samples directions towards an emitter, but these may miss concentrated high-contribution regions of the scattering distribution. We therefore combine all three strategies via MIS.

To apply the balance heuristic, we first convert the PDFs of all methods into the solid angle measure of the scattering direction ω_s . For CPNEE, we begin with $p(g_u | \omega_o, g_w)$ expressed in the local $\hat{\mathbf{u}}\hat{\mathbf{v}}\hat{\mathbf{w}}$ frame. The Jacobian factor required for the measure conversion is:

$$J_{\hat{\mathbf{u}}\hat{\mathbf{v}}\hat{\mathbf{w}} \rightarrow \omega_s} = \frac{1}{\|\hat{\mathbf{u}} \times \hat{\mathbf{v}}\|} \frac{\|\nabla f\|^2}{\|\omega_s \cdot \frac{(\hat{\mathbf{u}} \times \hat{\mathbf{v}})}{\|(\hat{\mathbf{u}} \times \hat{\mathbf{v}})\|}\|} \frac{1}{4|\omega_o \cdot \omega_s|}. \quad (12)$$

The first term corresponds to the change in density induced by transforming from the isotropic space to world space via $\mathbf{M}^T$. The second term accounts for the conversion between the area measure on the gradient plane and the solid angle measure of the normal $\mathbf{n}$. The third term converts between the normal $\mathbf{n}$ and the scattering direction ω_s .

4.5. Other Light Source Shapes

Our method can handle arbitrarily shaped lights by enclosing the light source within a bounding sphere and then applying the same procedure to sample scattering directions that intersect this sphere (Fig. 6). In this case, the shadow ray is not guaranteed to intersect the actual light, so if it misses, we simply return zero contribution. Naturally, the coverage ratio of the bounding sphere directly affects the efficiency of our algorithm.

