Efficient Monte Carlo Rendering of Implicit-Shaped Volumetric Emitters

Jiawei Huang
IDEA
huangjiawei@idea.edu.cn

Shaokun Zheng
Tsinghua University
zsk20@mails.tsinghua.edu.cn

Kun Xu
Tsinghua University
xukun@tsinghua.edu.cn

Yoshifumi Kitamura
Tohoku University
kitamura@riec.tohoku.ac.jp

Jiaping Wang
IDEA
jiapw.cg@gmail.com

Figure 1: SDF emitters used in production: moving cars illuminated by SDF-defined capped-torus emitters, distributed through the tunnel. The tori are modeled as a volumetric emitters using an implicit function. **Left:** Baseline BSDF importance sampling. **Center:** Our method enables importance-sampled next-event estimation using solid-angle sampling from the volumetric SDF emitters, and the tractable PDF seamlessly works with multiple importance sampling (MIS). **Right:** Reference. Baseline and our method are rendered with 256 samples per pixel. Our technique captures the soft volumetric glow and parallax-aware intensity with significantly reduced variance, enabling fast convergence at a fraction of the cost.

Abstract

We present a practical and unbiased sampling technique for volumetric light sources defined by signed distance functions (SDFs). SDFs can compactly represent complex procedural shapes and support Boolean operations for flexible composition. However, their use as emitters remains underexplored in photorealistic rendering due to the absence of efficient sampling strategies. Our key insight is to model the interior of an SDF as a uniform volume and to project volumetric samples onto the unit sphere centered at the shading point, yielding a solid-angle distribution that captures the emitter's relative spatial layout. We derive the directional probability densities for analytic primitives and generalize the formulation to arbitrary SDFs via Monte Carlo volume estimation. Leveraging robust sphere tracing, our method enables accurate and efficient sampling of SDF emitters without requiring explicit surface parameterization, voxelization, or precomputed tables. Compared

to baseline approaches such as uniform directional sampling or surface approximations, our technique achieves significantly lower variance and runtime in next-event estimation, while broadening the expressive power of light sources in rendering.

Keywords: Monte Carlo Rendering, Signed Distance Function, Next-Event Estimation, Importance Sampling

1. Introduction

In artistic and production rendering, light sources play a crucial role not only in illuminating scenes but also in conveying mood, form, and story. In physically-based rendering systems, light sources are typically modeled as emitters: objects that define how energy is radiated into the scene. While most systems support a variety of surface-based emitters such as points, polygons, or area meshes, they offer far fewer options for volumetric emission. This limitation becomes especially apparent when artists aim to represent emitters with rich internal structure or nontrivial spatial form. In such cases, mesh emitters are often used as a workaround, approximating complex volumetric behavior through surfaces. However, this approach introduces significant overhead in sampling, storage, and shading complexity, limiting both performance and creative flexibility.

This disparity is not due to any fundamental weakness of volumetric emitters, but rather to the challenges they pose in sampling and integration. Existing volumetric models tend to rely on discretized voxel grids, which trade fidelity for tractability. By contrast, signed distance functions (SDFs) offer a compact, expressive alternative for procedural modeling—supporting smooth blends, Boolean operations, and structural edits within a unified representation. Despite their widespread use for geometry definition and scene evaluation, SDFs remain almost entirely unexplored as emitters in Monte Carlo rendering frameworks, owing to the lack of efficient and unbiased sampling strategies over their implicit interiors.

We introduce SDF emitters: a new class of volumetric emitters defined directly by signed distance functions. Our core idea is to treat the interior of an SDF-defined region as a uniform emissive volume and to derive a solid-angle distribution at the shading point through spatial-to-directional projection. This yields accurate and efficient importance sampling without requiring meshing, voxelization, or precomputed tables. Seamlessly integrating into existing Monte Carlo pipelines, SDF emitters expand the expressive range of light modeling—enabling high-fidelity, artist-controllable illumination with minimal overhead and substantial variance reduction.

By integrating SDF emitters seamlessly into existing Monte Carlo rendering pipelines, our method enhances the expressive power and flexibility of light source modeling, with only small variance through effective importance sampling in next-event estimation.

In summary, the main contributions of this paper are

- The introduction of SDF emitters as a new class of volumetric light sources, enabling implicit geometries to participate in lighting without meshing or discretization;
- The formulation of an unbiased importance sampling technique based on solid-angle distributions projected from volumetric samples; and
- A practical strategy for integrating SDF emitters into Monte Carlo rendering, leveraging robust sphere tracing for intersection and evaluation.

2. Related Work

2.1. Procedural and Implicit Geometry in Rendering

Procedural and implicit representations offer a compact and flexible means of describing complex geometry. Unlike traditional mesh-based models, these methods define shapes through functions or composition rules, allowing efficient storage, evaluation, and manipulation. Such representations are widely used for geometric modeling in rendering, particularly for structured or fractal-like objects.

Signed distance functions (SDFs) are a common choice in this domain. They describe geometry implicitly through a signed scalar field and support constructive operations like unions and intersections. SDFs are broadly used in procedural modeling, differentiable rendering (25), and neural geometry representations (13; 24), and can be parameterized through networks to yield continuous, high-fidelity fields (18; 12).

In most of these applications, SDFs serve as surface proxies for visibility testing, shading, or optimization, valued for their compactness and efficient intersection routines. In contrast, we explore their use as volumetric emitters defined over interior regions, and develop a solid-angle–based directional sampling technique locally adapted to each shading point, without requiring meshing or voxelization.

2.2. Light Source Modeling and Volumetric Emitters

Various strategies exist for modeling light sources in physically based rendering. Analytic emitters like points, spheres, or rectangles benefit from closed-form intersection and sampling (2; 8; 20), but their simplicity limits the range of visual effects. Mesh lights, composed of tessellated surfaces, offer greater flexibility but increase memory and sampling complexity, as they effectively act as large aggregates of triangle emitters.

Volumetric emitters, such as those found in fire or fog, are typically modeled using density grids or proxy geometry (5). Sampling such media often involves integrating emission along rays, as in the works of Villemin *et al*. (22) and Simon *et al*. (17), which interpret emission in terms of directional contributions from volumetric regions.

Our approach builds on similar intuition by constructing directional sampling distributions from volume integrals. Unlike grid-based formulations, however, we operate directly on implicit volumes defined by SDFs. This allows us to retain geometric continuity while avoiding the discretization and memory costs typically associated with volumetric grids.

2.3. Importance Sampling for Direct Lighting

Accurate sampling of light sources is critical for reducing variance in Monte Carlo rendering. Standard techniques like next-event estimation (NEE) and multiple importance sampling (MIS) (21) combine samples from light and BSDF distributions, but they can still struggle in the presence of occlusion, complex emission profiles, or large light sets. To enable NEE for surface behind refracted surfaces, manifold exploration and polynomial systems are leveraged (9; 26; 7).

To better align sampling with the integrand, many methods operate in projected solid angle space (2), particularly for structured emitters like disks and rectangles (8; 20). Beyond geometric terms, product importance sampling strategies aim to sample the full product of illumination and material scattering; for instance, Diolatzis *et al*. (6) extended Linearly Transformed Cosines (LTC) (11) to guide paths by the product of light and BSDF. Our method complements these strategies by defining the normalized solid-angle distributions (p_L) for implicit volumetric emitters required by such frameworks, without requiring surface parameterization. This enables spatially adapted sampling for a broader class of emitter shapes and representations. In contrast to data-driven approaches (28) that often require pre-training, our technique remains purely algorithmic and lightweight, making it uniquely suitable for dynamic procedural content.

