

Practical product sampling for single scattering in media

Keven Villeneuve^{★,1,2} Adrien Gruson^{★,1} Iliyan Georgiev³ Derek Nowrouzezahrai¹

[★]equal contribution ¹McGill University, Canada ²SEED – Electronic Arts, Canada ³Autodesk, United Kingdom

Abstract

Efficient Monte-Carlo estimation of volumetric single scattering remains challenging due to various sources of variance, including transmittance, phase-function anisotropy, geometric cosine foreshortening, and squared-distance fall-off. We propose several complementary techniques to importance sample each of these terms and their product. First, we introduce an extension to equi-angular sampling to analytically account for the foreshortening at point-normal emitters. We then include transmittance and phase function via Taylor-series expansion and/or warp composition. Scaling to complex mesh emitters is achieved through an adaptive tree-splitting scheme. We show improved performance over state-of-the-art baselines in a diversity of scenarios.

CCS Concepts

• *Computing methodologies* → *Rendering*;

1. Introduction

Simulating single scattering in volumetric participating media is a long-standing problem in image synthesis where effective numerical solutions exist in only limited contexts. These include some recent importance sampling methods whose simplicity and efficacy have earned them notable adoption in the industry. However, the single-scattering problem involves many terms, yet existing importance sampling methods tend to focus on only a small subset of them, remaining susceptible to high estimation variance.

Computing single scattering often requires integrating the contribution of a point emitter – potentially attached to a surface – along a camera ray. Estimates of this contribution can suffer from variance due to various factors: geometric terms, transmittance along the path, phase function anisotropy. Here, techniques exist to individually importance sample some of these terms. Our goal is to approach a *full-product sampling* solution that incorporates variations due to their *composition*. Concretely, our contributions are:

- a new *point-normal* primitive to analytically sample the foreshortening and inverse-squared fall-off for an oriented emitter;
- an extension of point-normal sampling to include transmittance and/or phase function without pre-tabulation; and,
- improved performance compared to the state of the art, supporting scalable solutions for mesh lights.

Our approach is based on (non-uniform) sampling in the angular domain subtended by the point emitter and the ray [KF12]. Our derivations assume emitters contained in a homogeneous participating media; the resulting estimators remain unbiased for emitters outside the media and/or heterogeneous media, at the cost of extra variance.

2. Related work

Our work focuses on the contribution of a point emitter along a ray in scattering media. While analytical solutions exist in certain settings [PP09], we seek low-variance importance sampled Monte-Carlo estimators. Below, we summarize the most relevant prior art; we refer to Novak et al. [NGHJ18] for a thorough review.

Distance sampling. A classical approach for sampling distances in media is proportionally to transmittance [PJH16]. Equi-angular sampling instead targets the inverse squared distance fall-off of a given point source [KF12]. The two techniques can be combined via multiple importance sampling [VG95], which however corresponds to sampling from their mixture. Our work aims to develop techniques that directly treat the product of multiple contributing terms, such as the geometry factor (combining the inverse square fall-off with a cosine foreshortening at the emitter), transmittance (along with view and emission path segments), and phase function.

Product sampling. Many works address the importance sampling of a product of terms. Most recently, Hart et al. [HPM*20] composed individual sampling routines to approximate their product in a surface-illumination setting; we utilize their method for volumetric scattering. Bitterli et al. [BWP*20] improved the efficacy of importance resampling [TCE05], leveraging inter-pixel and inter-frame sample reuse in a well-founded manner. Our approach remains compatible with that reuse strategy.

Georgiev et al. [GKH*13] proposed an inversion- and tabulation-based method to importance sample volumetric paths of up to two bounces. For more bounces, path guiding is a general posterior sampling approach using cached radiance distributions in the scene. Herholz et al. [HZE*19] presented a volume path guiding method based on zero-variance random walks, whereas Deng et al. [DWWH20] adapted practical path guiding [MGN17] to media.

Emitter sampling. Analytic methods for sampling illumination from polygonal emitters exist [Arv95, UnFK13], but more complex emitter shapes require numerical solutions. Estevez and Kulla [EK18] proposed traversing a tree of light-source primitives, accounting for emitter orientation and distance through hierarchical importance sampling [MH97]. Yuksel [Yuk20] instead used a tailored importance function to avoid splitting during traversal. Vevoda et al. [VKK18] included a visibility factor that is continuously updated during emitter sampling using Bayesian statistics.