Algorithm 1 MIS (UNI, CPNEE)

```

1: function  $L(\mathbf{p}, \omega, \nabla f, \text{depth}, \text{seed})$ 
2:    $L_o \leftarrow L_e(\mathbf{p}_s^f)$ 
3:    $\mathbf{p}_s^f, \mathbf{g}_w, \nabla \mu \leftarrow \text{RAYMARCHGPIS}(\mathbf{p}, \omega, \nabla f)$ 
4:    $\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{w}} \leftarrow \text{COMPUTEUVW}(\omega)$ 
5:    $\mathbf{g}_v, \mathbf{p}_v \leftarrow \text{SAMPLESTDGAUSSIAN}()$ 
6:    $\nabla f_{\text{fix}} \leftarrow \mathbf{g}_v \hat{\mathbf{v}} + \mathbf{g}_w \hat{\mathbf{w}} + \nabla \mu$ 
7:    $\triangleright$  Perform CPNEE sampling
8:    $\mathbf{c}_l \leftarrow \text{COMPUTELIGHTCAP}(\mathbf{p}_s^f)$ 
9:    $\mathbf{c}_s \leftarrow \text{COMPUTESCATTERINGCIRCLE}(\mathbf{p}_s^f, \nabla f_{\text{fix}}, \omega, \hat{\mathbf{u}})$ 
10:  hasIssect, isMajorArc,  $\mathbf{g}_{u1}, \mathbf{g}_{u2} \leftarrow \text{INTERSECTCIRCLES}(\mathbf{c}_l, \mathbf{c}_s)$ 
11:  if hasIssect then
12:     $\mathbf{g}_u^{\text{cpnee}}, \mathbf{p}_u^{\text{cpnee}} \leftarrow$ 
13:       $\text{SAMPLETRUNCSTDGAUSS}(\mathbf{g}_{u1}, \mathbf{g}_{u2}, \text{isMajorArc})$ 
14:     $\mathbf{p}_{uv}^{\text{cpnee}} \leftarrow \mathbf{p}_u^{\text{cpnee}} \mathbf{p}_v$ 
15:     $\mathbf{n}^{\text{cpnee}} \leftarrow \text{NORMALIZE}(\nabla f_{\text{fix}} + \mathbf{g}_u^{\text{cpnee}} \hat{\mathbf{u}})$ 
16:     $\omega_s^{\text{cpnee}} \leftarrow \text{REFLECT}(\mathbf{n}^{\text{cpnee}}, \omega)$ 
17:     $L_e^{\text{cpnee}} \leftarrow \text{EVALEMISSION}(\mathbf{p}_s^f, \omega_s^{\text{cpnee}})$ 
18:     $\mathbf{p}_{uv}^{\text{uni}} \leftarrow \text{EVALSTDGAUSSIAN}(\mathbf{g}_u^{\text{cpnee}} \mathbf{p}_v)$ 
19:     $\rho^{\text{cpnee}} \leftarrow \text{EVALBSDFCOSINE}(\mathbf{p}_s^f, \omega, \omega_s^{\text{cpnee}}, \mathbf{n}^{\text{cpnee}})$ 
20:     $\mathbf{w}^{\text{cpnee}} \leftarrow \text{MISWEIGHT}(\mathbf{p}_{uv}^{\text{cpnee}}, \mathbf{p}_{uv}^{\text{uni}})$ 
21:     $L_o \leftarrow L_o + \mathbf{w}^{\text{cpnee}} \frac{\mathbf{p}_{uv}^{\text{uni}}}{\mathbf{p}_{uv}^{\text{cpnee}}} \rho^{\text{cpnee}} L_e^{\text{cpnee}}$ 
22:  else
23:     $\mathbf{p}_{uv}^{\text{cpnee}} \leftarrow 0$ 
24:     $\triangleright$  Perform Unidirectional sampling
25:     $\mathbf{g}_u^{\text{uni}}, \mathbf{p}_u^{\text{uni}} \leftarrow \text{SAMPLESTDGAUSSIAN}()$ 
26:     $\mathbf{p}_{uv}^{\text{uni}} \leftarrow \mathbf{p}_u^{\text{uni}} \mathbf{p}_v$ 
27:     $\nabla f^{\text{uni}} \leftarrow \nabla f_{\text{fix}} + \mathbf{g}_u^{\text{uni}} \hat{\mathbf{u}}$ 
28:     $\mathbf{n}^{\text{uni}} \leftarrow \text{NORMALIZE}(\nabla f^{\text{uni}})$ 
29:     $\omega_s^{\text{uni}} \leftarrow \text{REFLECT}(\mathbf{n}^{\text{uni}}, \omega)$ 
30:     $\rho^{\text{uni}} \leftarrow \text{EVALBSDFCOSINE}(\mathbf{p}_s^f, \omega, \omega_s^{\text{uni}}, \mathbf{n}^{\text{uni}})$ 
31:     $\mathbf{w}^{\text{uni}} \leftarrow \text{MISWEIGHT}(\mathbf{p}_{uv}^{\text{cpnee}}, \mathbf{p}_{uv}^{\text{uni}})$ 
32:    return  $L_o \leftarrow L_o + \mathbf{w}^{\text{uni}} \rho^{\text{uni}}$ 
33:   $L(\mathbf{p}_s^f, \omega_s^{\text{uni}}, \nabla f^{\text{uni}}, \text{depth} + 1, \text{seed} + 1)$ 

```

4.6. Generalization to Non-stationary GPISes

Our method is compatible with non-stationary GPISes in the same way as emitter GPIS NEE, as both rely on the 2D gradient distribution $\mathbf{g}_{uv}$ at the intersection. As discussed by Xu et al. [XBdEJ25], multi-resolution sparse convolution noise [LLD11] can be used to efficiently evaluate GPISes with spatially varying length scales. For general non-stationary anisotropy, if the spatial variation is sufficiently smooth, we can still obtain a good approximation of $\mathbf{g}_{uv}$ using local GPIS statistics. More precisely, this approximation remains accurate as long as the anisotropy is approximately constant over a region on the order of the length scale.