For scenes with many lights, hierarchical techniques like Lightcuts (23), cone trees (4), and light BVHs (14) improve scalability by clustering and pruning candidate emitters. More recently, sample reuse frameworks like ReSTIR (3) introduced spatio-temporal resampling, amortizing sampling costs across space and time to enable real-time rendering of direct lighting in complex scenes. Extensions like ReSTIR GI (15) and ReSTIR PT (16) further generalize this idea to indirect illumination.

Our method immediately fits into existing direct lighting pipelines, and can work seamlessly with these many-light sampling frameworks.

3. Formulation of SDF Emitter Sampling

Let $\phi\colon \mathbb{R}^3 \to \mathbb{R}$ be a signed distance function (SDF) defining an implicit surface:

$$\mathcal{S} = \{\mathbf{x} \in \mathbb{R}^3 \mid \phi(\mathbf{x}) = 0\}.$$

We define the corresponding solid region (i.e., the interior of the shape) as

$$\Omega = \{\mathbf{x} \in \mathbb{R}^3 \mid \phi(\mathbf{x}) \leq 0\}.$$

To model volumetric emission, we introduce the indicator function:

$$v(\mathbf{x}) = \begin{cases} 1, & \text{if } \mathbf{x} \in \Omega, \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

This describes a volume with uniform density contained within the implicit boundary defined by the SDF.

We denote the total volume of the region by

$$V := \int_{\mathbb{R}^3} v(\mathbf{x})\, d\mathbf{x}, \tag{2}$$

which will serve as a normalization factor throughout.

3.1. Volumetric Emission Model

We model the emitted radiance in a direction $\omega \in \mathbb{S}^2$ from a volumetric emitter by integrating an emissive density $\rho(\mathbf{x})$ along the ray from a shading point $\mathbf{x}_0$. The function $\rho : \mathbb{R}^3 \to \mathbb{R}$ defines the local radiative emission per unit length. The radiance is then given by

$$L_e(\omega) = \int_0^\infty v(\mathbf{x}_0 + t\omega) \cdot \rho(\mathbf{x}_0 + t\omega)\, dt, \tag{3}$$

where the indicator function v restricts integration to the emitter interior Ω. This integral reduces to a sum over ray–volume intersection intervals $[r_i^-, r_i^+]$:

$$L_e(\omega) = \sum_i \int_{r_i^-}^{r_i^+} \rho(\mathbf{x}_0 + r\omega)\, dr. \tag{4}$$

In the case of uniform emission (*i.e.* $\rho(\mathbf{x}) \equiv \text{const}$), we may normalize by total emitter volume V and set $\rho = 1/V$ to ensure unit total power. This yields the simplified form:

$$L_e(\omega) = \frac{1}{V} \sum_i (r_i^+ - r_i^-). \tag{5}$$

For general $\rho(\mathbf{x})$, we may estimate the radiance using Monte Carlo integration. Drawing one point $r^* \sim \mathcal{U}[r_i^-, r_i^+]$ per segment, where $\mathcal{U}$ denotes the uniform distribution, the unbiased estimator becomes:

$$\widehat{L}_e(\omega) = \sum_i (r_i^+ - r_i^-) \cdot \rho(\mathbf{x}_0 + r^*\omega). \tag{6}$$

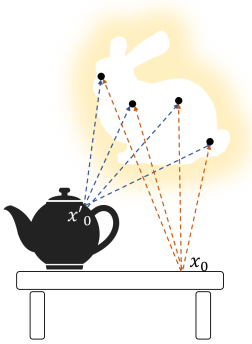

Figure 2: A spatial-to-directional projection: sampling a point $\mathbf{x} \in \Omega$ inside the emitter (white bunny) and projecting toward the shading point $\mathbf{x}_0$ yields a direction $\omega = \frac{\mathbf{x}-\mathbf{x}_0}{\|\mathbf{x}-\mathbf{x}_0\|}$. This projection naturally adapts the sampling distribution to the emitter's geometry as seen from $\mathbf{x}_0$. At a different point $\mathbf{x}'_0$, the induced distribution differs due to perspective.

When this is used in a Monte Carlo renderer, each sample's contribution must be weighted by the directional PDF $p(\omega)$, yielding a per-sample estimator:

$$\widehat{L}_e = \frac{1}{p(\omega)} \sum_i (r_i^+ - r_i^-) \cdot \rho(\mathbf{x}_0 + r^*\omega). \tag{7}$$

This formulation supports both constant and spatially varying emissive profiles within the same framework, enabling efficient importance sampling of arbitrarily complex volumetric emitters.

3.2. Projection-Based Directional Sampling

For SDF-defined volumes, directly generating a directionally adapted importance sample from a shading point is non-trivial due to the implicit nature of the geometry. However, it is often easier to sample a *spatial* point inside the volume using techniques like rejection sampling or analytic inversion.

Our key insight is to bridge this gap using a spherical projection from spatial samples to locally adapted directional ones (see Fig. 2). A point sampled uniformly inside the emitter induces a non-uniform distribution over directions at the shading point, which reflects the emitter's shape and relative position.

3.2.1 The spatial-to-directional projection

Let $\mathbf{x}_0 \in \mathbb{R}^3$ be the shading point. Given a point $\mathbf{x} \sim v(\mathbf{x})$, sampled uniformly within the volume V, the spatial distribution of $\mathbf{x}$ is

$$p(\mathbf{x}) = \frac{1}{V} v(\mathbf{x}) = \frac{1}{V}. \tag{8}$$

We define the corresponding direction by projecting to the unit sphere:

$$\omega = \frac{\mathbf{x} - \mathbf{x}_0}{\|\mathbf{x} - \mathbf{x}_0\|}, \quad \omega \in \mathbb{S}^2. \tag{9}$$

This defines a mapping from the volume to the unit sphere centered at $\mathbf{x}_0$, effectively converting uniform spatial samples into non-uniform direction samples that naturally account for the emitter's geometry.

Uniform sampling within V can be implemented through analytic inversion for primitive shapes or via rejection sampling for general SDFs. The resulting distribution over directions reflects both the geometry and spatial extent of the volume relative to the shading point.

3.2.2 The induced directional PDF

To derive the probability density function $p(\omega|\mathbf{x}_0)$ induced on the unit sphere, we perform a change of variables from Cartesian to spherical coordinates centered at $\mathbf{x}_0$:

$$\mathbf{x} = \mathbf{x}_0 + r\omega, \quad r \in [0, \infty), \quad \omega \in \mathbb{S}^2. \tag{10}$$

The volume element transforms as $d\mathbf{x} = r^2\, dr\, d\omega$, and the spatial PDF becomes:

$$f(r, \omega) = \frac{1}{V} v(\mathbf{x}_0 + r\omega) \cdot r^2. \tag{11}$$

Marginalizing over r gives the directional PDF:

$$p(\omega|\mathbf{x}_0) = \frac{1}{V} \int_0^\infty v(\mathbf{x}_0 + r\omega) \cdot r^2\, dr. \tag{12}$$

In practice, the volume V is often defined by a union of intervals $[r_i^-, r_i^+]$ along a ray. This leads to a discrete summation form:

$$p(\omega|\mathbf{x}_0) = \frac{1}{V} \sum_i \int_{r_i^-}^{r_i^+} r^2\, dr = \frac{1}{V} \sum_i \frac{(r_i^+)^3 - (r_i^-)^3}{3}. \tag{13}$$

This expression captures the contribution of each intersected segment of the emitter volume along the ray in direction ω.