3. Angular-domain importance sampling

We start by considering single scattering in media due to a point source \mathbf{p} at all locations $\mathbf{r}_t = \mathbf{x} + t \cdot \boldsymbol{\omega}$ along an eye ray $\mathbf{r} = (\mathbf{x}, \boldsymbol{\omega})$:

$$L = \int_{t_{\min}}^{t_{\max}} \frac{L_e(\mathbf{p}, \mathbf{r}_t) \rho(\mathbf{x}, \mathbf{r}_t, \mathbf{p}) T(\mathbf{x}, \mathbf{r}_t) T(\mathbf{r}_t, \mathbf{p}) G(\mathbf{r}_t, \mathbf{p})}{f(t)} dt, \quad (1)$$

where L_e is the emitted radiance, $\rho(\mathbf{x}, \mathbf{r}_t, \mathbf{p})$ is the phase function (including the scattering coefficient $\mu_s(\mathbf{r}_t)$), $T(\mathbf{x}, \mathbf{y}) = e^{-\int_x^y \mu_t(\mathbf{z}) dz}$ is the volumetric transmittance, and μ_t is the extinction coefficient. In the geometry factor $G(\mathbf{r}_t, \mathbf{p}) = N(\mathbf{p}, \mathbf{r}_t) / \|\mathbf{p} - \mathbf{r}_t\|^2$, the term $N(\mathbf{p}, \mathbf{r}_t)$ is the light-source cosine foreshortening if the point source \mathbf{p} is on a surface and 1 otherwise. The integration bounds $\{t_{\min}, t_{\max}\}$ are determined by the ray's entry and exit points.

Solving Eq. (1) requires numerical estimation in the general case. To that end, Kulla and Fajardo [KF12] proposed to first convert the integration to an angular domain, where θ is the angle subtended by the point source \mathbf{p} along the ray, as in Fig. 1, left. The derivative of the associated transformation $t(\theta) = h \tan \theta + t_h$ is $\frac{d}{d\theta} t(\theta) = \frac{h^2 + (h \tan \theta)^2}{h} = \frac{\|\mathbf{p} - \mathbf{r}_t\|^2}{h}$, where h is the (perpendicular) distance between the point \mathbf{p} and the ray \mathbf{r} , and t_h is the offset of \mathbf{p} 's projection from the ray origin \mathbf{x} . Applying this change of variable to Eq. (1), and also assuming uniform emission profiles, yields the integral

$$L = \int_{\theta_{\min}}^{\theta_{\max}} f(t(\theta)) \frac{d}{d\theta} t(\theta) d\theta = \frac{L_e}{h} \int_{\theta_{\min}}^{\theta_{\max}} \rho(\theta) T(\theta) N(\theta) d\theta, \quad (2)$$

where the squared distance $\|\mathbf{p} - \mathbf{r}_t\|^2$ term in the transformation derivative and in the G factor cancels out [KF12]. Moving to the angular domain thus eliminates variation due to this (now canceled) term; in turn, estimators for Eq. (2) will implicitly importance sample the inverse square of that distance. The cosine foreshortening, transmittance, and phase function are now also parameterized by the angle θ ; we define them below in Eqs. (4), (8) and (9). The integration domain can be clamped based on $\{t_{\min}, t_{\max}\}$ and the directional emission profile of the light source modeled by $N(\theta)$. The resulting angular bounds $\{\theta_{\min}, \theta_{\max}\}$ exclude regions where the emitter is back-facing the ray.

The general one-sample Monte-Carlo estimator for Eq. (2) has the form

$$\langle L \rangle = \frac{L_e}{h} \cdot \frac{\rho(\theta) T(\theta) N(\theta)}{p(\theta)}, \quad (3)$$

where $\theta \in [\theta_{\min}, \theta_{\max}]$ is sampled with density $p(\theta)$. Kulla and Fajardo [KF12] used a uniform density $p(\theta) = \frac{1}{\theta_{\max} - \theta_{\min}}$, and the resulting *equi-angular* sampling technique is a straightforward and




Figure 1: Single-scattering integration from a point source \mathbf{p} along a ray $(\mathbf{x}, \boldsymbol{\omega})$, parameterized by the angle θ subtended at the source (left). When \mathbf{p} is on a surface, a foreshortening term appears that equals the dot product between the (blue) θ -direction vector and the projection \mathbf{n}^\perp of the surface normal onto the plane through \mathbf{p} and the ray (top right). The phase function ρ can also be easily parameterized by θ (bottom right).

powerful method in the toolkit of volumetric rendering. However, this uniform distribution does not importance sample any of the remaining contribution terms in Eq. (3). In the remainder of this section, we present our main contribution, which comprises a series of methods to importance sample these terms through non-uniform sampling of θ .