5. Implementation and Results**5.1. Implementation**

We implemented our method in the Tungsten renderer [Bit18] for comparison against Xu et al. [XBdEJ25]; full source code is available at <https://github.com/dartmouth-vcl/gpis-cpnee-tungsten>. Algorithm 1 gives pseudocode for the MIS(UNI, CPNEE) combination; additionally incorporating emitter NEE is straightforward.




Figure 6: When using bounding spheres to handle arbitrarily shaped light sources, sampling efficiency is influenced by the tightness of the bound. While CPNEE provides a clear advantage over emitter NEE for the cube and bunny lights, for the thin triangle light (orange inset), our method exhibits 1.5 $\times$ as much error as emitter NEE. Even with such conservative bounds, however, CPNEE ultimately outperforms emitter NEE as GPIS roughness decreases (bottom plate).

We use MIS with the one-sample model [Vea97] to provide equal-sample comparison, i.e., at each path vertex, we first choose one strategy out of the existing three (Unidirectional, NEE, CPNEE) with equal probability, then weight the strategy with the balance heuristic to estimate the direct illumination. We then continue the path via unidirectional sampling.

The distribution of $\mathbf{g}$ is different depending on the applied memory model [SdEBJ24; XBdEJ25]. Xu et al. [XBdEJ25] found that their Renewal Half+ model – which neglects correlations of gradient components $\mathbf{g}_{uv}$ between two consecutive GPIS path vertices – leads to almost identical results as Renewal+ while providing more latitude for importance sampling. We therefore adopt Renewal Half+, but our approach naturally extends to other memory models. In Algorithm 1, Renewal Half+ is reflected in two places: Line 3 conditions the ray march on the incoming gradient ∇f , and the final recursive call on Line 31 forwards ∇f^{uni} to the next bounce. Under Renewal, Line 3 would not condition on ∇f ; since this affects only the distribution of $\mathbf{g}_w$ and not $p(\mathbf{g}_{uv})$, Line 5 onward require no modification. Under Renewal+, Line 5 and Line 12 would additionally condition the sampling of $\mathbf{g}_v$ and $\mathbf{g}_u$ on the incoming ∇f .

We use relative mean squared error (relMSE), defined as $\frac{1}{n} \sum_{i=1}^n (I_i - \hat{I}_i)^2 / (I_i^2 + \varepsilon)$, where I_i and $\hat{I}_i$ are the corresponding pixels in the reference and noisy images, respectively; we set $\varepsilon = 1 \times 10^{-6}$. We render the ground truth images using MIS(UNI, NEE) for isotropic GPISes and MIS(UNI, CPNEE) for anisotropic GPISes; in the anisotropic scenario, NEE cannot converge within a reasonable time budget. We show experiment results for stationary GPISes with a squared exponential covariance kernel; while we have not tested this, we would expect CPNEE to provide similar benefits on non-stationary GPISes.

5.2. Results

We present extensive results across multiple scenes – including isotropic and anisotropic, as well as surface-like and volume-like 3D GPISes – to demonstrate the effectiveness and robustness of our CPNEE method. We further show that our approach integrates

seamlessly with existing methods via MIS and supports arbitrary light source shapes enclosed within a sphere. We include each full OpenEXR image in the supplementary materials with an interactive web viewer to allow comparison of the error images against the reference.

We rendered all results as equal-sample count comparisons on an Intel i7-13700K (16 cores, 24 threads). Because CPNEE relies purely on simple, closed-form geometric operations, its overhead is low compared to the cost of GPIS ray marching. We confirmed this by running an equal-time comparison on the scene from Fig. 1. CPNEE (top right) in isolation takes 4%–10% fewer samples than emitter NEE (top left); while with MIS (bottom left vs. bottom right) sample counts are within 2% of each other. This negligible per-sample cost is vastly outweighed by variance reductions of 92× in highly anisotropic cases and 2.4× in isotropic ones, confirming that our equal-sample comparisons reliably proxy equal-time performance.