To understand why our projection-based sampling is effective, consider the relationship between our derived PDF and the ideal PDF. For a uniform volumetric emitter, the ideal PDF is proportional to the radiance, which is the line integral of the volume indicator: $p^\star(\omega) \propto L(\omega) = \sum_i (r_i^+ - r_i^-)$.

Our method samples based on the volume projected onto the sphere, yielding a PDF proportional to the difference of cubes (Eq. 13): $p(\omega) \propto \sum_i ((r_i^+)^3 - (r_i^-)^3)$. Factorizing the cubic term for a single segment gives:

$$(r^+)^3 - (r^-)^3 = (r^+ - r^-) \cdot \left((r^+)^2 + r^+ r^- + (r^-)^2\right). \tag{14}$$

For an emitter at a distance D where the thickness is small relative to D (i.e., $r^{\pm} \approx D$), the second term approaches $3D^2$. Consequently, $p(\omega) \approx C \cdot (r^+ - r^-) \propto p^\star(\omega)$.

This implies that for distant emitters, our sampling distribution converges to the ideal proportional-to-radiance distribution. In the near field, the quadratic factor $((r^+)^2 + \dots)$ gently over-weights deeper parts of the volume, which is a safe heuristic that avoids variance spikes from undersampling the emitter's core.

3.3. Integration Strategies for Common SDF Types

Having established the general projection-based framework, we now turn to the practical strategies used to compute volume integrals and ray intersections for different classes of SDFs. We focus on two representative cases: analytic primitives and general composite SDFs.

3.3.1 Analytic Sampling of Primitives

For common analytic SDF primitives such as *spheres*, *cubes*, and *tori*, both the total volume V and the ray–volume intersection intervals $[r_i^-, r_i^+]$ can be computed in closed form. These primitives serve as ideal testbeds for validating our method under controlled, low-variance conditions.

Sphere. Consider a sphere of radius R centered at the origin, defined by the SDF

$$\phi(\mathbf{x}) = \|\mathbf{x}\| - R.$$

The volume is given by $V = \frac{4}{3}\pi R^3$, and a ray $\mathbf{x}_0 + r\omega$ intersects the volume by solving the quadratic

$$\|\mathbf{x}_0 + r\omega\|^2 = R^2,$$

which yields:

$$r^{\pm} = -(\omega \cdot \mathbf{x}_0) \pm \sqrt{(\omega \cdot \mathbf{x}_0)^2 - \|\mathbf{x}_0\|^2 + R^2}.$$

If a real solution exists, the ray intersects the sphere in a single segment $[r^-, r^+]$.

Importantly, *uniform volume sampling* is fully analytic. to sample a point uniformly inside the sphere, draw radius $r = R \cdot u^{1/3}$, direction $\omega \sim \mathcal{U}(\mathbb{S}^2)$, and construct $\mathbf{x} = \mathbf{x}_{\text{center}} + r \cdot \omega$.

Torus. For more complex primitives, such as a torus defined by

$$\phi(\mathbf{x}) = \left(\sqrt{x^2 + y^2} - R\right)^2 + z^2 - r^2,$$

the volume can also be computed analytically as $V = 2\pi^2 R r^2$. Ray–volume intersections are found by solving a quartic polynomial in r. While intersection and evaluation are exact, *uniform sampling within the torus volume* is nontrivial and typically requires rejection sampling.

Despite this, our projection-based method remains unbiased. The induced directional PDF can be computed via Eq. (13), and performs well even when rejection sampling is used for interior points.

3.3.2 Generalization to arbitrary SDFs

SDFs are capable of representing complex shapes formed via Boolean compositions, procedural modeling, or learned neural representations. While analytic formulae for V and ray intersections are not generally available in such cases, our framework remains applicable.

Volume Estimation. The volume V can be estimated using Monte Carlo integration over a bounding region. Since $v(\mathbf{x}) \in \{0, 1\}$, the integration reduces to estimating the fraction of points inside the volume, yielding fast convergence.

Ray–Volume Intersections. The intersection intervals $[r_i^-, r_i^+]$ can be computed using robust sphere tracing, which leverages the signed distance property to advance along rays while avoiding surface penetration.

Unbiasedness. Estimating V with a Monte Carlo method introduces only a scalar normalization error:

$$\hat{L}_e(\omega) = \frac{V}{\hat{V}} \cdot L_e(\omega), \quad \frac{\hat{L}_e(\omega)}{L_e(\omega)} = \frac{V}{\hat{V}}. \tag{15}$$

This scalar factor does not affect the shape of the directional distribution, preserving the unbiased nature of the estimator and allowing seamless integration with other lighting strategies.

4. Practical Implementation Details

Our SDF emitter is designed for seamless integration into existing Monte Carlo path tracers. We implemented the method in a proprietary GPU production renderer, and also provide the implementation using LUISACOMPUTE's DSL (27), available in the supplemental material. Below, we highlight key aspects of the implementation, with a focus on robustness, and compatibility with common rendering pipelines.

4.1. Directional-Safe Stepping Inside Volumes

In standard sphere tracing (10), ray advancement is governed by the signed distance value $\phi(\mathbf{x})$, which guarantees a conservative step size while outside the implicit surface: that is, when $\phi(\mathbf{x}) > 0$, the ray can safely advance by

$\Delta t = \phi(\mathbf{x})$ without risk of penetrating geometry. However, this guarantee does not hold once the ray enters the interior of a solid volume.

Inside the volume, where $\phi(\mathbf{x}) < 0$, the value $|\phi(\mathbf{x})|$ continues to represent the minimal distance to the surface, but not necessarily along the ray direction. More precisely, the surface point minimizing $\|\mathbf{x} - \mathbf{y}\|$ such that $\phi(\mathbf{y}) = 0$ does not in general lie along the ray direction $\boldsymbol{\omega}$. As a result, stepping by $|\phi(\mathbf{x})|$ in direction $\boldsymbol{\omega}$ may overshoot the nearest intersection, particularly under grazing incidence, leading to missed surface crossings and systematic undercounting of entry and exit events. A naïve remedy of simply using $|\phi(\mathbf{x})|$ inside the solid tacitly assumes that the field is a perfect unit-Lipschitz distance in every direction. However, in practice, this assumption rarely holds once constructive-solid-geometry, smooth blends, or voxel-reconstructed SDFs are involved, so the overshoot problem reappears.

To prevent such overstepping, we adopt a *directional-safe* stepping rule: while inside the volume ($\phi < 0$), we restrict the step size by

$$\Delta t = \min(-\phi(\mathbf{x}), \delta_{\max}),$$

where $\delta_{\max}$ is a user-defined maximum step size, chosen to be no greater than the thickness of the smallest volumetric feature to be resolved. This clamp ensures that even in worst-case directional misalignment scenarios, the ray progresses conservatively without skipping over narrow structures or internal components.

