EUROGRAPHICS 2025 / A. Bousseau and A. Dai
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 44 (2025), Number 2

Linearly Transformed Spherical Distributions for Interactive Single Scattering with Area Lights

Aakash KT[1], Ishaan Shah[1], P. J. Narayanan[1]

[1]International Institute Of Information Technology, Hyderabad (IIIT-H)

Figure 1: *We present a semi-analytic method for interactive single-scattering in homogeneous media with polygonal area lights. Our method achieves biased but noise-free renderings and better performance compared to Volume-RIS [LWY21]. This figure shows equal-time renderings of two scenes, Sponza (left) with four area lights and a spherical light shaft (right) with two area lights inside the sphere. Our core unshadowed method (left) is completely noise-free, while Volume-RIS exhibits noise especially near the area light. Our method with the ratio estimator formulation for shadows (right) can handle volumetric shadows due to light shafts, rendering plausible light beams. Note that the we denoise the entire image for Volume-RIS and only shadows of our method for the figure on the right.*

Abstract

Single scattering in scenes with participating media is challenging, especially in the presence of area lights. Considerable variance still remains, in spite of good importance sampling strategies. Analytic methods that render unshadowed surface illumination have recently gained interest since they achieve biased but noise-free plausible renderings while being computationally efficient. In this work, we extend the theory of Linearly Transformed Spherical Distributions (LTSDs) which is a well-known analytic method for surface illumination, to work with phase functions. We show that this is non-trivial, and arrive at a solution with in-depth analysis. This enables us to analytically compute in-scattered radiance, which we build on to semi-analytically render unshadowed single scattering. We ground our derivations and formulations on the Volume Rendering Equation (VRE) which paves the way for realistic renderings despite the biased nature of our method. We also formulate ratio estimators for the VRE to work in conjunction with our formulation, enabling the rendering of shadows. We extensively validate our method, analyze its characteristics and demonstrate better performance compared to Monte Carlo single-scattering.

CCS Concepts

• ***Computing methodologies*** → ***Ray tracing;*** *Reflectance modeling;*

1. Introduction

Monte Carlo (MC) path tracing supported by GPU hardware has gained a lot of interest in offline and real-time rendering. The combination of state-of-the-art importance sampling, hardware accelerated ray tracing and efficient denoising has facilitated high-quality direct lighting with a small number of samples per pixel (spp), even with area lights. The visible noise from MC however increases in the presence of participating media. Monte Carlo direct lighting or single scattering in participating media with area lights remains challenging at low spp which is necessary for interactive and real-time applications. Beyond MC methods, analytic methods that can plausibly approximate direct lighting hold promise in many situations. These methods are typically combined with MC resulting in overall reduced noise, while being computationally efficient. The reduction in noise or lower variance comes at the cost of bias, which is often tolerable, especially for real-time applications.

The most well-known analytic method for direct lighting with area lights uses Linearly Transformed Cosines (LTCs), which are a type of Linearly Transformed Spherical Distributions (LTSDs) [HDHN16]. LTCs are used to analytically approximate unshadowed surface illumination from area lights and combined with MC ratio estimators [HHM18] to account for the shadows. The combination of LTCs with ratio estimators preserves high-frequency details when used with denoising, as opposed to the denoising of MC direct lighting which suffers from over-blurring. While the LTC formulation and method can be used for direct lighting in surface-based scenes, it cannot be used for single scattering in participating media.

In this paper, we develop a new formulation based on LTSDs to analytically render unshadowed single scattering in infinite homogeneous participating media with polygonal area lights. At the core of single scattering is a spherical in-scattering integral and a spherical surface integral. These integrals have a form that is similar to the surface rendering equation where LTCs are applicable. However, they differ in two aspects: (1) the presence of an additional transmittance term and (2) usage of phase functions instead of Bi-directional Reflectance Distribution Functions (BRDF) in the in-scattering integral, making the application of LTSDs non-trivial. We present a simple and practical approximation of the transmittance term that can be analytically computed at render time. While this approximation already enables analytic evaluation of the spherical surface integral, the presence of a phase function in the in-scattering integral makes such a direct application non-trivial. We analyze LTSDs and determine the exact reasons that prevent their application to phase functions and propose a new and efficient LTSD precomputation approach.

An analytic solution to the in-scattering integral is however not sufficient. This integral is nested in a line integral which forms the Volume Rendering Equation (VRE). Analytic evaluation of the VRE is not straightforward due to this nesting and the fact that the in-scattering integral can only be evaluated at individual points on the media. We show that the line integral can instead be evaluated using quadrature rules. Although not analytic, they are well behaved across pixels thanks to their deterministic nature and produce noise-free renderings. Finally, we also formulate ratio estimators for the line and surface integral to render shadows in a separate denoised pass. Since only the shadow pass is denoised, renderings preserve high-frequency details, similar to the combination of LTCs with ratio estimators for surface rendering.

We comprehensively validate our method against Monte Carlo single-scattering references and show that it produces high-quality plausible single scattering renderings. We also compare our method with Volume-RIS [LWY21] for unshadowed and shadowed configurations of both methods. Figure 1 shows two scenes rendered using our method compared to renderings of Volume-RIS. The left section of this figure shows unshadowed renderings where our method is completely noise-free. The right section shows shadowed renderings using denoised ratio estimators on a light shaft scene, demonstrating that our approach is capable of rendering volumetric shadows.

In summary, the following are the contributions of this paper:

- A new LTSD fitting approach for phase functions which paves the way for analytic evaluation of the in-scattering integral.
- A noise-free semi-analytic formulation of the unshadowed Volume Rendering Equation (VRE).
- Ratio-estimator formulation of the VRE to handle shadows.

2. Related Work

In this section we discuss related works which are grouped by the approach used to solve the VRE.

2.1. Stochastic Methods

Monte-Carlo integration can naively be applied to solve the VRE. However, better importance sampling strategies are needed to ensure better convergence. Sampling proportional to the transmittance [PJH16] is a good strategy, however it disregards the contribution of the light term which increases variance. Kulla et al. [KF12] propose a method of sampling the initial distance which considers the incoming light from a point light source. Villeneuve et al. [VGGN21] further propose a method to sample proportional to the product of transmittance, the phase function and the orientation of a point light source. Both of the above methods can be applied to estimate the illumination due to a polygonal light source by randomly sampling a point on the light source at the cost of additional variance. Instead of sampling proportional to the product of phase function and light one can use Resampled Importance Sampling (RIS) [TCE05]. RIS operates by generating a set of random initial candidates and chooses one of them by weighing them according to the target PDF. In the case of single scattering, one can sample the distance with a simple sampling strategy (for eg. proportional to the transmittance) and use RIS to obtain a sample which is distributed according to the VRE. RIS is at the core of Volume-ReSTIR [LWY21] which is the state of the art method to render volumes.

2.2. Analytic Methods

Previous work in this area has mainly focused on obtaining closed-form solutions to single scattering in a homogeneous medium under the influence of punctual lights. Sun et al. [SRNN05] propose a combination of analytic evaluation and precomputation to achieve real-time rendering of single scattering under isotropic point lights. They also extend their formulation to render single scattering from distant complex illumination such as environment and achieve real-time frame rates. However, the interpolation of the precomputed data introduces artifacts in the rendering. Pegoraro et al. [PP09] propose an analytical solution for isotropic phase functions which does not rely on any pre-computation and storage and produces more accurate results. However it is not suited for real-time applications. Further efforts to optimize the method [PSP09b, PSS11] and relax the condition of isotropic phase functions and lighting [PSP09a, PSP10] were also made. It is important to note that all of the analytic methods do not support shadows and additional methods need to be employed to add shadow information.

Our method works in a similar fashion, in that the core method itself does not support shadows and we show how ratio estimators can be formulated and combined for shadow support. We would like to point out that any other method that can render volumetric shadows [WR08, Wym11, BSA10, CBDJ11] could potentially be used as well. Like the works of [ED10, KSE14], our method is based on the principles of volumetric scattering (VRE) and thus produces plausible renderings. Finally, similar to the usage of LTCs for better unbiased importance sampling of area lights for surface based scenes [Pet21, SKN23], our LTSD formulation could be used for better sampling in participating media instead.