3.1. Analytical point-normal sampling

When the point source \mathbf{p} is on a surface, the foreshortening factor $N(\mathbf{p}, \mathbf{r}_t) = N(\theta)$ is the cosine of the angle between that surface's normal \mathbf{n} and the vector $\mathbf{r}_t - \mathbf{p}$. In the local frame $(\mathbf{p}, \boldsymbol{\omega}, \mathbf{h})$, where \mathbf{h} is the direction from \mathbf{p} toward the ray, the coordinates of that normalized vector are $(\cos \theta, \sin \theta, 0)$. Since the last coordinate is zero, the sought cosine equals the dot product between the 2D vectors $\boldsymbol{\theta} = (\cos \theta, \sin \theta)$ and the projection $\mathbf{n}^\perp = (\mathbf{n} \cdot \mathbf{h}, \mathbf{n} \cdot \boldsymbol{\omega})$ of the normal onto the plane $\boldsymbol{\omega}\mathbf{h}$ (see Fig. 1, top right):

$$N(\theta) = \boldsymbol{\theta} \cdot \mathbf{n}^\perp = (\mathbf{n} \cdot \mathbf{h}) \cos \theta + (\mathbf{n} \cdot \boldsymbol{\omega}) \sin \theta. \quad (4)$$

Our aim is to sample an angle θ proportionally to $N(\theta)$:

$$p_N(\theta) = a \cos \theta + b \sin \theta, \quad a = \frac{\mathbf{n} \cdot \mathbf{h}}{n}, \quad b = \frac{\mathbf{n} \cdot \boldsymbol{\omega}}{n}, \quad n = \int_{\theta_{\min}}^{\theta_{\max}} N(\theta) d\theta, \quad (5)$$

$$P_N(\theta) = a (\sin \theta - \sin \theta_{\min}) - b (\cos \theta - \cos \theta_{\min}), \quad (6)$$

where $p_N(\theta)$ and $P_N(\theta)$ are respectively the desired PDF and CDF. We can sample an angle θ from this density by transforming canonical uniform variates $\xi \in [0, 1)$ using the inverse CDF: $\theta = P_N^{-1}(\xi)$. Symbolic inversion of Eq. (6) yields the expression

$$\theta = \arctan \left(\frac{|a|c \pm \operatorname{sgn}(a)bd}{-bc \pm d|a|} \right), \quad (7)$$

where $c = \xi + a \sin \theta_{\min} - b \cos \theta_{\min}$ and $d = \sqrt{a^2 + b^2 - c^2}$. This expression provides two θ values, only one of which is between θ_{\min} and θ_{\max} and is thus our valid, retained sample.

3.2. Approximate product sampling

In homogeneous media we can evaluate the transmittance over the entire path length, $t + \|\mathbf{p} - \mathbf{r}_t\| = h \tan \theta + t_h + h \sec \theta$, as a function of θ (see Fig. 1, left). For the Henyey-Greenstein phase function ρ_{HG} , which is already a function of angle, the θ reparameterization is straightforward (see Fig. 1, bottom right), completing the definition of the terms in the Monte Carlo estimator (3):

$$T(\theta) = e^{-\mu_t(h(\tan \theta + \sec \theta) + t_h)} = e^{-\mu_t h(\tan \theta + \sec \theta)} e^{-\mu_t t_h}, \quad (8)$$

$$\rho(\theta) = \rho_{\text{HG}}\left(\theta + \frac{\pi}{2}\right) = \frac{1}{4\pi} \frac{1 - g^2}{(1 + g^2 + 2g \sin \theta)^{3/2}}, \quad (9)$$

where we have used $\cos(\theta + \pi/2) = -\sin \theta$ to simplify the expansion of ρ_{HG} , parameterized by the anisotropy g . Note that in $T(\theta)$ only the left exponential term is a function of θ , and it also depends on a single parameter, $\mu_t h$. The right exponential is a scaling factor that would cancel out should $T(\theta)$ be normalized to a valid PDF.

While the foreshortening $N(\theta)$ can be sampled analytically, the above two terms do not admit analytic solutions in θ and require approximation. (Note that ρ_{HG} admits analytic sampling only over the *sphere*, facilitated by the $\sin \theta$ factor in the Jacobian determinant of the change of variables from solid angle to spherical coordinates.) To that end, we derive analytically integrable polynomial approximations of the transmittance and phase function for efficient importance sampling.