We evaluate our method on a modified version of the classic Veach MIS scene, which features four metallic plates with varying roughness. Seyb et al. [SdEBJ24] show that a heightfield GPIS with noise amplitude σ and lateral correlation length ℓ corresponds to a Beckmann roughness of $\alpha = \sigma\sqrt{2}/\ell$. We render the four plates using 3D GPISes while adhering to this mapping. In Fig. 7, the roughness values from top to bottom are $\alpha = 0.5, 0.1, 0.03, 0.001$.

The top two rows correspond to isotropic GPISes, while the bottom two rows show anisotropic GPISes. For anisotropic cases, we apply the roughness mapping to the GPIS parameters along the two lateral axes, and increase the length scale by a factor of 100× along the macro-normal direction to introduce anisotropy. We show the rendering of unidirectional, emitter NEE, and our CPNEE method in the first and the third rows. The second and the fourth rows show a comprehensive set of MIS combinations between the strategies.

For the anisotropic GPISes (third row) that CPNEE was designed for, it outperforms emitter NEE and UNI throughout the scene as expected. Interestingly, our method also demonstrates an overall benefit in the isotropic case (first row) compared to both UNI and NEE. In particular, CPNEE still performs significantly better for slabs with glossy appearances illuminated by a large light source (green inset). This is because CPNEE importance samples directions on the narrow arc-like cvGDFs, and the corresponding scattering circles often have an intersection with the light source. The scattering circle and light cap might not always intersect, which leads to variance. Such misses occur more frequently for small lights and rough surfaces, where we see the advantages of CPNEE diminish. For configurations involving glossy surfaces with small lights or rough surfaces with large lights, analytical reasoning is harder. Nevertheless, in this scene, our CPNEE method consistently outperforms UNI and NEE in most scenarios. Notably, the worst-case behavior typically arises from UNI or NEE, never from CPNEE.

The second and the fourth rows of Fig. 7 compare various MIS combinations of the strategies using the balance heuristic. We show each two-way MIS combination as well as three-way MIS across all techniques. In practice, while three-way MIS may slightly increase variance in certain boundary regions compared to CPNEE alone, this defensive, parameter-free combination remains robust across a wide range of configurations.

Additionally, in Fig. 8, we demonstrate the practicality of our method on a more complex dragon scene with global illumination. We show results for both surface-like and volume-like GPISes, using isotropic and anisotropic covariance kernels in the first and second rows, respectively. In the second row, the noise is compressed along the direction pointing into the screen to approach a more volume-like GPIS. Our results show that CPNEE works robustly across various types of GPISes. We observe a smaller but still consistent improvement for the isotropic case. Although CPNEE is not designed for volume-like GPISes with broad normal distributions, it still effectively complements UNI and emitter NEE; three-way MIS performs comparably to two-way MIS in this regime. Notably, in Fig. 1, emitter NEE fails completely in the anisotropic settings for both surface and volume-like GPISes. The MIS(UNI, NEE) combination (bottom left) is more robustness, but since NEE performs so poorly here, it mostly gets down-weighted while providing little more than overhead. In contrast, CPNEE substantially improves over UNI in the surface-like case, so defensive three-way MIS (bottom right) outperforms two-way MIS (bottom left).

6. Limitation and Future Work

As shown in Fig. 6, sampling efficiency degrades when the bounding sphere fits the emitter loosely. For flat or elongated emitters, a large portion of the spherical bound is empty, leading to many rejected samples and degraded performance. One way to alleviate this issue is to approximate the emitter geometry with multiple proxy spheres. Developing efficient intersection strategies for the common case of quad- or triangle-shaped light sources is an interesting direction for future work. For planar lights, it may be possible to replace the bounding sphere with a bounding disk and intersect it with the set of possible scattering directions (an elliptical 3D cone). This approach resembles earlier cone tracing techniques [Ama84].