This modification has no effect outside the volume, where the standard step $\Delta t = \phi$ remains valid. In combination with root bracketing and bisection refinement, this directional-safe rule yields a traversal algorithm that is both robust and unbiased, capable of accurately recovering all ray–volume intersections even in complex or fragmented geometries.

To our knowledge, this interior-specific correction to signed distance stepping has not been previously formalized in the literature. While similar ideas have been heuristically employed in real-time applications and shader code, we present it here as a principled and necessary component of accurate volumetric integration.

We utilizes a single tight axis-aligned bounding box (AABB) for each emitter. Once the ray brackets the volume interface during sphere tracing, we perform a few iterations of bisection to refine the intersection root to high precision, ensuring robust entry/exit determination.

4.2. Controllable Emitter Parameters

In digital content creation (DCC) pipelines, a *light source* typically refers to an artist-facing object with parameters like *intensity* and *transform*, abstracting over the underlying emitter model. To ensure our SDF emitters integrate naturally into this paradigm, we support these controls directly.

Intensity. We normalize emission by the volume V of the SDF-defined region. The user-defined intensity then acts as a scalar multiplier on the sampled radiance, preserving consistent brightness across emitters of different shapes or scales. For non-uniform emission, we scale each segment's contribution by the evaluated local color before normalization.

Transform. We support arbitrary diagonal scaling, rotation, and translation. Let the world-space transform be represented as:

$$\mathbf{x} = \mathbf{R}\,\mathbf{S}\,\mathbf{x}' + \mathbf{t},$$

where $\mathbf{S}$ is a diagonal scale matrix, $\mathbf{R}$ is a rotation matrix, and $\mathbf{t}$ is a translation vector. The transformed SDF is evaluated as:

$$\phi'(\mathbf{x}) = \lambda_{\min}(\mathbf{S}) \cdot \phi\left(\mathbf{S}^{-1}\mathbf{R}^{-1}(\mathbf{x} - \mathbf{t})\right), \qquad (16)$$

where $\lambda_{\min}(\mathbf{S})$ is the minimum singular value of $\mathbf{S}$, used conservatively to ensure a valid signed distance bound for safe sphere tracing.

During sampling and PDF evaluation, all geometry is interpreted in local (pre-transform) space. Volume sampling is performed in the unit emitter domain, then mapped to world space using the full transform. The directional PDF is corrected by the Jacobian determinant of the spatial transform:

$$p'(\omega) = \frac{1}{|\det(\mathbf{S})|} \cdot p(\omega'), \qquad (17)$$

where ω' is the direction computed in the emitter's local frame.

This formulation supports non-uniform scaling and maintains physical correctness in sampling and evaluation, while exposing artist-friendly parameters at the light source level.

5. Experiments and Results

We perform a series of experiments to evaluate our method. All experiments were conducted on a PC equipped with an Intel i7-9700K CPU, 56,GB DDR4 RAM, and an NVIDIA RTX 4070 GPU. Unless otherwise specified, all images are rendered at a resolution of 1024×1024. The comparisons are always in equal-spp manner, as the computational overhead of our method remains moderate across scenarios.

5.1. Variance Reduction

We thoroughly evaluate the variance reduction capabilities of our proposed SDF emitter sampling technique across a diverse set of volumetric light sources and 3D scenes. Our experiments demonstrate that by correctly sampling the induced solid-angle distribution, our method consistently achieves significantly lower variance compared to traditional approaches, leading to higher-fidelity renders at equivalent sampling budgets. We perform the same rendering tasks (direct lighting only, to isolate the variance introduced by emitter sampling, while a production-level comparison with global illumination is shown in Fig. 1) with different approaches. Each approach uses the same SPP=64 budget. To avoid overexposure, we suppress the emission that is directly visible to the camera. We report the error measured in mean average percentage error (MAPE) and FLIP (1). We categorize our comparisons into two groups below.

5.1.1 Analytic and Simple Implicit Shapes

For this group, we evaluate volumetric emitters defined by analytic SDFs (e.g., sphere, torus, cube). These shapes allow for a direct comparison against baselines where alternative sampling strategies might theoretically be feasible. Here, we test with torus shape which is slightly more complex than other analytic primitives. We compare three distinct sampling methods for direct lighting:

- Uniform Directional Sampling (Baseline 1): In this fundamental baseline, samples are drawn uniformly over the unit sphere from the shading point. This approach is highly general but completely ignores the light source's specific geometry and extent, leading to high variance, especially for small or distant emitters.
- Emitter Surface Sampling (Baseline 2): For shapes where a parameterized surface exists (or can be easily derived, e.g., the boundary of a sphere or cube), we sample points uniformly over the emitter's surface area. While common for surface lights, this method is fundamentally suboptimal for volumetric emitters as it only considers emission from the boundary, not the interior. It also requires explicit surface parameterization, which is often unavailable for implicit shapes.
- Our SDF Emitter Sampling Method (Analytic): Our proposed technique, which samples directly from the analytically derived solid-angle distribution induced by the volumetric SDF.

Across both analytic test scenes our solid–angle sampler attains the lowest error without extra render time. In DRAGON, MAPE drops from 0.699 (uniform) and 0.078 (surface) to **0.024**, with FLIP falling from 0.448 and 0.054 to **0.021**; the ROOM scene exhibits the same pattern. Because Monte-Carlo error scales as $1/\sqrt{N}$, these improvements translate to more than $2\times$ faster convergence than surface sampling and two orders of magnitude over uniform sampling, yet the wall-clock cost remains essentially unchanged (4–6 s per render).

5.1.2 Complex and Composite Implicit Shapes

This group focuses on volumetric emitters defined by complex or composite SDFs (e.g., Boolean operations on primitives like a fractal shape or a 3D logo). For these shapes, direct surface parameterization or efficient surface sampling is generally intractable. Here, our method provides the crucial ability to sample directly from the volume. We compare three distinct approaches:

- Uniform Directional Sampling (Naive Baseline): As in Group 1, this acts as the most general fallback, where rays are sampled uniformly over the sphere from the shading point, and contributions are evaluated only if they intersect the emitter. This method makes no attempt to target the light source and thus serves as a lower bound on efficiency.
- Bounding Volume-Based Sampling (Approximate Our Method): In this baseline, we apply our proposed sampling framework not to the exact complex SDF, but to a simple bounding volume (e.g., an axis-aligned bounding box or sphere) that encloses the emitter. While this provides a form of importance sampling, it overestimates the emissive region, leading to wasted samples and higher variance compared to using the precise SDF. This demonstrates the benefit of leveraging the exact SDF query capabilities.
- Our SDF Emitter Sampling Method (Precise): Our full approach, which robustly identifies ray-volume intersections via sphere tracing within the precise composite SDF and samples from the resulting accurate solid-angle distribution.

The results are shown in Fig. 3. For composite emitters the advantage of sampling the *exact* SDF is even more pronounced, because any mismatch between proposal and true shape wastes samples on empty space. In the KITCHEN scene the uniform baseline suffers a MAPE of 0.945, bounding–box importance sampling trims this to 0.079, but our precise sampler reaches **0.013**: a $6\times$ improvement over the bound approach, and more than $70\times$ over uniform, with FLIP dropping from 0.761 and 0.083 to **0.020**. The CORRIDOR scene shows the same pattern. Crucially, evaluating the true SDF adds only moderate runtime overhead (less than 30% s across all methods), so the extra geometric fidelity translates directly into substantial variance reduction at negligible cost.