3. Preliminaries

We begin with a recap on the preliminaries of volume rendering which also helps establish the notation of our paper.

The outgoing radiance L at a point $\mathbf{y}$ in the direction ω_o in the presence of homogeneous participating media is given by the Volume Rendering Equation (VRE) [NGHJ18] as:

$$L(\mathbf{y},\omega_o) = \mu_s \int_0^x T(\mathbf{y},\mathbf{z})L_s(\mathbf{z},\omega_o)dz + T(\mathbf{y},\mathbf{x})L_c(\mathbf{x},\omega_o), \quad (1)$$

where μ_s is a spatially constant scattering coefficient. We refer to the first term in the above equation as the air-light integral, inspired by Sun et al. [SRNN05]. In the following, we will assume that the scene contains one area light A with $A(\mathbf{x})$ denoting the solid angle it subtends on the unit sphere around $\mathbf{x}$. Our method trivially extends to multiple area lights by summing up individual contributions due to the linearity of light transport.

The point $\mathbf{z} = \mathbf{y} - z\omega_o$ is a point in the medium and L_s is its in-scattered radiance. With single scattering, we need to explicitly account for the visibility and the transmittance in L_s with the integration domain over the solid angle of the area light:

$$L_s(\mathbf{z},\omega_o) = L_i \int_{A(\mathbf{z})} \rho(\mathbf{z},\omega_o,\omega_i)T_v(\mathbf{z},\omega_i)d\omega_i, \quad (2)$$

where ρ is the media's phase function which describes forward or back scattering. Similarly, the point $\mathbf{x} = \mathbf{y} - x\omega_o$ lies on a surface and its radiance L_c towards ω_o is given by:

$$L_c(\mathbf{x},\omega_o) = L_i \int_{A(\mathbf{x})} f(\mathbf{x},\omega_o,\omega_i)T_v(\mathbf{x},\omega_i)|n_\mathbf{x}\cdot\omega_i|d\omega_i. \quad (3)$$

where f is the Bi-Directional Reflectance Distribution Function (BRDF) describing the surface's reflectance and $n_\mathbf{x}$ is the normal vector at $\mathbf{x}$. In both equations, L_i denotes the spatially constant incoming radiance from the area light A and T_v is defined as:

$$T_v(\mathbf{x},\omega_i) = T(\mathbf{x},\mathbf{t}(\mathbf{x},\omega_i))V(\mathbf{x},\omega_i), \quad (4)$$

where $\mathbf{t}(\mathbf{x},\omega_i)$ is the ray-casting function that returns the intersection point of a ray towards ω_i from $\mathbf{x}$. The visibility $V = 1$ if the area light A is visible from $\mathbf{x}$ in direction ω_i or $V = 0$ otherwise.

The function T in Eq. 1 and Eq. 4 gives the transmittance (or attenuation) between two points in a homogeneous medium:

$$T(\mathbf{y},\mathbf{x}) = e^{-\mu_t||\mathbf{y}-\mathbf{x}||_2}. \quad (5)$$

3.1. Linearly Transformed Spherical Distributions

Linearly Transformed Cosines or LTCs are a class of Linearly Transformed Spherical Distributions (LTSD) as defined by Heitz et al. [HDHN16]. An LTSD is defined by a matrix M that maps a source distribution D_o to a target distribution D as:

$$D(\omega) = D_o(\omega')\frac{\partial\omega'}{\partial\omega} = D_o\left(\frac{M^{-1}\omega}{||M^{-1}\omega||}\right)\frac{M^{-1}\omega}{||M^{-1}\omega||^3} \quad (6)$$

Heitz et al. [HDHN16] also establish the following equivalence for integral of a LTSD over an arbitrary solid angle P:

$$\int_P D(\omega)d\omega = \int_{P_o} D_o(\omega')d\omega', \quad (7)$$

where $P_o = M^{-1}P$ is the transformation of the original solid angle by the LTSD matrix M.

LTCs use the clamped cosine distribution for $D_o(\omega') = \frac{1}{\pi}\max(n\cdot\omega', 0)$ and approximate $D(\omega) \approx f(\omega_o,\omega)|n\cdot\omega|$ i.e. the BRDF times cosine target. In this setting, the unshadowed surface radiance can be analytically approximated as:

$$\begin{aligned} L_c(\mathbf{x},\omega_o) &= L_i \int_{A(\mathbf{x})} f(\mathbf{x},\omega_o,\omega_i)|n_\mathbf{x}\cdot\omega_i|d\omega_i \\ &\approx L_i \int_{A(\mathbf{x})} D(\omega_i)d\omega_i \\ &\approx L_i \int_{A_o(\mathbf{x})} D_o(\omega_o)d\omega_o = L_i E(\mathbf{x}, A_o(\mathbf{x})), \end{aligned} \quad (8)$$

where $A_o(\mathbf{x}) = M^{-1}A(\mathbf{x})$ and E represents the cosine irradiance expression [BRW89].

LTC matrices are precomputed in a table for different parameter combinations of the target distribution and fetched during rendering. For an isotropic BRDF, a two-dimensional table parameterized by (α, θ) is precomputed, where α is the BRDF's roughness and θ is the elevation of ω_o. For an anisotropic BRDF, this table is four-dimensional parameterized by (α_x, α_y, θ, ϕ) [KHDN22], where similarly α_x, α_y and BRDF roughness in x and y and θ, ϕ are the spherical angles are of ω_o.

4. Semi-Analytic Single Scattering

Our goal is to develop a semi-analytic method to render unshadowed single-scattering in homogeneous media with polygonal area lights. Achieving this goal involves deriving expressions for (1) unshadowed surface radiance L_c, (2) unshadowed in-scattered radiance L_s and (3) the air-light integral.

To this end, we first derive a practical closed-form approximation of the transmittance at a point towards an area light (Sect. 4.1). This expression is useful to derive analytic expressions for L_c (Sect. 4.2) and L_s (Sect. 4.3) using appropriate LTSDs. Finally, in Sect. 4.4, we describe the usage of quadrature rules that use analytic expressions of L_c and L_s to evaluate the air-light integral.

It is straightforward to apply LTSDs for analytic surface radiance as we show in Sect. 4.2, which is not the case for analytic in-scattered radiance due to the usage of phase functions. In Sect. 5, we discuss the challenges in fitting LTSDs to phase functions, identify the core problem and propose our solution that is used in Sect. 4.3.

We ignore the visibility V in our derivations of analytic L_c and L_s. In Sect. 6, we formulate ratio estimators to take visibility into account.

4.1. Aggregated transmittance towards an area light

The aggregated transmittance $\overline{T}$ from a point $\mathbf{x}$ towards an area light A can be written as an integral of the point trasmittance (Eq. 5) over the area light's solid angle $A(x)$:

$$\overline{T}(\mathbf{x}) = \int_{A(\mathbf{x})} e^{-\mu_t ||\mathbf{t}(\mathbf{x},\omega_i)-\mathbf{x}||_2} d\omega_i. \tag{9}$$

To find an analytic approximation of the above equation, we take the exponent out of the integral as follows:

$$\overline{T}(\mathbf{x}) = e^{-\mu_t d} \int_{A(\mathbf{x})} d\omega_i = e^{-\mu_t d} A(\mathbf{x}), \tag{10}$$

where d is a statistical quantity related to A. We use $d = \min(A)$. This expression is a reasonable closed-form approximation, and requires only the minimum distance to the light and the solid angle subtended by it. As such, Eq. 10 can be computed on the fly at rendering time.