As we noted in the Veach scene, much of the remaining variance in our method occurs when the scattering circle misses the light cap. For degenerate heightfield-like GPISes, this is determined by the g_w sampled during ray marching. It might be possible to importance sample g_w to induce scattering circle–light cap intersections, but changing $p_{mc}(g_w)$ would require evaluating more complex terms in the estimator which currently cancel. For general GPISes we have further flexibility via the sampled g_v . Currently we importance sample g_v according to the GPIS distribution, $p_{mc}(g_v) = p(g_v) = \mathcal{N}(0, \sigma)$, which is known in closed form. To limit such misses, it may be possible to at least compute the minimum and maximum g_v that lead to scattering circle–light cap intersections, and sampling from the corresponding truncated Gaussian.

For general GPISes, an additional degree of freedom is the choice of the orthonormal basis $\mathbf{u}, \mathbf{v}$, as discussed in Sec. 4.3.1. These axes map from isotropic space back to world space as $\hat{\mathbf{v}}$ and $\hat{\mathbf{u}}$. Our method performs best when $|\hat{\mathbf{v}}|$ is small: a smaller $|\hat{\mathbf{v}}|$ limits how much the scattering circle’s center shifts with the sampled g_v , reducing the chance of a miss between the two circles. We conduct ablation studies on the effect of this axis choice in Fig. 9 across different GPIS configurations, showing that our method benefits more from horizontal roughness and is less advantageous for vertical roughness. We leave a more comprehensive investigation of the choice of $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$, particularly for the general 3D GPIS setting, to future work.




Figure 7: We compare our method with previous sampling strategies (UNI and NEE) in the Veach scene. In the first two rows, we use isotropic 3D GPISes. For a small light source with rough surfaces (orange), our method performs worse than NEE due to the relatively short circle segment and frequent misses between the two cones. Nevertheless, our method still provides benefits at locations with large light sources (green). In the last two rows, where we use a heightfield-like GPIS with a large length scale along the surface normal, our method provides increasing benefit and outperforms emitter NEE at every location in the scene. In practice, we recommend using a three-way MIS (highlighted in the red box) to accommodate a broad range of scene configurations. The three-way MIS result strikes a middle ground among the three methods while performing reasonably well in most cases.

MIS strikes a conservative balance: it often improves the worst case, but this comes at the cost of wasting effort on down-weighted strategies even when one may be the clear winner. While it is often difficult to know a-priori which strategies will perform best, our

experiments suggest a strong prior: emitter NEE excels with near-isotropic GPISes and small light sources, while CPNEE performs best with highly anisotropic GPISes and larger emitters. It may therefore be possible to extract further performance by dynamically




Figure 8: Results for the dragon scene. Top: isotropic 3D GPIS. Our method performs slightly better than UNI and NEE. Three-way MIS therefore shows a slight improvement over two-way MIS with UNI and NEE. Bottom: anisotropic GPIS with the global anisotropy direction pointing into the screen. Although we do not specifically target volume-like GPISes, our method still shows a significant improvement over Xu et al.’s NEE method (third column). With three-way MIS, we handle a broader range of GPISes more robustly (last column).

selecting strategies based on the local degree of anisotropy and emitter size.

Finally, we note that the Gaussian scattering circle that arises in our approach bares a striking resemblance to the Marschner hair model’s scattering lobe [MJC*03]. It may therefore be possible to adapt our technique for better NEE sampling for hair and curve geometry.

7. Conclusion

We introduced a novel NEE technique for Gaussian process implicit surfaces that overcomes the failure of existing methods on highly anisotropic and heightfield-like GPISes, while also providing benefits for general GPISes. By explicitly sampling the manifold of admissible directions restricted to the light source’s spherical cap and employing MIS, we successfully bridge the gap between efficiently rendering surface-like and volume-like stochastic geometry.

Acknowledgement

We thank the anonymous reviewers for their helpful suggestions and the members of the Dartmouth Visual Computing Lab for valuable discussions. This work was supported by NSF award 2440472.