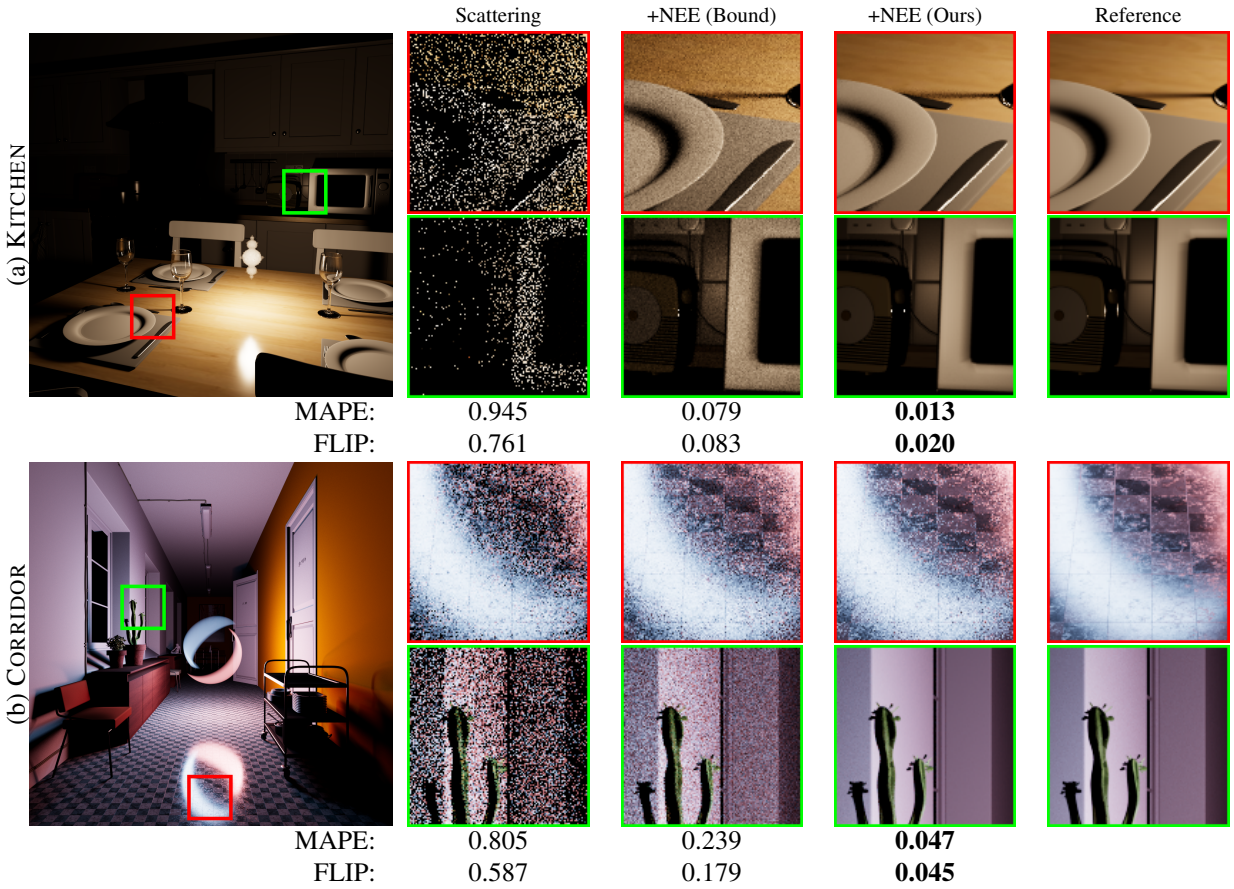

Figure 3: Comparison of different sampling techniques for complex SDF emitters (a: fractal-shaped lamp, b: logo in room). From left to right: BSDF sampling only, added Next-Event Estimation (NEE) using bounding volume approximation, NEE using our precise SDF emitter sampling, and high-sample-count reference. Our method (third column) achieves significantly lower noise and perceptual error (FLIP) than both the naive uniform and bounding volume-based approaches, demonstrating its effectiveness in targeting complex emission volumes.

5.1.3 Non-uniform Emission

As discussed, our method naturally extends to support non-uniform volumetric emission without introducing bias. In Fig. 4, we evaluate its variance performance under two types of non-uniformity. In Fig. 4 (a), the emitter uses spatially varying color with uniform intensity, while in Fig. 4 (b), the intensity varies and falls off with distance to the surface. For comparison, we also render the same scene with uniform emission.

We observe an increase in error metrics (approx. 50% higher MAPE) for non-uniform emitters compared to their uniform counterparts. This behavior is expected: our sampling PDF is derived solely from the spatial geometry of the emitter (the indicator function $v(\mathbf{x})$) and assumes a constant emission density. When the emission $\rho(\mathbf{x})$ varies spatially, especially with high-frequency noise or strong contrast, the integrand $L_e(\omega) = \int v \cdot \rho$ diverges from our geometric PDF $p(\omega) \propto \int v$.

The resulting mismatch in weights leads to increased variance. However, because our method still robustly identifies the valid solid angle subtended by the volume, it remains significantly more efficient than uniform sphere sampling, which fails to capture the emitter's boundary entirely.

5.2. Volume Estimation Accuracy

To evaluate the accuracy and practical cost of our Monte Carlo volume estimation procedure, we test a variety of representative implicit shapes with increasing complexity:

- A torus, for which the exact volume is known analytically;
- A union of two disjoint spheres;
- A procedurally defined fractal shape;
- A stylized 3D logo.