4.2. Analytic surface radiance

Applying LTCs to obtain an analytic expression for L_c seems intuitive at first glance, due to similar forms of Eq. 3 and Eq. 8. However, LTCs can only be applied in the absence of the transmittance term. We thus split the transmittance out as a separate integral, similar to the split sum strategy used in image-based lighting [Deb05]:

$$L_c(\mathbf{x},\omega_o) = L_i \int\limits_{A(\mathbf{x})} T(\mathbf{x},\omega_i) d\omega_i \int\limits_{A(\mathbf{x})} f(\mathbf{x},\omega_o,\omega_i) |n_{\mathbf{x}} \cdot \omega_i| d\omega_i. \tag{11}$$

In the above equation, the integral on the right is now in a form where LTCs can be applied. Using also the result from Eq. 10, we get the expression for analytic surface radiance $\overline{L_c}$:

$$\overline{L_c}(\mathbf{x},\omega_o) = L_i \overline{T}(\mathbf{x}) E(\mathbf{x}, A_o(\mathbf{x})), \tag{12}$$

where E is the cosine irradiance.

Figure 2: *Equal-time (~ 40 ms) comparison of our semi-analytic unshadowed single scattering with Volume-RIS. Thanks to our use of analytic in-scattered and surface radiance (Eqs. 4.3, 4.2) and deterministic quadrature for the air-light integral (Eq. 15), we obtain noise-free renderings.*

4.3. Analytic in-scattered radiance

To derive an analytic expression for the in-scattered radiance, we similarly split transmittance out of Eq. 2 and use the result from Eq. 10:

$$L_s(\mathbf{z},\omega_o) = L_i \overline{T}(\mathbf{z}) \underbrace{\int_{A(\mathbf{z})} \rho(\mathbf{z},\omega_o,\omega_i) d\omega_i}_{\text{Phase Function Integral}}. \tag{13}$$

Define S to be the analytic expression of the phase function integral, we get the final expression for analytic in-scattered radiance $\overline{L_s}$:

$$\overline{L_s}(\mathbf{z},\omega_o) = L_i \overline{T}(\mathbf{z}) S(\mathbf{z}; A(\mathbf{z})). \tag{14}$$

S approximates the phase function integral by fitting LTSD matrices to the phase function. We show in Sect. 5 that the choice of source distribution D_o to compute these matrices is non-trivial: we cannot simply use a clamped cosine distribution, like done in LTCs.

4.4. Air-light integral

We now turn our attention to the air-light integral in Eq. 1. We observe that in the presence of the analytic in-scattered radiance $\overline{L_s}$, the air-light integral is a simple one-dimensional integral. In such a low dimensional setting, it is beneficial to use quadrature rules for better convergence.

We thus approximate the outgoing radiance in Eq. 1 using the Riemann sum, a formulation used also in ray-marching. With N partitions of the air-light integral, noise-free unshadowed single scattering $\overline{L}$ is computed as:

$$\overline{L}(\mathbf{y},\omega_o) = \underbrace{\mu_s \sum_{i=1}^{N} T(\mathbf{y},\mathbf{z}_i^*) \overline{L_s}(\mathbf{z}_i^*,\omega_o) \Delta z}_{\overline{\mathcal{A}}} + T(\mathbf{y},\mathbf{x}) \overline{L_c}(\mathbf{x},\omega_o), \tag{15}$$

where $\overline{\mathcal{A}}$ is the semi-analytic solution to the air-light integral, $\Delta z = z_i - z_{i-1}$ and $z_i^* \in [z_{i-1}, z_i]$. An added benefit of using quadrature rules is that they do not produce noise in the final renderings, thanks to their deterministic nature. Fig. 2 shows an equal-time comparison with Volume-RIS, demonstrating the noise-free nature of our method.

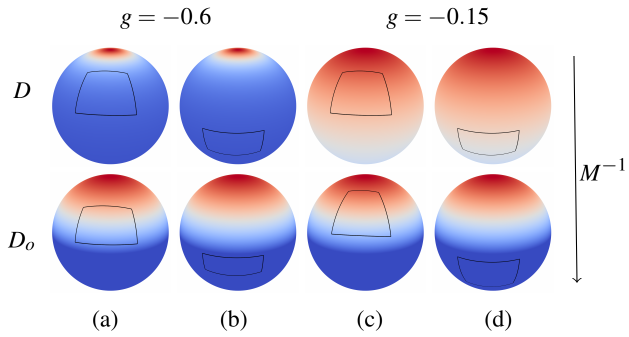

Figure 3: *Approximating HG target ($D(\omega) = \rho(\omega_o, \omega)$, top row) with a clamped-cosine source ($D_o(\omega') = max(n \cdot \omega', 0)$, bottom row) with a LTSD matrix M. (d) A solid angle with non-zero integral in D is transformed to have zero integral in D_o, demonstrating that fitting to clamped-cosine fails for near isotropic values of g.*

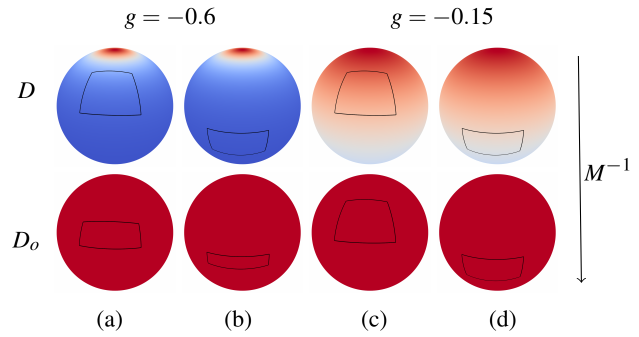

Figure 4: *Approximating HG target ($D(\omega) = \rho(\omega_o, \omega)$, top row) with a uniform sphere source ($D_o(\omega') = 1/4\pi$, bottom row) with a LTSD matrix M. (b) A solid angle with a zero integral in D is transformed to have non-zero integral in D_o, demonstrating that fitting to uniform sphere fails for highly directional values of g.*

5. Fitting Linearly Transformed Spherical Distributions to Phase Functions

In this section, we describe our approach to fit LTSDs to phase functions. The resultant LTSD matrices are used to compute S for the analytic in-scattered radiance (Eq. 14).

The equation for L_s has a similar form to Eq. 8 albeit with a major difference: the BRDF is replaced by a phase function. In our work, we use the Henyey-Greenstein (HG) [HG41] phase function which has the following form:

$$\rho(\mathbf{x}, \omega_o, \omega_i) = \rho(\cos\theta) = \frac{1}{4\pi} \frac{1 - g^2}{(1 + g^2 - 2g\cos\theta)^{3/2}}, \tag{16}$$

where $\cos\theta = (\omega_o \cdot \omega_i)$ and $g \in [-1, 1]$ is the asymmetry parameter controlling the scattering (forward or back scattering). Unlike BRDFs, phase functions are defined on the entire sphere.

We work in a co-ordinate frame where ω_o is aligned with the z-axis. In this setting, the only parameter that controls the phase function's shape is g. We thus need to compute a one-dimensional table of matrices parameterized by g and approximate $D(\omega) \approx \rho(\omega_o, \omega)$ for the target. In the next subsections, we discuss different choices of the source distribution D_o to compute these matrices and their failure cases.

5.1. Clamped-cosine distribution for D_o

Fig. 3 visualizes the usage of clamped cosine distribution on the upper hemisphere for D_o. The top row of this figure shows the distribution D with a solid angle defined by a black boundary. The bottom row shows this solid angle and the distribution transformed with the corresponding LTSD matrix. While a clamped-cosine source distribution works well for highly directional lobes of Henyey-Greenstein, it fails as its lobes approach isotropic. The failure case is demonstrated at $g = -0.15$ when the solid angle lies in the lower hemisphere of D with a non-zero integral (Fig. 3(d)). After transformation, this solid angle still remains lower and will get clipped, giving a value of zero. This gives incorrect analytic integration for near isotropic values of g. Fig. 5 (b), top row shows the failure case of using LTSD matrices fitted to clamped-cosine for rendering, demonstrating the inability to render isotropic scattering.