References

- [AK90] ARVO, J. R. and KIRK, D. “Particle transport and image synthesis”. *Computer Graphics (Proceedings of SIGGRAPH)* 24.4 (Sept. 1, 1990). ISSN: 0097-8930. DOI: [10/dtp6gd2](https://doi.org/10/dtp6gd2).
- [Ama84] AMANATIDES, J. “Ray tracing with cones”. *Computer Graphics (Proceedings of SIGGRAPH)* 18.3 (July 1, 1984). ISSN: 0097-8930. DOI: [10/b64smh8](https://doi.org/10/b64smh8).
- [BD16] BILLEN, N. and DUTRÉ, P. “Line sampling for direct illumination”. *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 35.4 (July 1, 2016). ISSN: 0167-7055. DOI: [10/f84z2h2](https://doi.org/10/f84z2h2).
- [Bit18] BITTERLI, B. *Tungsten Renderer*. 2018. URL: <https://github.com/tunabrain/tungsten/>.




Figure 9: In this GPIS heightfield scene, each plate is split into regions with higher roughness horizontally (top half) and vertically (bottom half). CPNEE provides larger benefits for horizontally rough plates (upper; orange and blue insets) than vertically rough ones (lower; green and purple insets). For the orange and green insets we additionally visualize the cvGDFs and scattering distributions $p(\omega_s)$, with the light cap’s projection overlaid. The more concentrated cvGDFs in the vertically rough cases explain the limited benefit; rotating the $\mathbf{u}, \mathbf{v}$ axes by 90° mitigates this.

- [BS63] BECKMANN, P. and SPIZZICHINO, A. *The Scattering of Electromagnetic Waves from Rough Surfaces*. NY: Pergamon Press, 1963 2.
- [CT81] COOK, R. L. and TORRANCE, K. E. “A reflectance model for computer graphics”. *Computer Graphics (Proceedings of SIGGRAPH)* 15.3 (Aug. 1981). ISSN: 0097-8930. DOI: [10/br5ps62](https://doi.org/10/br5ps62).