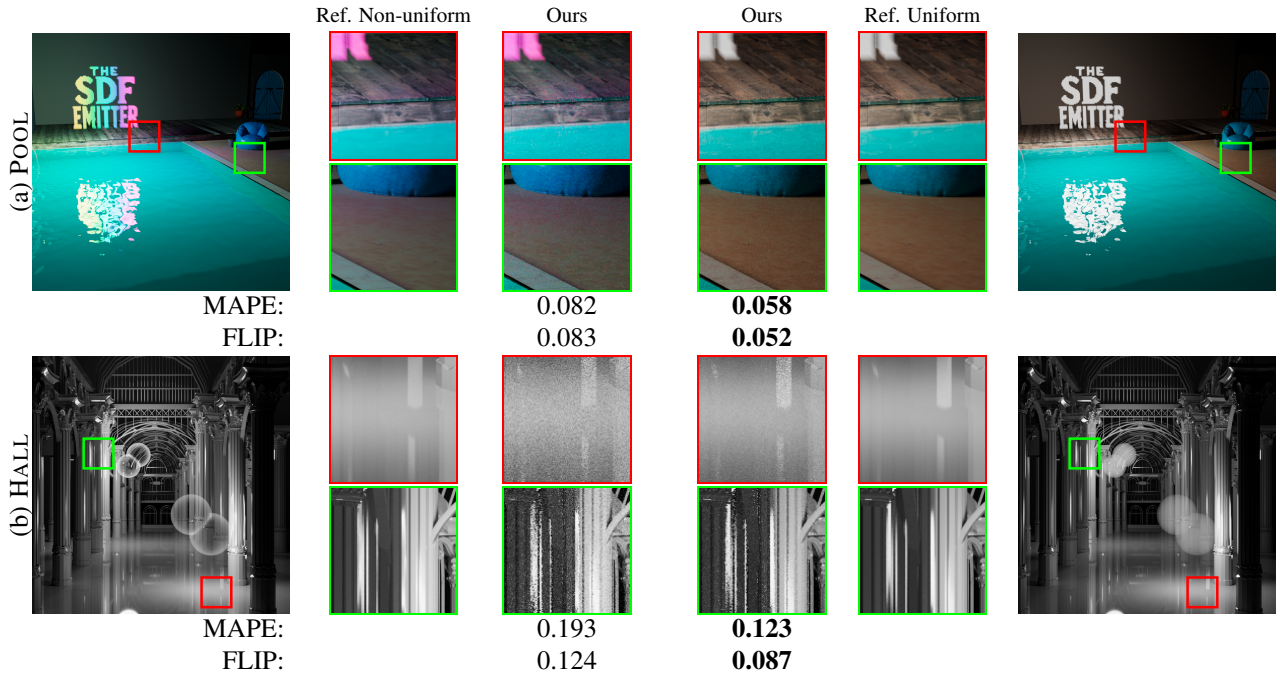

Figure 4: Comparison of rendering results using SDF emitter with non-uniform (left) versus uniform (right) volumetric emission. In each scene, we render two variants: one with uniform intensity and one with spatially varying emission. (a): an SDF extruded from 2D SDF map, with modulated emission color; (b): two Boolean combined spheres, with distance-based intensity falloff. Despite the added complexity of non-uniformity, our method maintains low error under the same sampling budget, with only a modest increase in error metrics.

For each shape, we estimate the enclosed volume by uniformly sampling points within a tight bounding box and applying the binary indicator function described in § 3. We record the estimated volume as a function of sample count, and report wall-clock timings for all evaluations.

Fig. 5 presents the convergence behavior across the four shapes. For the TORUS and TWOSPHERES cases, we compare against exact analytic volumes. For the FRACTAL and LOGO[1] cases, we estimate ground-truth values using 2^{40} samples. Despite the high sample ceiling, all shapes exhibit rapid convergence, with relative errors falling below 0.1% by approximately 2^{18} samples. This level of precision is sufficient to preserve perceptual consistency in rendered brightness across all tested scenes.

To assess runtime overhead, we also measured the total estimation time using 2^{31} samples. Across all shapes, the GPU execution time remains below 2 milliseconds. This preprocessing cost is on par with alias table construction for environment lights, and can be amortized across frames or shared across emitter instances. As a result, our method enables fast and accurate volume normalization with negligible impact on overall performance.

[1]SDF courtesy https://www.shadertoy.com/view/4sl3zn

5.3. Lighting Results

Compactness of SDF emitters. Signed–distance functions (SDFs) encode geometry implicitly, so a single analytic formula, or a tiny neural network with only a few dozen weights, can describe shapes that would require thousands of mesh triangles. Fig. 6 (a) and (b) juxtaposes a bunny emitter stored as a polygon mesh versus an SDF[2]. The rendered appearance is very similar, yet the SDF representation is two orders of magnitude smaller in memory, needs no importance sampling overhead (alias table or tree-based representation), and can be edited procedurally at virtually no extra cost.

Volumetric versus surface emission. Fig. 6(c) and (d) use a simpler luminous sphere to reveal a fundamental behavioural difference between *surface* and *volumetric* lights. For the surface mesh light, radiance is uniform across the emitting area; therefore every point on the projected disk contributes equally and the resulting shadow shows a classic, evenly blurred penumbra. By contrast, our SDF emitter is filled with uniform *volumetric* emission. A shading ray aimed through the centre of the sphere travels a longer path

[2]SDF courtesy https://www.shadertoy.com/view/3lyyWc

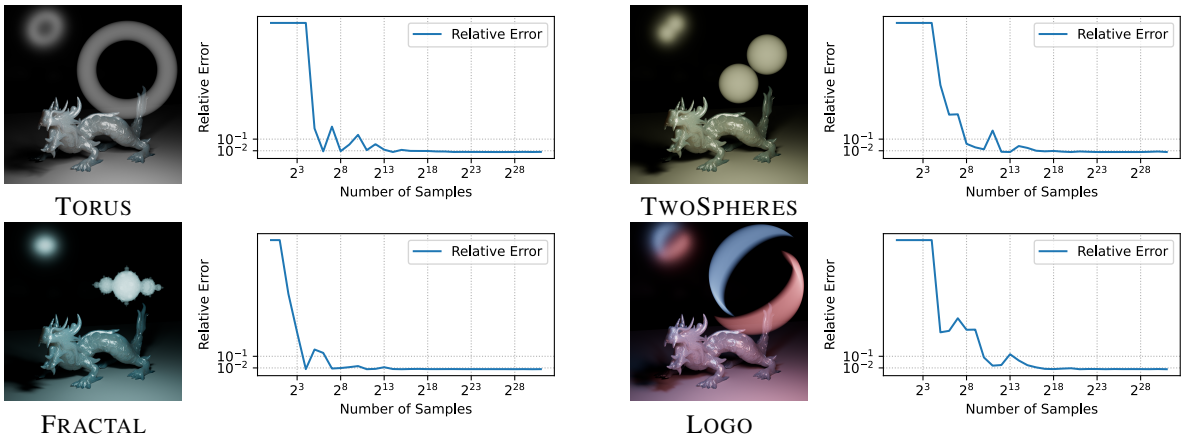

Figure 5: Convergence plots of volume estimation for four representative implicit shapes. For each case, the relative error is plotted against the number of samples used during uniform point sampling inside a bounding box. For TORUS and TWOSPHERES we calculate the analytical volume as reference, while the reference of the others are calculated via Monte Carlo estimation using significant amount of samples. Although the estimation procedure uses up to 2^{31} samples, we observe that all examples converge to less than 1% error by approximately 2^{18} samples. Our implementation calculate 2^{31} samples within 2 milliseconds, which achieves an 0.1% error. This rapid convergence demonstrates the efficiency and practical viability of our volume normalization strategy, even for highly complex or fragmented implicit geometries.

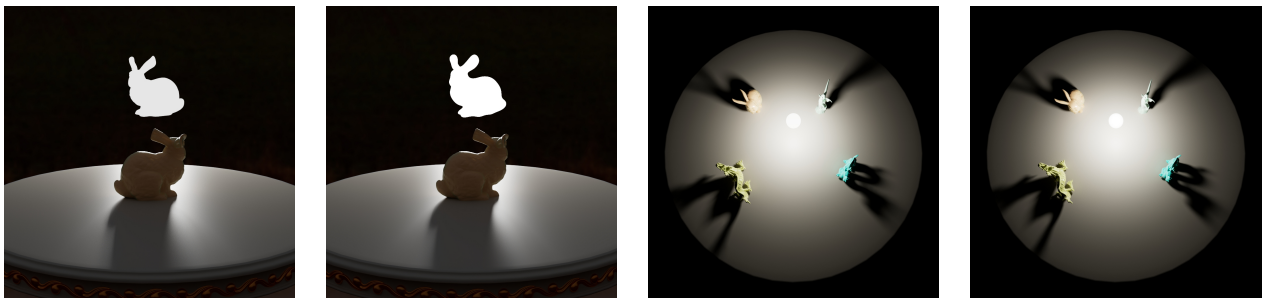

(a) Mesh emitter (b) SDF emitter (c) Surface emitter (d) SDF emitter

Figure 6: Comparison between surface and volumetric SDF emitters. (a, b) A bunny-shaped emitter represented as a triangle mesh versus an SDF. The SDF is a SIREN network (19) with dozens of weights. Both produce similar appearance. (c, d) A sphere-shaped emitter represented as an area emitter and as a volumetric SDF emitter. Despite identical geometry, the shadow sharpness differs: the volumetric SDF produces a more focused shadow due to spatially varying emission along different ray paths. This highlights the expressive lighting behavior unique to volumetric emitters. Readers are referred to supplemental material for clearer comparison.

inside the volume than one skimming the rim, so the effective radiant power varies with direction as the line-integral length $L(\omega)$. Consequently the solid-angle distribution is *denser* toward the centre, yielding a sharper umbra and a penumbra that fades more rapidly (Fig. 6 (d)). This directional modulation is captured automatically by our sampling PDF and incurs no additional cost—yet it produces shadow silhouettes that cannot be replicated by merely adjusting the intensity of a surface light.