Polynomial expansion. Our first approximation is a Taylor expansion around $\theta_0 = 0$; a function $f(\theta)$ is approximated by

$$\mathcal{T}_f(\theta) = \sum_{k=0}^{\text{order}} \frac{f^{(k)}(0)}{k!} \theta^k. \quad (10)$$

We consider f being either ρ or T . The expansion can be multiplied with $N(\theta)$ (4) to model the product with the emitter cosine foreshortening. The PDF and CDF of that product are respectively

$$P_{N*f}(\theta) = \frac{N(\theta) \mathcal{T}_f(\theta)}{C_{N*f}(\theta_{\min}, \theta_{\max})}, \quad P_{N*f}(\theta) = \frac{C_{N*f}(\theta_{\min}, \theta)}{C_{N*f}(\theta_{\min}, \theta_{\max})}, \quad (11)$$

where $C_{N*f}(a, b) = \int_a^b N(\theta) \mathcal{T}_f(\theta) d\theta$. Analytical expressions for C_{N*f} exist for $f \in \{\rho, T\}$, which can be obtained through integration by parts or a computer algebra system. The full derivation of those expressions can be found in the supplemental document.

Sampling still requires inverting the CDF P_{N*f} , which is generally infeasible analytically. Fortunately, the numerical Newton-Raphson inversion technique is ideally suited to the monotonic shape of CDFs. We employ an additional interval bisection to safeguard against out-of-range solution estimates due to small CDF derivatives, as described by Ureña and Georgiev [UG18].

We considered performing the Taylor expansion around $\theta_0 = (\theta_{\min} + \theta_{\max})/2$, but the improved accuracy did not offset the increased expansion complexity. The choice of expansion order also impacts the approximation quality, especially away from the expansion point: higher orders provide better approximations but at an increased computational cost. Even a costly order-14 approximation \mathcal{T}_f can significantly overshoot the ground truth and also produce negative values. For $f \in \{\rho, T\}$, we found that such gross inaccuracies occur for angles $\theta \in [\theta_{\text{clamp}}, \theta_{\text{max}}]$ where the original

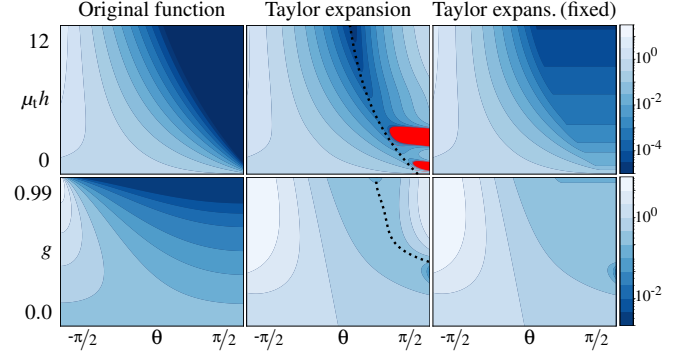


Figure 2: Plots of the transmittance (8) (top) and phase function (9) (bottom) and their Taylor expansions. The angle θ varies horizontally and the configuration parameters, $\mu_t h$ and g respectively, vary vertically. For high θ values (not too often encountered in practice), the expansions (middle) can overshoot the ground truth or have negative values (in red). To that end, for each of the two functions we find the θ_{clamp} values where these issues begin occurring, which we plot as dotted curves. For each configuration (i.e. scanline) we then replace the Taylor expansion at $\theta > \theta_{\text{clamp}}$ by a constant function: the expansion value at θ_{clamp} (right).

function g actually has a simple, near-constant shape. We thus set $\mathcal{T}_f(\theta)$ to a constant function with value $\mathcal{T}_f(\theta_{\text{clamp}})$ inside the interval $[\theta_{\text{clamp}}, \theta_{\text{max}}]$. The value of θ_{clamp} depends on the geometric configuration and the expansion order. For order-6 expansion of transmittance and phase function, we have fitted curves for θ_{clamp} as functions of the configuration parameters (see Fig. 2):

$$\theta_{\text{clamp}, T} = e^{0.210824 - 0.15974 \mu_t h}, \quad (12)$$

$$\theta_{\text{clamp}, \rho} = 18.82 - 93.9g + 184.2g^2 - 160.2g^3 + 51.77g^4. \quad (13)$$

We have obtained these curves by first manually identifying good θ_{clamp} values for many configurations (i.e. vertical coordinates in Fig. 2) and then seeking for the best-fit expression across a range of simple function parameterizations. This constant-function fix improves the expansion accuracy without increasing its order.