5.2. Uniform spherical distribution for D_o

Next, consider using the uniform spherical distribution for D_o as shown in Fig. 4, which follows the layout of Fig. 3. This choice

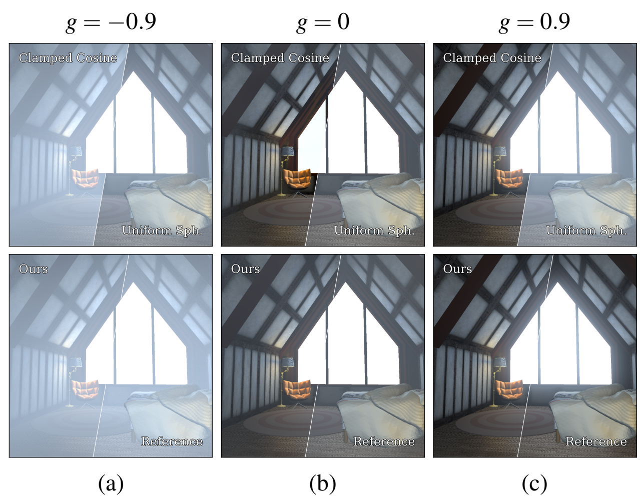

Figure 5: *Renderings using our approach with LTSD matrices fitted to clamped-cosine source and uniform spherical source (top row). The bottom row shows renderings with our method and the proposed fitting approach in Sect. 5. Fitting to clamped cosine source is unable to render isotropic scattering ((b), top row) while fitting to uniform spherical source is unable to render directional scattering ((a), (c), top row). Our fitting approach renders all effects accurately, in accordance to reference ray-tracing (bottom row).*

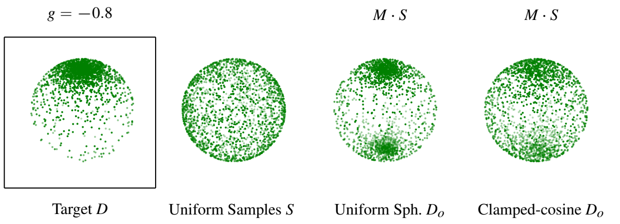

(a) (b) (c) (d)

Figure 6: *Fitting LTSDs to the Henyey-Greenstein phase function at $g=-0.8$ with uniform spherical and clamped-cosine source distributions. We randomly sample and visualize directions from (a) Henyey-Greenstein distribution, (b) uniform spherical distribution S. In (c), we visualize samples S transformed by a LTSD matrix fitted to uniform spherical source and in (d), we visualize S transformed by a LTSD matrix fitted to clamped cosine source. LTSD transforms are symmetric across the horizon and cannot capture a asymmetric distribution over the sphere.*

fails in the opposite case: at highly directional lobes of Henyey-Greenstein, and works well for near isotropic lobes. For example, at $g=-0.6$ with the solid angle in the lower hemisphere (Fig. 4(b)). The true integral has a value of zero, but after transformation, it evaluates to a non-zero value. In effect, this choice fails to capture forward and back scattering of phase functions. The result of rendering with matrices fitted to uniform spherical source is shown in Fig. 5(a), (c) top row, demonstrating that forward and back scatter effects are not reproduced.

5.3. Symmetric nature of Linearly Transformed Spherical Distributions

The core issue is that LTSDs are *symmetric* across the horizon plane. More generally, these transformations are point-symmetric across the origin; however in our case, the rotational symmetry of Henyey-Greenstein results in plane symmetry. The symmetric nature is illustrated in Fig. 6 by fitting different source distributions D_o to a HG target D with $g=-0.8$. The figure also visualizes directions S sampled uniformly on the sphere (Fig. 6 (b)). Fig. 6 (c) shows S transformed by a LTSD matrix fitted to uniform sphere source and Fig. 6 (d) similarly shows S transformed by a LTSD matrix fitted to clamped cosine source. Note that symmetry exists between the upper and lower hemisphere, and the transformed distribution in (c) and (d) do not match the target. This is precisely the challenge in fitting LTSDs to target distributions defined on the entire sphere, as they cannot capture asymmetric distributions on it.

5.4. Fitting to symmetric distributions on the unit sphere

Drawing from the analysis above, we propose the following. For a given value of g, we fit two LTSD matrices M_u and M_l. M_u transforms samples of HG in the upper hemisphere while M_l transforms samples of HG in the lower hemisphere. Effectively, we split the phase function integral in Eq. 13 as:

$$\int_{A^u(\mathbf{z})} \rho(\mathbf{z},\omega_o,\omega_i)d\omega_i + \int_{A^l(\mathbf{z})} \rho(\mathbf{z},\omega_o,\omega_i)d\omega_i, \tag{17}$$

where $A^u(\mathbf{z})$ is the solid angle in the upper hemisphere and $A^l(\mathbf{z})$ is the solid angle in the lower hemisphere. Given this, we can compute S in Eq. 13 using M_u and M_l as:

$$S(\mathbf{z},A(\mathbf{z})) = \int_{A^u_o(\mathbf{z})} D^u_o(\omega')d\omega' + \int_{A^l_o(\mathbf{z})} D^l_o(\omega')d\omega', \tag{18}$$

where $A^u_o(\mathbf{z}) = M_u A^u(\mathbf{z})$ and $A^l_o(\mathbf{z}) = M_l A^l(\mathbf{z})$.

The source distribution to which these matrices are fitted depends on the value of g:

$$\text{for } M_u: \quad D^u_o(\omega') = \begin{cases} \max(\omega'\cdot n, 0), & \text{if} \quad g<0 \\ 1-\max(\omega'\cdot n, 0), & \text{if} \quad g\geq 0 \end{cases},$$

$$\text{for } M_l: \quad D^l_o(\omega') = \begin{cases} 1-\max(-\omega'\cdot n, 0), & \text{if} \quad g<0 \\ \max(-\omega'\cdot n, 0), & \text{if} \quad g\geq 0 \end{cases}, \tag{19}$$

where $n=(0,0,1)$. In effect, we fit to a clamped cosine source if the HG peak is in the upper hemisphere for M_u and if not, we fit to a inverse clamped cosine source. For M_l we fit the reverse. We found that renderings for forward and back scattering are better when using inverse clamped cosine to fit the hemisphere where the HG peak is not present.

The integral over an arbitrary solid angle of the clamped cosine is well known and used above in Eq. 8, Eq. 12 (cosine irradiance). Similarly, the integral over an arbitrary solid angle A_o of the inverse clamped cosine distribution is:

$$\begin{aligned} \int_{A_o} (1-\max(n\cdot\omega',0))d\omega' &= \int_{A_o} d\omega' - \int_{A_o} \max(n\cdot\omega',0)d\omega' \\ &= A_o - E(A_o), \end{aligned} \tag{20}$$

that is, the cosine irradiance subtracted from the solid angle.

Combining Eq. 18, Eq. 19 and Eq. 20 we obtain an analytic expression for S:

$$S(\mathbf{z},A(\mathbf{z})) = \begin{cases} a_u \cdot E(\mathbf{z},A^u_o(\mathbf{z})) + \\ a_l \cdot [A^l_o(\mathbf{z}) - E(\mathbf{z},A^l_o(\mathbf{z}))] & \text{if} \quad g<0 \\ \\ a_u \cdot [A^u_o(\mathbf{z}) - E(\mathbf{z},A^u_o(\mathbf{z}))] + \\ a_l \cdot E(\mathbf{z},A^l_o(\mathbf{z})) & \text{if} \quad g\geq 0 \end{cases} \tag{21}$$

where E is the cosine irradiance [BRW89] as before, a_u is the amplitude of HG the upper hemisphere and a_l is its amplitude in the lower hemisphere. These amplitudes are defined as an integral over a given solid angle and are trivially precomputed, by integrating Henyey-Greenstein over the upper hemisphere for a_u and the lower hemisphere for a_l.

6. Ratio Estimators for Visibility

With our approach, using the formulation from Sect. 4 and LTSD matrices from Sect. 5, we can render biased but noise-free unshadowed single scattering. We now formulate ratio estimators to render shadows, following the work of Heitz et al. [HHM18].