Although SDFs could be converted to surface meshes prior to rendering, retaining the implicit representation during rendering provides practical benefits. Sampling SDF emitters avoids the construction and traversal of complex mesh-based light sampling structures and is often faster than sampling heavy mesh lights with comparable geometric detail. Moreover, SDFs naturally support volumetric emission and continuous shape manipulations. In practice, this yields lighting effects such as directionally modulated shadows and volumetric glows that surface emitters cannot reproduce by intensity reweighting alone, while remaining compact, procedural, and straightforward to integrate.

6. Discussion

6.1. Surface vs. Volumetric Emission Sampling

For diffuse area lights, sampling uniformly over the emitter's solid angle (2; 20) leads to constant Monte Carlo weights. That's because the geometry term $G = \frac{\cos\theta}{\|\mathbf{x}-\mathbf{x}_0\|^2}$ appears both in the radiance expression and in the Jacobian from solid angle to surface area. The result is fast convergence with even a few samples. However, in an emissive volume, radiance in a direction ω depends on how far that ray travels through the emitter:

$$\frac{d\Phi}{d\omega} \propto L(\omega), \quad L(\omega) = \sum_i (r_i^+ - r_i^-).$$

This means the optimal PDF should be proportional to $L(\omega)$. But uniform-solid-angle sampling doesn't account for this, causing high variance especially for thin or distant volumes. We sample a point uniformly inside the volume and project it to a direction. This yields a PDF (Eq. (13)) that increases with the cube of segment length—closely tracking $L(\omega)$ in most practical cases. When the emitter is far from the surface, our PDF approaches the ideal $p^\star(\omega) \propto L(\omega)$. Even nearby, it rejects directions that miss the volume and still favors longer paths, reducing variance. This explains the consistent speedups in Fig. 7.

6.2. Limitations and Failure Cases

While our method performs robustly across a wide variety of emitter geometries, certain classes of implicit shapes and scenarios pose challenges worth highlighting.

High-Frequency or Noisy SDFs. In Fig. 8(a), we test a procedurally modulated implicit surface with high-frequency perturbations. Such SDFs require very small sphere tracing steps to avoid overshooting during ray–volume intersection, resulting in increased traversal cost. In our current implementation, this leads to noticeably higher render times, although image quality remains consistent. We note that hierarchical SDF acceleration (e.g., grid-based min/max trees) could substantially alleviate this issue.

Low Volume-to-Bounding-Box Ratio. Emitters with sparse or thin structures incur high sampling overhead. In Fig. 8(c), a box-frame SDF fills only a small fraction of its bounding volume, causing a low acceptance rate in rejection sampling and thus higher computational cost. Tight bounding boxes and stratified sampling reduce this effect, but a more efficient sampling method remains an open problem.

Lack of Surface Parameterization. As with most implicit representations, SDFs do not provide inherent 2D surface parameterization. However, in a production rendering context, artists typically rely on UV coordinates to drive shader graphs for procedural textures, emission maps, or surface variation. Without UVs, applying fine-grained material control to the emitter surface becomes difficult, or requires ad hoc workarounds, such as baking textures into volumes.

7. Conclusion and Future Work

We introduced SDF emitters, a new class of light source in physically-based rendering, along with a practical method for efficient rendering. By turning spatial volume sampling into solid-angle importance sampling, our approach enables complex SDF-based lights to be used directly in path tracing—without meshing, discretization, or precomputed tables. The result is both unbiased and efficient, with substantial variance reduction in scenes where conventional strategies struggle.

There's room to extend this work in several directions. One is cosine-weighted sampling, which would better match typical BRDFs near surfaces. Another is incorporating solid-angle–to–projected-area corrections to account for emitter orientation. Both are standard ideas in light sampling, but adapting them to arbitrary SDFs is nontrivial, and not essential for most of the lighting tasks. Our focus has been on building a general, reliable foundation. Optimizations can follow.

References

[1] P. Andersson, J. Nilsson, P. Shirley, and T. Akenine-Möller. Visualizing Errors in Rendered High Dynamic Range Images. In *Eurographics Short Papers*, May 2021. 7

[2] J. Arvo. Stratified sampling of spherical triangles. In *Proceedings of the 22nd annual conference on Computer graphics and interactine techniques 1995*, pages 437–438, 1995. 2, 3, 11

[3] B. Bitterli, C. Wyman, M. Pharr, P. Shirley, A. Lefohn, and W. Jarosz. Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting. In *ACM Transactions on Graphics*, number 4, 2020. 3

[4] A. Conty Estevez and C. Kulla. Importance sampling of many lights with adaptive tree splitting. *Proc. ACM*, 1(2), Aug. 2018. 3

[5] C. Crassin, F. Neyret, S. Lefebvre, and E. Eisemann. Gigavoxels: ray-guided streaming for efficient and detailed voxel rendering. *Proceedings of the 2009 symposium on interactive 3D graphics and games*, pages 15–22, 2009. 3

[6] S. Diolatzis, A. Gruson, W. Jakob, D. Nowrouzezahrai, and G. Drettakis. Practical product path guiding using linearly transformed cosines. *Computer Graphics Forum (Proceedings of Eurographics Symposium on Rendering)*, 39(4):23–33, 2020. 3

[7] Z. Fan, J. Guo, Y. Wang, T. Xiao, H. Zhang, C. Zhou, Z. Chen, P. Hong, Y. Guo, and L.-Q. Yan. Specular polynomials. *ACM Trans. Graph.*, 43(4), July 2024. 3

[8] M. N. Gamito. Solid angle sampling of disk and cylinder lights. *Computer Graphics Forum*, 35(4):25–36, 2016. 2, 3