Polynomial interpolation. An alternative way to construct a polynomial PDF approximation for $f(\theta)$ is via interpolation. Hart et al. [HPM*20] proposed to build a quadratic Bezier PDF interpolant by evaluating f at three locations: θ_{\min} , $(\theta_{\min} + \theta_{\max})/2$, and θ_{\max} . The corresponding CDF polynomial is then cubic and analytically invertible. We use double precision for the inversion to avoid potential loss of precision in cases where the coefficient of θ^3 is small.

Warp composition. The two aforementioned approximation schemes have mild requirements on the given function f : differentiability for Taylor expansion and evaluability for interpolation. They could thus be used to importance sample not only the individual contribution terms in the radiance estimator (3) but also their full product. However, high efficiency requires a low-order approximation, which is not sufficiently expressive to provide effective importance sampling of that product. Instead, we follow Hart et al. [HPM*20] to perform approximate product sam-

Algorithm 1: Approximate full-product importance sampling for the radiance estimator (3) through the warp composition $\mathcal{B}_h(N * \mathcal{T}_f)$, where $\xi \in [0, 1]$ is a random number, and one of f and h is transmittance T and the other one is phase function ρ . In the code, P_{clamp} is the probability of sampling θ in the region $[\theta_{\min}, \theta_{\text{clamp}}]$ proportionally to the product $N * \mathcal{T}_f$. For the rest of the domain, $[\theta_{\text{clamp}}, \theta_{\max}]$, \mathcal{T}_f is forced to be constant (e.g. see Fig. 2), so we apply only our analytic point-normal sampler N .

```

1: function SAMPLEPRODUCT( $\theta_{\min}, \theta_{\max}, \theta_{\text{clamp}}, f, h, \xi$ )
2:    $P_{\text{clamp}} \leftarrow P_{N*f}(\theta_{\text{clamp}})$   $\leftarrow P_{\text{clamp}} = \int_{\theta_{\min}}^{\theta_{\text{clamp}}} p_{N*f}(\theta) d\theta$ 
3:   if  $P_{\text{clamp}} > \xi$  then  $\downarrow$  numerical inversion of  $P_{N*f}(\theta)$  (11)
4:      $\theta_0 \leftarrow \text{NewtonRaphson}(\theta_{\min}, \theta_{\text{clamp}}, \xi/P_{\text{clamp}})$ 
5:      $p_0 \leftarrow p_{N*f}(\theta) \cdot P_{\text{clamp}}$ 
6:   else
7:      $\theta_0 \leftarrow \text{SamplePointNormal}(\theta_{\text{clamp}}, \theta_{\max}, \frac{\xi - P_{\text{clamp}}}{1 - P_{\text{clamp}}}) \leftarrow \text{Eq. (7)}$ 
8:      $p_0 \leftarrow p_N(\theta) \cdot (1 - P_{\text{clamp}})$   $\leftarrow \text{Eq. (5)}$ 
9:      $(\theta_1, p_1) \leftarrow \text{WarpBezier}(h, \frac{\theta_0 - \theta_{\min}}{\theta_{\max} - \theta_{\min}})$   $\leftarrow$  rescale  $\theta_0$  to  $[0, 1]$ 
10:  return  $(\theta_1, p_0 \cdot p_1)$ 
    
```

pling by composing sampling techniques. Specifically, we apply our Taylor-expansion scheme to the product of foreshortening and transmittance/phase and warp that through a Bezier curve modeling phase/transmittance. Algorithm 1 shows pseudo-code of this approach. Our guideline is to apply the more accurate Taylor-expansion approximation to the term (transmittance or phase) with larger variation. For example, with a highly anisotropic phase function one would warp the product $N * \mathcal{T}_\rho$ through a Bezier curve \mathcal{B}_T to obtain the composition $\mathcal{B}_T(N * \mathcal{T}_\rho)$.

4. Results

We benchmark combinations of the techniques from the previous section: analytic point-normal N sampling, Taylor-expansion \mathcal{T}_f sampling of $f \in \{T, \rho\}$, sampling the product $N * \mathcal{T}_f$, Bezier-interpolation \mathcal{B}_f sampling of $f \in \{T, \rho\}$, as well as composition-based approximate product sampling through Bezier warping, e.g. the product $N * \rho$ through $\mathcal{B}_\rho(N)$ or $N * \rho * T$ through $\mathcal{B}_T(N * \mathcal{T}_\rho)$. We measure the symmetric mean absolute percentage error (SMAPE): $E = \frac{1}{P} \sum_{i=1}^P \frac{|r_i - e_i|}{|r_i| + |e_i|}$, where r_i and e_i are the reference and estimated values for the i -th pixel, respectively.