- [dEBWZ23] D'EON, E., BITTERLI, B., WEIDLICH, A., and ZELTNER, T. "Microfacet theory for non-uniform heightfields". *ACM SIGGRAPH Conference Papers*. Los Angeles, CA: ACM Press, 2023. doi: [10/kn652](#).
- [HCWY26] HUANG, M., CUI, Y., WANG, B., and YAN, L. *Macrofacet Theory for Gaussian Process Statistical Surfaces*. 2026. arXiv: [2603.00280 \[cs.GR\]](#) 2.
- [HdE14] HEITZ, E. and d'EON, E. "Importance sampling microfacet-based BSDFs using the distribution of visible normals". *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 33.4 (2014). doi: [10/f6fgph3](#).
- [HDF15] HANIKA, J., DROSKE, M., and FASCIONE, L. "Manifold next event estimation". *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 34.4 (2015). doi: [10/f7mbxk2](#).
- [HHdE16] HEITZ, E., HANIKA, J., d'EON, E., and DACHSBACHER, C. "Multiple-scattering microfacet BSDFs with the Smith model". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 35.4 (July 2016). doi: [10/f89kkm2](#).
- [ICG86] IMMEL, D. S., COHEN, M. F., and GREENBERG, D. P. "A radiosity method for non-diffuse environments". *Computer Graphics (Proceedings of SIGGRAPH)* 20.4 (Aug. 1986). ISSN: 0097-8930. doi: [10/dmjm9t3](#).
- [JAM*10] JAKOB, W., ARBREE, A., MOON, J. T., BALA, K., and MARSCHNER, S. "A radiative transfer framework for rendering materials with anisotropic structure". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 29.4 (July 2010). ISSN: 07300301. doi: [10/ftr7st2](#).
- [Kaj86] KAJIYA, J. T. "The rendering equation". *Computer Graphics (Proceedings of SIGGRAPH)* 20.4 (Aug. 1986). ISSN: 0097-8930. doi: [10/cvf53j3](#).
- [Lew84] LEWIS, J.-P. "Texture synthesis for digital painting". *Proceedings of the 11th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '84. New York, NY, USA: Association for Computing Machinery, Jan. 1, 1984. ISBN: 978-0-89791-138-2. doi: [10/bh7mdz3](#).
- [LLD11] LAGAE, A., LEFEBVRE, S., and DUTRE, P. "Improving Gabor noise". *IEEE Transactions on Visualization and Computer Graphics* 17.8 (Aug. 2011). ISSN: 1077-2626. doi: [10/cr3qwx7](#).
- [LLDD09] LAGAE, A., LEFEBVRE, S., DRETTAKIS, G., and DUTRÉ, P. "Procedural noise using sparse Gabor convolution". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 28.3 (July 27, 2009). ISSN: 0730-0301, 1557-7368. doi: [10/cvv5s73](#).
- [LRPB23] LUCAS, S., RIBARDIERE, M., PACANOWSKI, R., and BARLA, P. "A Micrograin BSDF Model for the Rendering of Porous Layers". *ACM SIGGRAPH Asia Conference Papers*. SA '23. New York, NY, USA: Association for Computing Machinery, Dec. 11, 2023. ISBN: 979-8-4007-0315-7. doi: [10/gtb6cz2](#).
- [MJC*03] MARSCHNER, S. R., JENSEN, H. W., CAMMARANO, M., WORLEY, S., and HANRAHAN, P. "Light scattering from human hair fibers". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 22.3 (July 1, 2003). ISSN: 0730-0301. doi: [10/b4cprd10](#).
- [ON94] OREN, M. and NAYAR, S. K. "Generalization of Lambert's reflectance model". *Annual Conference Series (Proceedings of SIGGRAPH)*. New York, NY, USA: ACM Press, 1994. ISBN: 978-0-89791-667-7. doi: [10/dxrtcv2](#).
- [RW06] RASMUSSEN, C. E. and WILLIAMS, C. K. I. *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. Cambridge, Mass.: MIT Press, 2006. ISBN: 978-0-262-18253-9. URL: <https://gaussianprocess.org/gpml/> 2.
- [SdEJ24] SEYB, D., d'EON, E., BITTERLI, B., and JAROSZ, W. "From microfacets to participating media: A unified theory of light transport with stochastic geometry". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 43.4 (July 2024). doi: [10/gt5nh91-3,7,8](#).
- [SJ19] SALESIN, K. and JAROSZ, W. "Combining point and line samples for direct illumination". *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 38.4 (July 2019). ISSN: 0167-7055, 1467-8659. doi: [10/gf6rx62](#).
- [SWZ96] SHIRLEY, P., WANG, C., and ZIMMERMAN, K. "Monte Carlo techniques for direct lighting calculations". *ACM Transactions on Graphics* 15.1 (Jan. 1996). ISSN: 0730-0301. doi: [10/ddgbgg2](#).
- [Vea97] VEACH, E. "Robust Monte Carlo Methods for Light Transport Simulation". PhD thesis. Stanford University, Dec. 1997 7.
- [VG95] VEACH, E. and GUIBAS, L. J. "Optimally combining sampling techniques for Monte Carlo rendering". *Annual Conference Series (Proceedings of SIGGRAPH)*. Vol. 29. ACM Press, Aug. 1995. ISBN: 978-0-89791-701-8. doi: [10/d7b6n42](#).
- [WMLT07] WALTER, B., MARSCHNER, S. R., LI, H., and TORRANCE, K. E. "Microfacet models for refraction through rough surfaces". *Rendering Techniques (Proceedings of the Eurographics Symposium on Rendering)*. Eurographics Association, June 2007. ISBN: 978-3-905673-52-4. doi: [10/gfz4kg2](#).
- [XBdEJ25] XU, K., BITTERLI, B., d'EON, E., and JAROSZ, W. "Practical Gaussian process implicit surfaces with sparse convolutions". *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 44.6 (Dec. 4, 2025). ISSN: 0730-0301. doi: [10.1145/37633291-4,7,10](#).