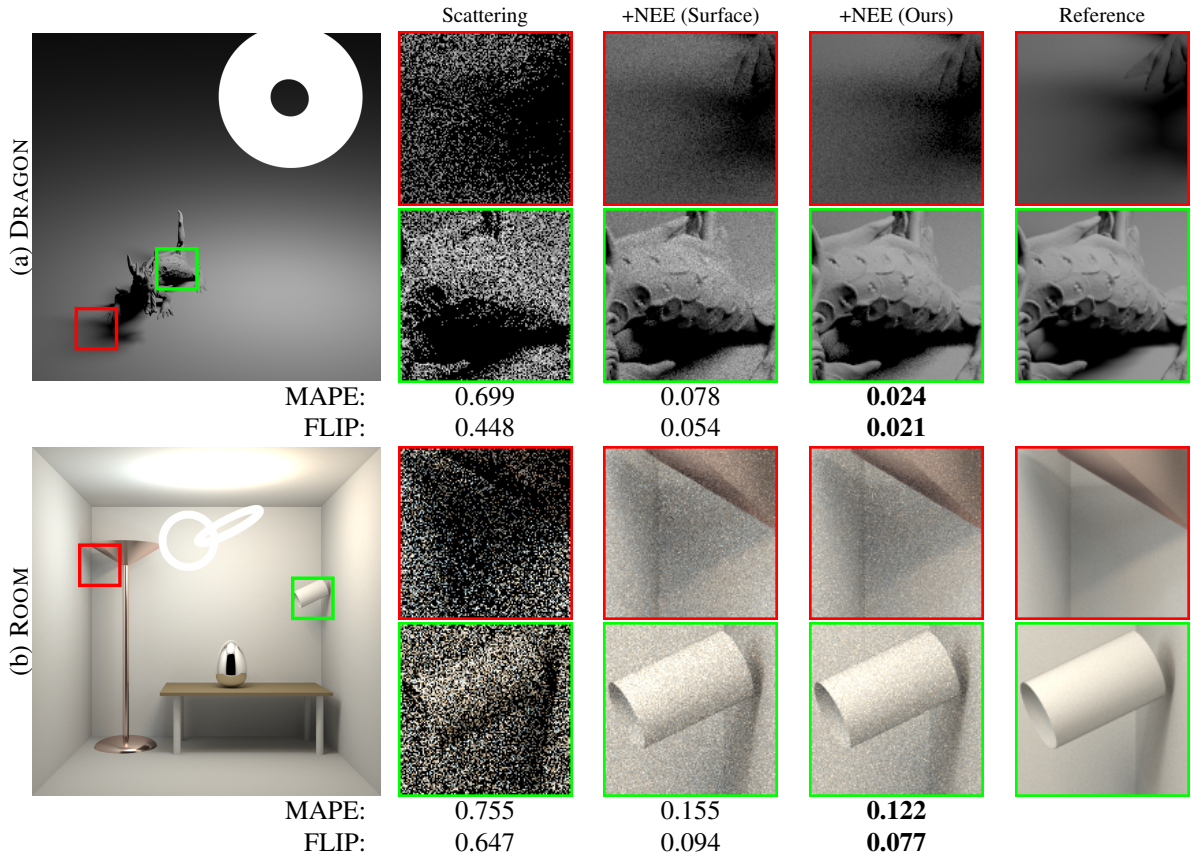

Figure 7: Comparison of rendering quality for analytical SDF emitters, using three sampling strategies: uniform scattering (baseline), next-event estimation (NEE) with surface-area sampling, and NEE with our solid-angle sampling. NEE methods are combined with scattering via multiple importance sampling (MIS). The naive approach samples directions by uniformly sampling the SDF surface and projecting to solid angle, which ignores the interior structure. In contrast, our method enables analytically exact directional sampling over the volume. We report MAPE and FLIP (lower is better).

(a) (b) (c) (d)

Figure 8: Limitations of our method in challenging SDF scenarios. (a–b): High-frequency perturbations in the SDF (a) require much finer sphere tracing steps than a smooth surface (b), leading to a significant increase in render time (60s vs. 12s). (c–d): Emitters with low volume-to-bounding-box ratios, such as a thin frame (c), suffer from inefficient rejection sampling compared to solid volumes (d), doubling the computation time (20s vs. 10s).

[9] J. Hanika, M. Droske, and L. Fascione. Manifold next event estimation. *Computer Graphics Forum*, 34(4):87–97, 2015.

[10] J. C. Hart. Sphere tracing: A geometric method for the an-

tialiased ray tracing of implicit surfaces. *The Visual Computer*, 12:527–545, 1996. 5

[11] E. Heitz, J. Dupuy, S. Hill, and D. Neubelt. Real-time polygonal-light shading with linearly transformed cosines. *ACM Transactions on Graphics*, 35(4):1–8, 2016. 3

[12] Y. Jiang, D. Ji, Z. Han, and M. Zwicker. Sdfdiff: Differentiable rendering of signed distance fields for 3d shape optimization. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1251–1261, 2020. 2

[13] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 2019. 2

[14] P. Moreau, M. Pharr, and P. Clarberg. Dynamic many-light sampling for real-time ray tracing. In *Proceedings of High Performance Graphics 2019*. Eurographics Association, 2019. 3

[15] Y. Ouyang, S. Liu, M. Kettunen, M. Pharr, and J. Pantaleoni. Restir gi: Path resampling for real-time path tracing. *Computer Graphics Forum*, (8):17–29, 2021. 3

[16] Y. Ouyang, S. Liu, M. Kettunen, M. Pharr, and J. Pantaleoni. Restir gi: Path resampling for real-time path tracing. Number 8, pages 17–29, 2021. 3

[17] F. Simon, J. Hanika, T. Zirr, and C. Dachsbacher. Line integration for rendering heterogeneous emissive volumes. *Computer Graphics Forum*, 36(4):101–110, 2017. 3

[18] V. Sitzmann, J. N. Martel, A. W. Bergman, D. B. Lindell, and G. Wetzstein. Implicit neural representations with periodic activation functions. *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pages 7462–7473, 2020. 2

[19] V. Sitzmann, J. N. Martel, A. W. Bergman, D. B. Lindell, and G. Wetzstein. Implicit neural representations with periodic activation functions. In *Proc. NeurIPS*, 2020. 10

[20] C. Ureña, M. Fajardo, and A. King. An area-preserving parametrization for spherical rectangles. *Computer Graphics Forum*, 32(4):59–66, 2013. 2, 3, 11

[21] E. Veach. *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, 1997. 3

[22] R. Villemin and C. Hery. Practical illumination from flames. *Journal of Computer Graphics Techniques*, 2(2):142–155, 2013. 3

[23] B. Walter, S. Fernandez, A. Arbree, K. Bala, M. Donikian, and D. P. Greenberg. Lightcuts: A scalable approach to illumination. *ACM Transactions on Graphics*, 24(3):1098–1107, 2005. 3

[24] P. Wang, L. Liu, Y. Liu, C. Theobalt, T. Komura, and W. Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *Proceedings of the 35th International Conference on Neural Information Processing Systems*, pages 27171–27183, 2021. 2

[25] L. Yariv, J. Gu, Y. Kasten, and Y. Lipmen. Volume rendering of neural implicit surfaces. *arXiv preprint arXiv:2106.12052*, pages 4805–4815, 2021. 2

[26] T. Zeltner, I. Georgiev, and W. Jakob. Specular manifold sampling for rendering high-frequency caustics and glints. *ACM Trans. Graph.*, 39(4), Aug. 2020. 3

[27] S. Zheng, Z. Zhou, X. Chen, D. Yan, C. Zhang, Y. Geng, Y. Gu, and K. Xu. Luisarender: A high-performance rendering framework with layered and unified interfaces on stream architectures. *ACM Trans. Graph.*, 41(6), Nov. 2022. 5

[28] J. Zhu, Y. Bai, Z. Xu, S. Bako, E. Velázquez-Armendáriz, L. Wang, P. Sen, M. Hašan, and L.-Q. Yan. Neural complex luminaires: Representation and rendering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 40(4), 2021. 3