Point lights. We begin with an equal-time comparison on scenes with point and point-normal emitters and with isotropic and anisotropic ($g = 0.9$) phase functions, shown in Fig. 3. We observe that in all results our techniques produce the best results, demonstrating the benefit of product sampling. For the isotropic-phase point-normal setting we also tried the $\mathcal{B}_T(N)$ technique, but the more accurate product sampling $N * \mathcal{T}_T$ yielded consistently lower variance. Figure 4 shows plots of the variance of our techniques relative to the equal-sample variance of equi-angular sampling [KF12], for a range of ray-emitter configurations. Our techniques deliver major variance reduction in most cases.

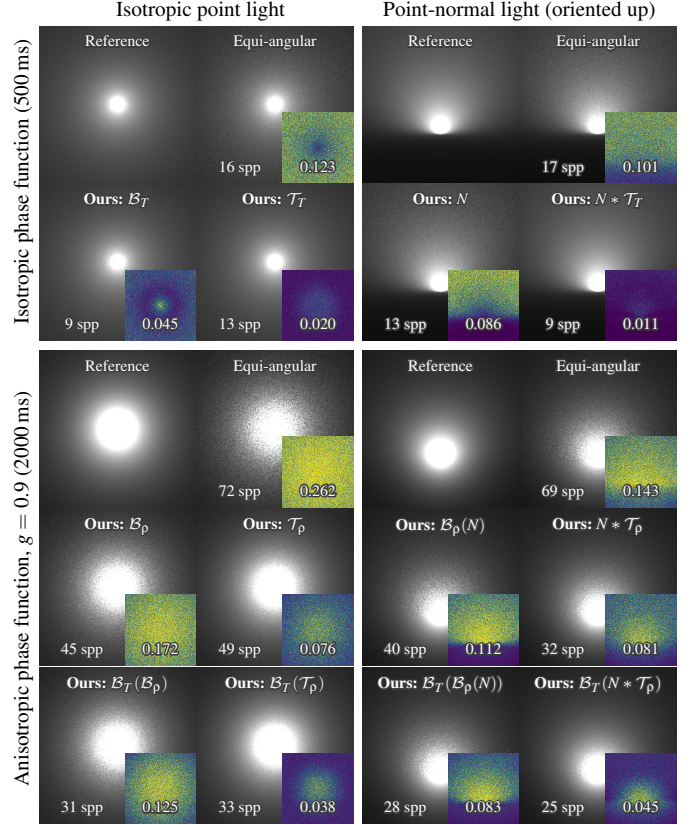


Figure 3: Equal-time comparisons on different point-emitter primitives and phase functions. In every scenario our product sampling and composition achieves the best results. See the description of the technique notation at the top of Section 4.




Figure 4: Variance plots of our techniques relative to equal-sample equi-angular sampling [KF12] under different emitter configurations (top-row illustrations), for isotropic (middle row) and anisotropic ($g = 0.9$; bottom row) phase functions. Our samplers substantially reduce variance, except when our product approximation introduces more error than a uniform (a.k.a. equi-angular) distribution. One such case is when rays span clamping angle regions, where our samplers perform similarly to equi-angular sampling (i.e. the variance ratio is around 1). Another case is when rays are very close to the emitter (bottom scanlines in the plots).




Figure 5: Equal-sample renders of a quad light scaled down by $2\times$ in each column (retaining the same total power). For a given point-normal on the quad, our $N * \mathcal{T}_T$ technique importance samples the product of all contribution terms in this isotropically scattering media. However, the point-normal is sampled uniformly, and the resulting variance can be high when the emitter is large.

Rectangular light. Our next example uses a (textured) rectangular emitter in an isotropically scattering media, shown in Fig. 5. We first sample a point-normal uniformly on the emitter surface, then employ our $N * \mathcal{T}_T$ technique. Larger area lights exhibit higher variance due to the unformed surface-area sampling; as the emitter area shrinks, we approach our technique’s ideal (point-normal) configuration.

Mesh lights. More complex (e.g. mesh) lights require a better surface sampling strategy. We apply the adaptive tree splitting of Kulla et al. [KCSG18] with their linear geometric fall-off factor and a splitting factor of 0.04. Once a triangle has been selected, we sample a point uniformly on it. Figure 6 shows an equal-sample comparison on a scene with an isotropic phase function, and Fig. 7 shows an equal-time comparison with an anisotropic phase function ($g = 0.8$). In both scenes, our techniques achieve appreciable level of variance reduction at the cost of slightly longer rendering time or smaller sample counts, respectively.