Since the final radiance L is given by a sum of two terms (Eq. 1, Eq. 15), we need to formulate two ratio estimators: one for the

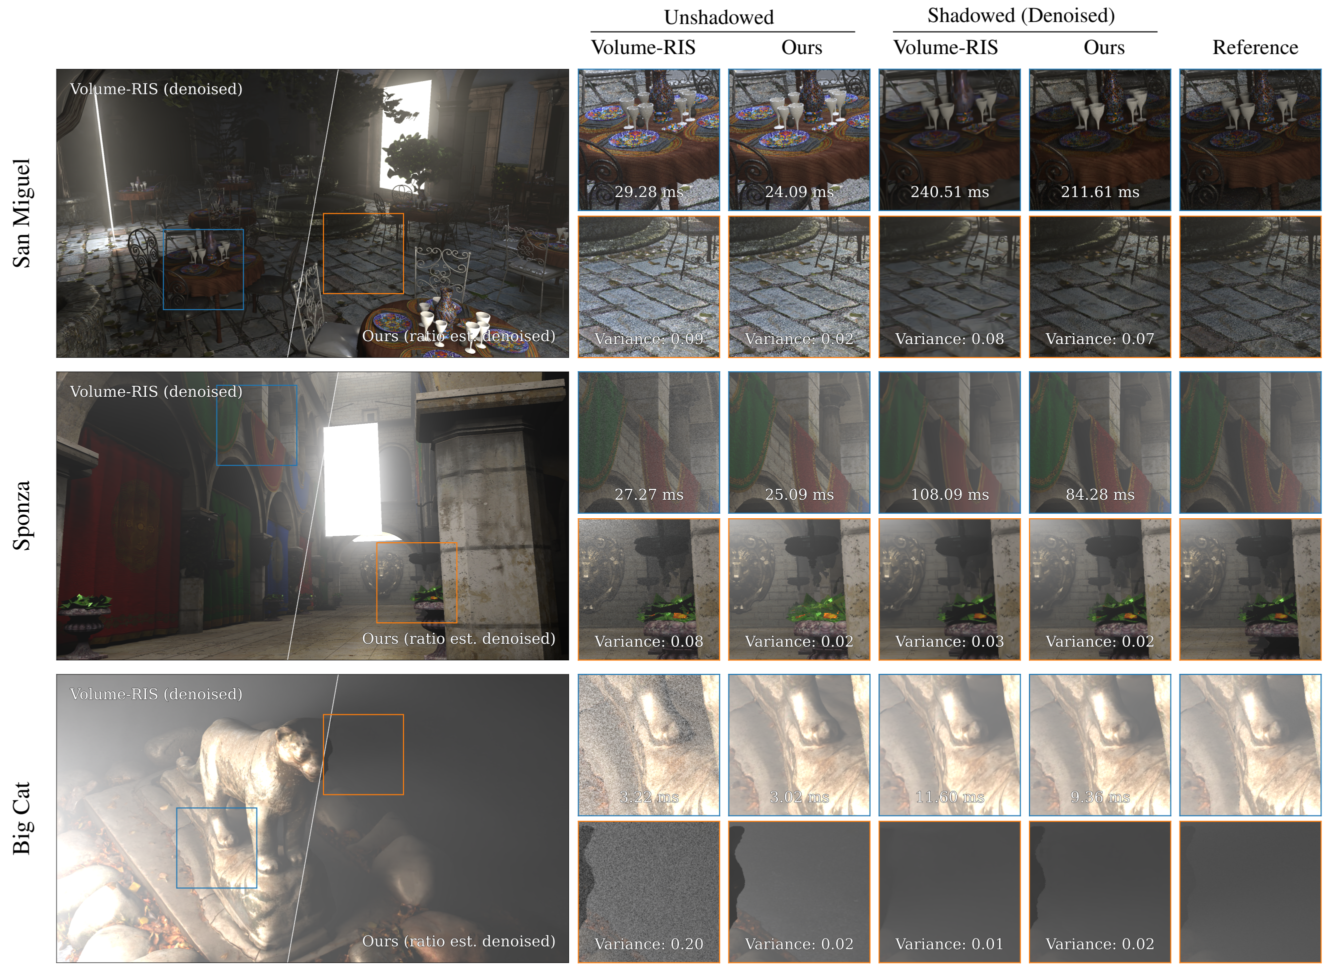

Figure 7: *We perform equal-time comparisons on three scenes of our method with Volume-RIS, for both unshadowed and shadowed configurations. We used denoised Volume-RIS and denoised ratio estimators of our method in the shadowed comparison. We also show the variance, which is error between the rendered image and it's fully converged counterpart. Our method is preserves all visual details, thanks to the semi-analytic unshadowed part while also achieving lower variance in most cases. All scenes are rendered with $N = 8$ quadrature samples.*

air-light integral and one for the surface radiance. A ratio estimator for the air-light integral can be formulated as:

$$R_{\bar{\mathcal{A}}} = \frac{\cancel{\mu_s}\int_0^x T(\mathbf{y},\mathbf{z})\left[L_i \int_{A(\mathbf{z})} \rho(\mathbf{z},\omega_o,\omega_i) T_v(\mathbf{z},\omega_i)\, d\omega_i\right] dz}{\cancel{\mu_s}\int_0^x T(\mathbf{y},\mathbf{z})\left[L_i \int_{A(\mathbf{z})} \rho(\mathbf{z},\omega_o,\omega_i) T(\mathbf{z},\omega_i)\, d\omega_i\right] dz}, \tag{22}$$

where the numerator uses T_v (transmittance with visibility) and the denominator uses T (only transmittance). We can similarly formulate a ratio estimator $R_{\bar{\mathcal{C}}}$ for the surface integral. The semi-analytic radiance $\overline{L}$ in Eq. 15 can be modified with these ratio estimators to render shadows:

$$\overline{L}(\mathbf{y},\omega_o) = \overline{\mathcal{A}} \cdot \langle R_{\bar{\mathcal{A}}} \rangle + T(\mathbf{y},\mathbf{x})\overline{L_c}(\mathbf{x},\omega_o) \cdot \langle R_{\bar{\mathcal{C}}} \rangle, \tag{23}$$

where $\langle . \rangle$ denotes a MC estimator. A benefit of this formulation is that the noise due to ratio estimators can be denoised independently. This denoising does not over-blur the final render thanks to its combination with noise-free semi-analytic radiance $\overline{L}$. More specifically, we separately denoise the numerator and denominator of $R_{\bar{\mathcal{A}}}$ and $R_{\bar{\mathcal{C}}}$ before evaluating Eq. 23

Fig. 7 shows an equal-time comparison of our method with denoised ratio estimators with denoised Volume-RIS. Denosing the output of Volume-RIS leads to blurring of details, most visible in the red curtain in Sponza and the floor in San Miguel. Our method preserves these details thanks to the semi-analytic unshadowed contribution.