Newton solver analysis. We validate the applicability of our semi-analytic solution in cases where an iterative Newton-Raphson solver is used to invert the sampling CDF. In Fig. 8, we visualize the average number of iterations taken per pixel across four scenes. The results show that relatively few iterations are needed in practice: 3-4 steps for all scenes but the one in Fig. 7 (5.7 steps).

4.1. Discussion

We now summarize our findings with takeaways to guide the application of our various proposed solutions:

- Analytic point-normal sampling is a lightweight improvement to equi-angular sampling when dealing with area light sources.
- To sample the product of point-normal with either transmittance or phase function, the most effective option is to use our semi-analytical Taylor expansion scheme, i.e. $N * \mathcal{T}_T$ or $N * \mathcal{T}_\rho$ respectively. This scheme outperforms the Bezier-warp alternative, respectively $\mathcal{B}_T(N)$ or $\mathcal{B}_\rho(N)$.
- Sampling the full product makes most sense in optically thick and highly anisotropic media, where the preferred solution is our $\mathcal{B}_T(N * \mathcal{T}_\rho)$ sampler. When the scattering anisotropy is not

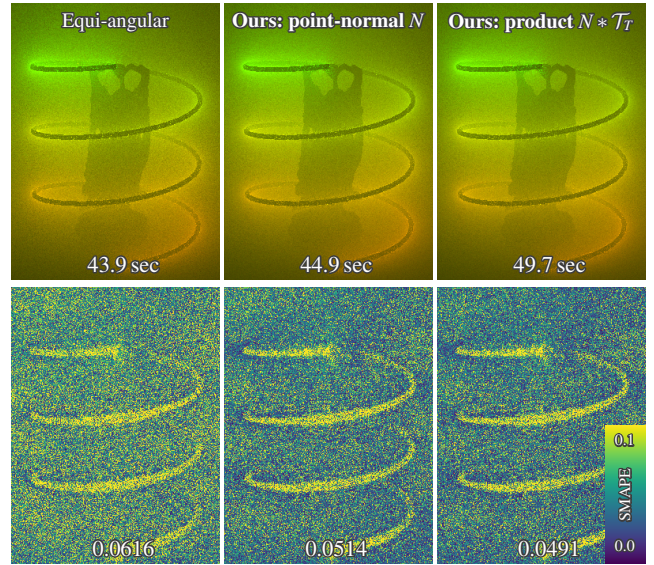


Figure 6: Equal-sample (1 spp) comparison on a scene with a mesh light and isotropic phase function. Bottom row shows false-color SMAPE images. Our point-normal and product sampling techniques achieve the lowest error.

extreme, a good alternative is $\mathcal{B}_\rho(N * \mathcal{T}_T)$ which uses the more accurate Taylor approximation for the transmittance.

- When using mesh lights, a good supplemental emitter sampling approach is essential for reducing the variance due to spatial emission variation. We employ adaptive tree splitting [KCSG18] with reasonable results.

5. Conclusion

We presented a suite of sampling techniques to reduce the variance of Monte-Carlo single-scattering estimation in media. Our point-normal primitive generalizes equi-angular sampling by additionally accounting for cosine foreshortening, while our Taylor-expansion approach addresses variations due to phase-function anisotropy or transmittance.

Our method is compatible with existing techniques, including an adaptation of warp-composition sampling [HPM*20] to the volumetric regime, to consistently outperform the state-of-the-art in diverse settings. We release our full source code to facilitate straightforward integration into physics-based renderers.

Specializing our approaches directly to planar lights is an interesting avenue for future work. So too would be an extension to heterogeneous media, perhaps using an adaptive generalization of our Taylor-based sampling approach. Finally, our method assumes direct light source visibility, i.e. without intermediate interactions with refractive medium boundaries. It would be interesting to revisit techniques that handle such interactions explicitly [WZHB, Hol15].




Figure 7: Equal-time (300 sec) comparison between our full-product sampling technique and equi-angular sampling on a scene with a complex mesh light and an anisotropic phase function ($g = 0.8$). Our technique significantly reduces the variance. The right half of each image shows false-color SMAPE.




Figure 8: False-color visualizations of the average number of Newton-Raphson steps taken per pixel by the Taylor-expansion technique used in the corresponding figure.

Acknowledgements

The scene in Fig. 7 is courtesy of matthodgesmusic. We acknowledge funding from Autodesk and the Natural Sciences and Engineering Research Council of Canada through (CRDPJ 543410-19).

References

- [Arv95] ARVO J. R.: Stratified sampling of spherical triangles. In *Annual Conference Series (Proceedings of SIGGRAPH)* (1995), ACM Press, pp. 437–438. doi:10/bdrqbf. 2
- [BWP*20] BITTERLI B., WYMAN C., PHARR M., SHIRLEY P., LEFOHN A., JAROSZ W.: Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting. doi:10/gg8xc7. 1
- [DWWH20] DENG H., WANG B., WANG R., HOLZSCHUCH N.: A practical path guiding method for participating media. *Computational Visual Media* 6 (Mar. 2020). doi:10.1007/s41095-020-0160-1. 1
- [EK18] ESTEVEZ A. C., KULLA C.: Importance sampling of many lights with adaptive tree splitting. 25:1–25:17. doi:10/ggh89v. 2
- [GKH*13] GEORGIEV I., KŘIVÁNEK J., HACHISUKA T., NOWROUZEZAHRAI D., JAROSZ W.: Joint importance sampling of low-order volumetric scattering. 1–14. doi:10/gbd5qs. 1
- [Hol15] HOLZSCHUCH N.: Accurate computation of single scattering in participating media with refractive boundaries. *Computer Graphics Forum* 34, 6 (Sept. 2015), 48–59. URL: https://hal.inria.fr/hal-01083246, doi:10.1111/cgf.12517. 5
- [HPM*20] HART D., PHARR M., MÜLLER T., LOPES W., MCGUIRE M., SHIRLEY P.: Practical product sampling by fitting and composing warps. vol. 39. doi:10.1111/cgf.14060. 1, 3, 5
- [HZE*19] HERHOLZ S., ZHAO Y., ELEK O., NOWROUZEZAHRAI D., LENSCH H. P. A., KŘIVÁNEK J.: Volume path guiding based on zero-variance random walk theory. doi:10.1145/3230635. 1

- [KCSG18] KULLA C., CONY A., STEIN C., GRITZ L.: Sony Pictures Imageworks Arnold. 29:1–29:18. doi:10/gfjkn7. 5
- [KF12] KULLA C., FAJARDO M.: Importance sampling techniques for path tracing in participating media. 1519–1528. doi:10/f35f4k. 1, 2, 4
- [MGN17] MÜLLER T., GROSS M., NOVÁK J.: Practical path guiding for efficient light-transport simulation. 91–100. doi:10/gbnvrs. 1
- [MH97] MCCOOL M. D., HARWOOD P. K.: Probability trees. In *Proceedings of the Graphics Interface 1997 Conference* (May 1997), pp. 37–46. 2
- [NGHJ18] NOVÁK J., GEORGIEV I., HANIKA J., JAROSZ W.: Monte Carlo methods for volumetric light transport simulation. 551–576. doi:10/gd2jq. 1
- [PJH16] PHARR M., JAKOB W., HUMPHREYS G.: *Physically Based Rendering: From Theory to Implementation*, 3rd ed. Morgan Kaufmann, 2016. 1
- [PP09] PEGORARO V., PARKER S. G.: An analytical solution to single scattering in homogeneous participating media. 329–335. doi:10/c9zhxn. 1
- [TCE05] TALBOT J. F., CLINE D., EGBERT P.: Importance Resampling for Global Illumination. In *Rendering Techniques (Proceedings of the Eurographics Symposium on Rendering)* (2005), Eurographics Association, pp. 139–146. doi:10/gfzsm2. 1
- [UG18] UREÑA C., GEORGIEV I.: Stratified sampling of projected spherical caps. *Computer Graphics Forum (Proceedings of EGSR)* 37, 4 (2018). 3
- [UnFK13] UREÑA C., FAJARDO M., KING A.: An area-preserving parametrization for spherical rectangles. 59–66. doi:10/gfzncr. 2
- [VG95] VEACH E., GUIBAS L. J.: Optimally combining sampling techniques for Monte Carlo rendering. In *Annual Conference Series (Proceedings of SIGGRAPH)* (1995), vol. 29, ACM Press, pp. 419–428. doi:10/d7b6n4. 1
- [VKK18] VÉVODA P., KONDAPANENI I., KŘIVÁNEK J.: Bayesian online regression for adaptive direct illumination sampling. doi:10.1145/3197517.3201340. 2
- [WZHB] WALTER B., ZHAO S., HOLZSCHUCH N., BALA K.: Single scattering in refractive media with triangle mesh boundaries. 1. doi:10/bq66rh. 5
- [Yuk20] YUKSEL C.: Stochastic lightcuts for sampling many lights. *IEEE Transactions on Visualization and Computer Graphics* (2020). doi:10.1109/TVCG.2020.3001271. 2