7. Implementation

We implement our semi-analytic single scattering approach with CUDA and OptiX [PBD*10] for hardware accelerated ray-tracing. We refer the reader to the supplementary document for the pseudocode of our rendering algorithm. We also implement Volume-RIS, which is the method of Lin et al. [LWY21] *without* spatial and temporal reuse. This choice was made for easier evaluation, and as such

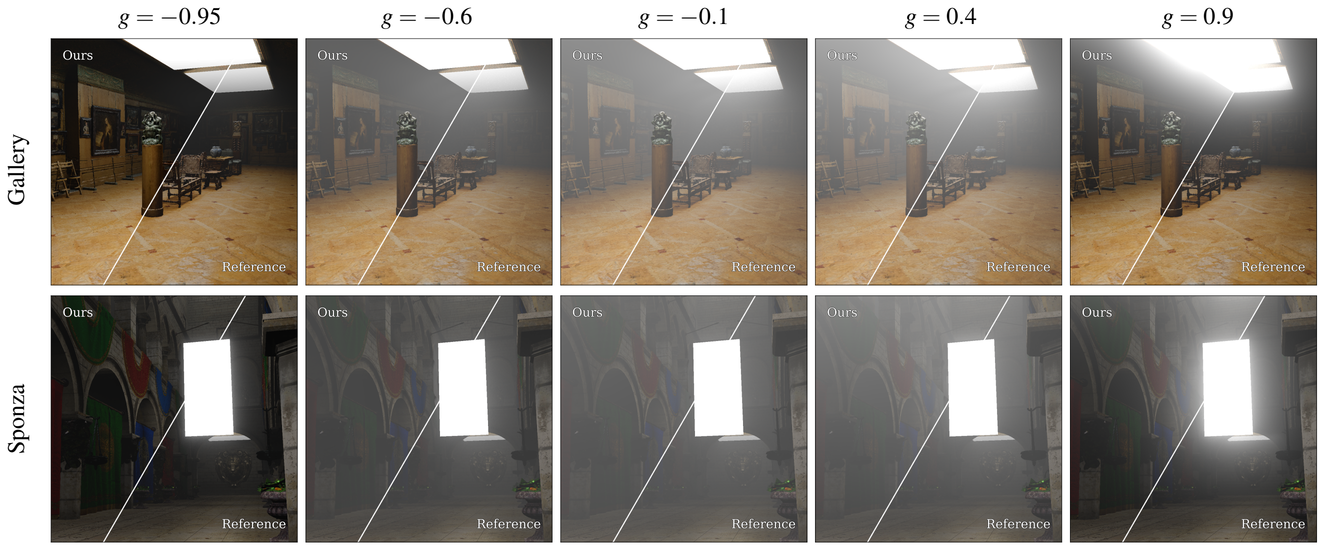

Figure 8: *Varying values of g rendered using our method and a 2K spp reference. Our method plausibly renders forward and back scattering effects.*

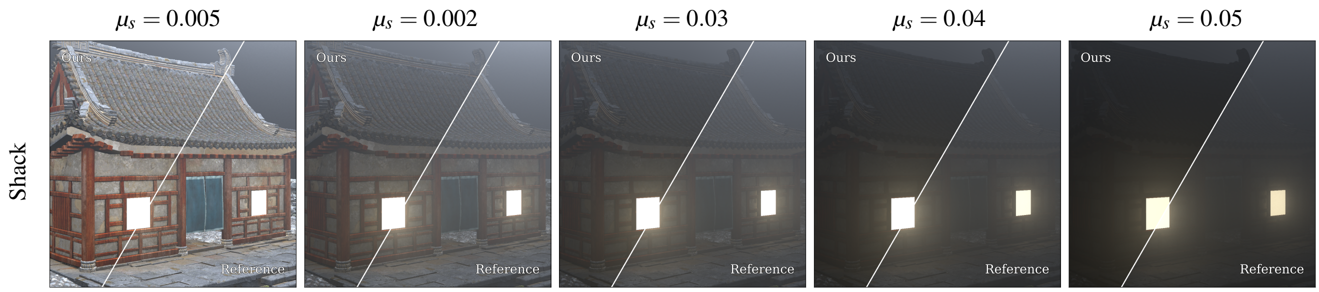

Figure 9: *Varying values of the scattering coefficient μ_s rendered using our method and a 2K spp reference. Our method plausibly renders density variations.*

any form of reuse will also benefit our method (by applying reuse to ratio-estimators), proportionally scaling the render quality and run-times.

Our implementation for fitting LTSDs to phase functions builds on the fitting approach and the code of KT et al. [KHDN22]. We modify the code to sample phase functions instead of a BRDF. We also need sampling routines for the two source distributions: clamped-cosine and inverse clamped cosine (Eq. 19). The sampling procedure for the former is well-known, and we derive a sampling procedure for the latter in the supplementary document.

Since Henyey-Greenstein is symmetric across $g = 0$, we make the following optimization to store the LTSD matrices. For a value of $g \in [-1, 0]$, we fit matrices using the clamped cosine source. This is done for 128 equally spaced values of g and stored in a table $\mathcal{T}_1$. For a value of $g \in [0, 1]$, we fit matrices using the inverse clamped cosine source, and similarly store in a 128 dimensional table $\mathcal{T}_2$. For both cases, we discard samples in the lower hemisphere of HG during fitting. The matrices M_u and M_l are fetched as follows:

$$M_u = \begin{cases} \mathcal{T}_1(g) & \text{if} \quad g < 0 \\ \mathcal{T}_2(g) & \text{if} \quad g \geq 0 \end{cases}$$
$$M_l = \begin{cases} \text{diagonal}(1,1,-1) \cdot \mathcal{T}_2(|g|) & \text{if} \quad g < 0 \\ \text{diagonal}(1,1,-1) \cdot \mathcal{T}_1(-g) & \text{if} \quad g \geq 0 \end{cases} \tag{24}$$

This avoids repeated storage, requiring only two one-dimensional tables of size 128 to be precomputed and stored. The fitting pre-computation takes about six hours on a NVIDIA RTX 3090 GPU.

8. Results & Comparisons

In this section, we validate the renderings of our method, discuss its characteristics and compare with Volume-RIS [LWY21]. Since Volume-RIS is unbiased, we use it for large number of samples per pixel (spp) reference comparisons as well.

Our rendering results are shown on various scenes with varying

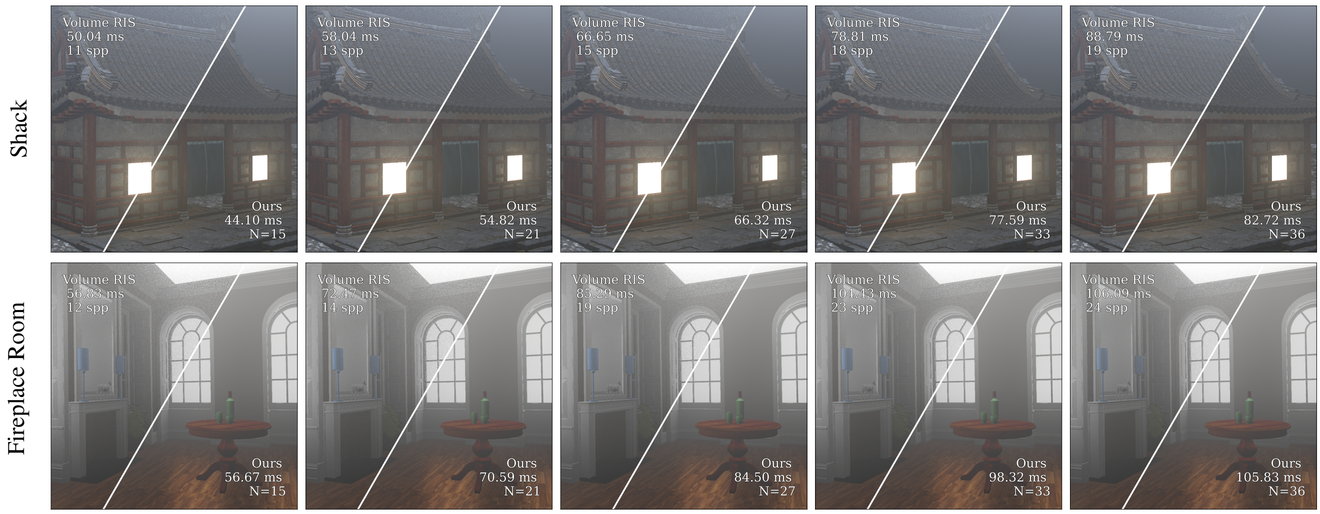

Figure 10: *Equal-time renderings of Volume-RIS compared with our renderings for various values of N (no. of quadrature samples), on two scenes. For all values of N, our method achieves a noise-free result, with bias being more for lower values. N thus serves as a quality-speed tradeoff slider of our method.*

number of area lights, geometric and texture detail. We show both unshadowed and shadowed renderings and our scenes also help evaluate volumetric shadows with light shafts (Fig. 1 right, Fig. 7 Big Cat). All images are rendered at a resolution of 1920×1080 or 1080×1080 on a workstation with NVIDIA RTX 3090 GPU, unless otherwise specified.

8.1. Validation

We begin by validating our method with ground truth ray traced 2K spp reference. We evaluate renderings at different parameter values of the medium and setting $N = 500$ (no. of quadrature samples). Fig. 8 shows renderings of our method for different values of $g = \{-0.95, -0.6, -0.1, 0.4, 0.9\}$ on two different scenes, compared with the reference. This validation serves to show that our method plausibly renders backward (negative g) and forward scattering g (positive g), thanks to our fitting approach in Sect. 5. Fig. 9 similarly evaluates renderings of our method for different values of μ_s against the reference. Note that μ_s is a spectrally varying quantity that is supported by our method as well. However, for validation purposes, it suffices to use a constant values across the spectrum. Our renderings faithfully captures the density variation in the medium.

8.2. Comparisons

In Fig. 7 we perform equal-time comparisons of our renderings with Volume-RIS on three scenes, in both unshadowed and shadowed configurations. Our unshadowed method visually performs better than Volume-RIS thanks to no noise. In the shadowed configuration, we denoise Volume-RIS and the ratio estimator in our method. Note that for Volume-RIS, we perform denoising separately on the medium and the surface for a fair comparison. The renderings of Volume-RIS exhibits blurring from the denoiser, especially of high-frequency details (floor in San Miguel, curtain in Sponza). This is not the case for our method, since we get exact details from our semi-analytic unshadowed part. The bottom row of this figure and the teaser figure show the renderings on a shadow shaft scene, demonstrating that our ratio estimators work well for volumetric shadows too.

We also show the variance for each configuration in Fig. 7. The variance of a method is computed as the error between the current render and its converged counterpart. This error estimates ignores the bias and allows us to measure variance. In most cases, we achieve lower variance than Volume-RIS.

8.3. Characteristics

Number of Quadrature samples. We evaluate equal-time renderings of our method compared to Volume-RIS for different values of N in Fig. 10. For all values of N, our method renders a noise-free result, thanks to the use of analytic evaluations of Sect. 4.2, Sect. 4.3 and the use of quadrature rule in Sect. 4.4. The bias increases for lower values of N, but the time taken to render at these values approaches real-time values. Thus, N can serve as a quality-speed tradeoff slider of our method, which is useful for gaming and lookdev applications. Equal-time renderings of Volume-RIS have significant noise at lower N which quickly startes to disappear for higher N.

Bias. Fig. 11 shows the renderings of our method for $N = 500$ and a false-color difference with ray-traced 2K spp reference. This figure shows the bias of our method, due to the approximation of LTSDs and aggregated transmittance approximation (Eq. 10). The bias of our method is acceptable for many applications and the renderings are plausible since our method is derived directly from the VRE.

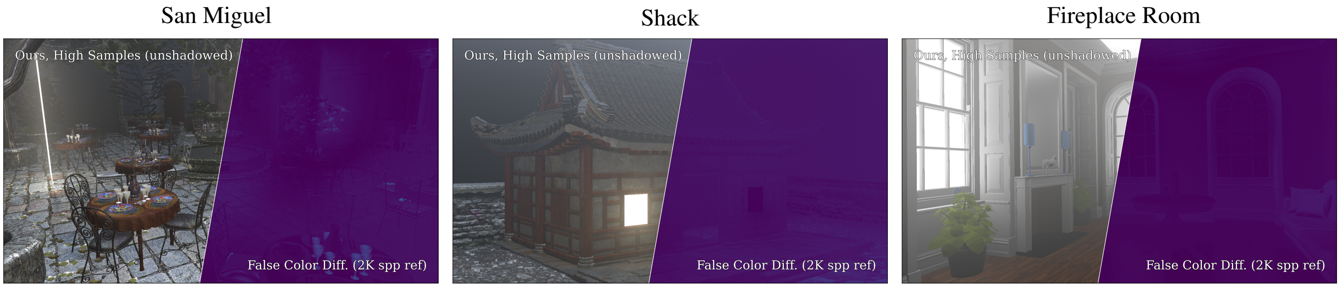

Figure 11: *We show bias of our method with false color images, computed by taking a difference with a 2K spp reference. .*

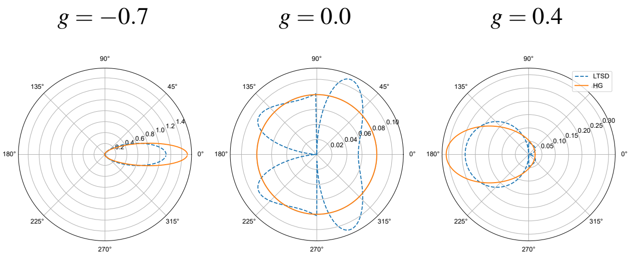

Figure 12: *Polar plots of the Henyey-Greenstein phase function (orange solid line) and our LTSD approximation from Sect. 5.2 (blue dotted line). Our approximation follows the overall shape of HG, but is discontinuous across the horizon and doesn't match HG values for individual angles (similar to previous LTC fits not exactly matching the GGX BRDF). We thus do not recommend to use it for applications such as importance sampling (a similar observation was drawn by KT et al. [KHDN22]). However, our approximation results in plausible rendering, owing to it being integrated over the solid angle of the area light.*

Temporal Consistency. Our method is temporally consistent, inspite of no temporal information sharing, as shown in our supplementary video. This is because our renderings have no noise and our approximations smoothly vary across frames. This is also true with shadows, where the ratio estimator is denoised. Since denoising is done only on the ratio part and combined with our analytic result, it hides most temporal inconsistencies [HHM18].

9. Conclusion, Limitations & Future work

We presented a method for semi-analytic interactive and unshadowed single scattering in homogeneous media with area lights. We formulated our method in a way that permitted the application of LTSDs for analytic evaluation. We showed that applying LTSDs to the in-scattered radiance is not straightforward due to the presence of a phase function that is defined on the entire sphere. This required analysis on the failure cases of LTSDs and identification of the core problem that prevents their application. We then proposed a solution that accurately fits phase functions and plausibly renders forward and back scattering effects. We also provided details to efficiently fit LTSDs to phase function using our approach, taking advantage of symmetries. Finally, we thoroughly validated our method's renderings and discussed its various characteristics. We also showed equivalent to better quality results compared to equal-time renderings of Volume-RIS.

Limitations. The biggest limitation of our method is that it is not physically correct, even though we base our formulation on the VRE. Although the resulting renderings are plausible, they may contain subtle artefacts: Fig. 8, rightmost shows an example over-brightening due to the transmittance approximation. Evaluation with quadrature of the air-light could also cause ringing artefacts, for example, if samples for two adjacent rays cross a light boundary, and should be carefully handled.

For a scene with a large number of area lights, the LTSD evaluation starts to get expensive and may not hold benefit over pure Monte Carlo evaluation, especially with a good importance sampling strategy like Volume-RIS. A minor drawback of our method is also that we now have two hyper-parameters controlling the render quality: (1) the number of quadrature samples and (2) the number of ratio estimator samples. This can be problematic for artists as these parameters, like the number of samples in traditional MC, are not intuitive.

In our current formulation, we have not considered emissive media and our method assumes homogeneous infinite media - thus it cannot handle bounded or heterogeneous media.

Finally in Fig. 12 we show polar plots of the HG phase function in orange solid lines and our fits in blue dotted lines. Our fits match the overall shape of the Henyey-Greenstein phase function, it may not be exact for most individual angles. We thus do not recommend to use our LTSD fits for sampling applications. Furthermore, our fit is not continuous across the horizon for near isotropic values of g. This is akin to previous fits shown in Heitz et al. [HDHN16] and KT et al. [KHDN22], where the LTC fits do not exactly match the GGX BRDF. We note that a similar recommendation was made by KT et al. [KHDN22].

Future Work. In the future, we would like to minimize the discontinuities and the deviations of our fits from the Henyey-Greenstein phase function. We would also like to explore the use of other phase functions like the SGGX phase function [HDCD15]. This should be possible with the fitting approach presented, but needs to be carefully explored. We would further like to derive a better approximation of the aggregated transmittance than the one in Eq. 10. It may be possible to use Stokes theorem for this. The extension to bounded and heterogeneous media is an interesting future direction, along with exploring the extension to non-exponential media.

Acknowledgements

We would like to thank Eurographics 2025 reviewers and previous SIGGRAPH Asia 2024 reviewers for their detailed comments, which led to substantial improvements in the paper. We would also like to acknowledge BlenderKit, which was a source of free 3D models for scenes used in this paper. Finally, we would like to thank Harshavardhan P. for help with result generation. Aakash KT was funded by the KCIS PhD fellowship.

References

[BRW89] BAUM D. R., RUSHMEIER H. E., WINGET J. M.: Improving radiosity solutions through the use of analytically determined form-factors. *SIGGRAPH Comput. Graph. 23*, 3 (jul 1989), 325–334. doi:10.1145/74334.74367. 3, 6

[BSA10] BILLETER M., SINTORN E., ASSARSSON U.: Real time volumetric shadows using polygonal light volumes. In *Proceedings of the Conference on High Performance Graphics* (Goslar, DEU, 2010), HPG '10, p. 39–45. 3

[CBDJ11] CHEN J., BARAN I., DURAND F., JAROSZ W.: Real-time volumetric shadows using 1d min-max mipmaps. In *Symposium on Interactive 3D Graphics and Games* (New York, NY, USA, 2011), I3D '11, p. 39–46 PAGE@7. doi:10.1145/1944745.1944752. 3

[Deb05] DEBEVEC P.: Image-based lighting. In *ACM SIGGRAPH 2005 Courses* (New York, NY, USA, 2005), SIGGRAPH '05, p. 3–es. doi:10.1145/1198555.1198709. 4

[ED10] ENGELHARDT T., DACHSBACHER C.: Epipolar sampling for shadows and crepuscular rays in participating media with single scattering. In *Proceedings of the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (New York, NY, USA, 2010), I3D '10, p. 119–125. doi:10.1145/1730804.1730823. 3

[HDCD15] HEITZ E., DUPUY J., CRASSIN C., DACHSBACHER C.: The sggx microflake distribution. *ACM Trans. Graph. 34*, 4 (jul 2015). doi:10.1145/2766988. 10

[HDHN16] HEITZ E., DUPUY J., HILL S., NEUBELT D.: Real-time polygonal-light shading with linearly transformed cosines. *ACM Trans. Graph. 35*, 4 (jul 2016). doi:10.1145/2897824.2925895. 2, 3, 10

[HG41] HENYEY L. G., GREENSTEIN J. L.: Diffuse radiation in the Galaxy. *93* (Jan. 1941), 70–83. doi:10.1086/144246. 5

[HHM18] HEITZ E., HILL S., MCGUIRE M.: Combining analytic direct illumination and stochastic shadows. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (New York, NY, USA, 2018), I3D '18. doi:10.1145/3190834.3190852. 2, 6, 10

[KF12] KULLA C., FAJARDO M.: Importance sampling techniques for path tracing in participating media. *Computer Graphics Forum 31*, 4 (2012), 1519–1528. 2

[KHDN22] KT A., HEITZ E., DUPUY J., NARAYANAN P. J.: Bringing linearly transformed cosines to anisotropic ggx. *Proc. ACM Comput. Graph. Interact. Tech. 5*, 1 (may 2022). doi:10.1145/3522612. 3, 8, 10

[KSE14] KLEHM O., SEIDEL H.-P., EISEMANN E.: Prefiltered single scattering. In *Proceedings of the 18th Meeting of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (New York, NY, USA, 2014), I3D '14, p. 71–78. doi:10.1145/2556700.2556704. 3

[LWY21] LIN D., WYMAN C., YUKSEL C.: Fast volume rendering with spatiotemporal reservoir resampling. *ACM Trans. Graph. 40*, 6 (dec 2021). doi:10.1145/3478513.3480499. 1, 2, 7, 8

[NGHJ18] NOVÁK J., GEORGIEV I., HANIKA J., JAROSZ W.: Monte Carlo methods for volumetric light transport simulation. *Computer Graphics Forum (Proceedings of Eurographics - State of the Art Reports) 37*, 2 (May 2018). doi:10/gd2jqq. 3

[PBD*10] PARKER S. G., BIGLER J., DIETRICH A., FRIEDRICH H., HOBEROCK J., LUEBKE D., MCALLISTER D., MCGUIRE M., MORLEY K., ROBISON A., ET AL.: Optix: a general purpose ray tracing engine. *Acm transactions on graphics (tog) 29*, 4 (2010), 1–13. 7

[Pet21] PETERS C.: Brdf importance sampling for polygonal lights. *ACM Trans. Graph. 40*, 4 (July 2021). doi:10.1145/3450626.3459672. 3

[PJH16] PHARR M., JAKOB W., HUMPHREYS G.: *Physically Based Rendering: From Theory to Implementation*, 3 ed. Cambridge, MA, 2016. 2

[PP09] PEGORARO V., PARKER S. G.: An analytical solution to single scattering in homogeneous participating media. *Computer Graphics Forum 28*, 2 (2009), 329–335. doi:https://doi.org/10.1111/j.1467-8659.2009.01372.x. 3

[PSP09a] PEGORARO V., SCHOTT M., PARKER S. G.: An analytical approach to single scattering for anisotropic media and light distributions. In *Proceedings of Graphics Interface 2009* (CAN, 2009), GI '09, p. 71–77. 3

[PSP09b] PEGORARO V., SCHOTT M., PARKER S. G.: Reduced dual-formulation for analytical anisotropic single scattering. In *HPG Posters* (2009). 3

[PSP10] PEGORARO V., SCHOTT M., PARKER S. G.: A closed-form solution to single scattering for general phase functions and light distributions. *Computer Graphics Forum 29*, 4 (2010), 1365–1374. 3

[PSS11] PEGORARO V., SCHOTT M., SLUSALLEK P.: A mathematical framework for efficient closed-form single scattering. In *Proceedings of Graphics Interface 2011* (Waterloo, CAN, 2011), GI '11, p. 151–158. 3

[SKN23] SHAH I. N., KT A., NARAYANAN P. J.: Combining resampled importance and projected solid angle samplings for many area light rendering. In *SIGGRAPH Asia 2023 Technical Communications* (New York, NY, USA, 2023), SA '23. doi:10.1145/3610543.3626165. 3

[SRNN05] SUN B., RAMAMOORTHI R., NARASIMHAN S. G., NAYAR S. K.: A practical analytic single scattering model for real time rendering. *ACM Trans. Graph. 24*, 3 (jul 2005), 1040–1049. doi:10.1145/1073204.1073309. 2, 3

[TCE05] TALBOT J. F., CLINE D., EGBERT P.: Importance resampling for global illumination. In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques* (Goslar, DEU, 2005), EGSR '05, p. 139–146. 2

[VGGN21] VILLENEUVE K., GRUSON A., GEORGIEV I., NOWROUZEZAHRAI D.: Practical Product Sampling for Single Scattering in Media. In *Eurographics Symposium on Rendering - DL-only Track* (2021), Bousseau A., McGuire M., (Eds.). doi:10.2312/sr.20211290. 2

[WR08] WYMAN C., RAMSEY S.: Interactive volumetric shadows in participating media with single-scattering. In *2008 IEEE Symposium on Interactive Ray Tracing* (2008), pp. 87–92. doi:10.1109/RT.2008.4634627. 3

[Wym11] WYMAN C.: Voxelized shadow volumes. In *Proceedings of the ACM SIGGRAPH Symposium on High Performance Graphics* (New York, NY, USA, 2011), HPG '11, p. 33–40. doi:10.1145/2018323.2018329. 3