

Unbiased Emission and Scattering Importance Sampling For Heterogeneous Volumes

Wei-Feng Wayne Huang
Walt Disney Animation Studios

Peter Kutz*
Adobe
Walt Disney Animation Studios

Yining Karl Li
Walt Disney Animation Studios

Matt Jen-Yuan Chiang
Walt Disney Animation Studios



Figure 1: A torch embedded in thin anisotropic heterogeneous mist. Equal-time comparison of a conventional null-collision approach (left), incorporating our emission sampling strategy (middle), and additionally combining with our scattering sampling strategy via MIS (right).

Abstract

We present two new distance-sampling methods for production volume path tracing. We extend the null-collision integral formulation to efficiently gather heterogeneous volumetric emission, achieving higher-quality results. Additionally, we propose a tabulation-based approach to importance sample volumetric in-scattering through a spatial guiding data structure. Our methods improve the sampling efficiency for scenarios where low-order heterogeneous scattering dominates, which tends to cause high variance renderings with existing null-collision methods.

CCS Concepts

• Computing methodologies → Rendering; Ray tracing.

Keywords

volume rendering, path tracing

ACM Reference Format:

Wei-Feng Wayne Huang, Peter Kutz, Yining Karl Li, and Matt Jen-Yuan Chiang. 2021. Unbiased Emission and Scattering Importance Sampling For Heterogeneous Volumes. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks (SIGGRAPH '21 Talks)*, August 09-13, 2021. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3450623.3464644>

1 Introduction

Volume rendering based on the null-collision formulation of volumetric light transport has recently gained popularity in film production. Null-collision techniques are characterized by unbiased random walks that make rendering high-order scattering effects

in a heterogeneous volume, such as cloudscares, computationally feasible. Furthermore, the unbiased nature of these techniques frees users from the responsibility of adjusting ray-marching visual-quality controls such as step size and opacity threshold.

While various null-collision-based sampling strategies have been developed over the years [Novák et al. 2018], naive distance sampling with a probability density function (PDF) proportional to only transmittance makes certain artistically desirable effects extremely inefficient to render. These effects include highly emissive volumes with low extinction (e.g. fire and flames), and direct-lighting-dominated thin anisotropic media (e.g. god rays and light shafts). We propose two novel sampling strategies for improving emission sampling and scattering sampling; these strategies complement existing null-collision algorithms to make challenging production cases like the one described above practical to render, as shown in Figure 1.

2 Approach

The null-collision integral formulation of the volume rendering equation that evaluates the radiance at location \mathbf{x} in direction ω within distance d is

$$L(\mathbf{x}, \omega) = \int_0^d \bar{T}(\mathbf{x}, \mathbf{y}) (\mu_a(\mathbf{y})L_e(\mathbf{y}, \omega) + \mu_s(\mathbf{y})L_s(\mathbf{y}, \omega) + \mu_n(\mathbf{y})L(\mathbf{y}, \omega)) d\mathbf{t}, \quad (1)$$

where $\mathbf{y} = \mathbf{x} - \mathbf{t} \times \omega$. By adding null-collision particles to homogenize the originally heterogeneous volume, a distance along the ray can be analytically sampled with a PDF proportional to the combined transmittance \bar{T} formed by a constant combined extinction coefficient $\bar{\mu}$, which is the sum of volume coefficients of all three types of medium events: absorption μ_a , scattering μ_s and null-collision μ_n . These medium events can then be selectively evaluated using Monte Carlo estimators with probabilities P_\star , where subscript \star represents absorption events a , scattering events s or null-collision events n .

2.1 Emission Sampling

Conventionally, $\bar{\mu}$ is chosen as the majorant of $\mu_t = \mu_a + \mu_s$ and medium events are selected with probabilities $P_\star = \frac{\mu_\star}{\bar{\mu}}$. This poses

*The presented work was conducted while this author was employed at Walt Disney Animation Studios.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGGRAPH '21 Talks, August 09-13, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8373-8/21/08.

<https://doi.org/10.1145/3450623.3464644>




Figure 2: Free-path sample along emissive volume and PDF integral over emission energy distribution grid.

a problem for highly emissive volumes with low extinction, where the sampler is likely to step over potentially highly emissive regions as illustrated in Figure 1 (left). Kutz et al. [2017] pointed out that $\bar{\mu}$, P_a , P_s , and P_n are in fact arbitrary parameters as long as we counter-balance their contributions with sample weights. We chose $\bar{\mu}$ to be the maximum of the majorant of μ_t and the majorant of $\mu_a \times L_e$ so that we take smaller steps in highly emissive regions. To avoid excessive memory access in a single voxel, we also make sure $\bar{\mu}$ is smaller than the average voxel size. Additionally, since absorption and null-collision events don't spawn new rays, we set $P_a = P_n = 1$, such that we gather emission at every free-path sample to acquire higher quality results along a single ray (Algorithm 1). In Algorithm 2, we demonstrate how we evaluate the PDF of a next-event-estimation sample from an emissive volume by integrating probabilities along the ray through a coarser version of the volume (an emission-energy-distribution grid) [Villemin and Hery 2013] in solid angle measure [Simon et al. 2017]. Our modified emission-tracking algorithm and its PDF estimator are illustrated in Figure 2.

Algorithm 1	Algorithm 2
<pre> 1: function EVALUATEEMISSION(\mathbf{x}, ω, d) 2: $w \leftarrow 1, L_e \leftarrow 0$ 3: repeat 4: $\Delta t \leftarrow -\frac{\ln(1-\xi)}{\bar{\mu}}$ 5: $\mathbf{x} \leftarrow \mathbf{x} - \Delta t \times \omega$ 6: $L_e \leftarrow L_e + w \times \frac{\mu_a(\mathbf{x}) \times L_e(\mathbf{x})}{\bar{\mu}}$ 7: $w \leftarrow w \times \frac{\mu_n(\mathbf{x})}{\bar{\mu}}$ 8: until ($t \leftarrow t + \Delta t$) > d 9: return L_e 10: end function </pre>	<pre> 1: function PDFEMISSION(\mathbf{x}, ω) 2: $p = 0$ 3: for voxel v along ray(\mathbf{x}, ω) do 4: $[t_0, t_1] \leftarrow v$ entry/exit 5: $p \leftarrow p + \frac{(t_1^3 - t_0^3)}{3} \times \text{pdf}(v)$ 6: end for 7: return p 8: end function </pre>

2.2 Scattering Sampling

With the emission sampling introduced in the previous section, we now propose a technique to importance sample the remaining source term in Equation 1: the product of μ_s and in-scattered radiance L_s . To learn this term, we use coarsely distributed probe spheres generated as a pre-process during renderer startup. For each probe sphere, we learn guiding data by looping over light sources and summing up a *scattering sample weight* s , which approximates the integral of the product of $\mu_s(\mathbf{y})$, incoming radiance $L(\mathbf{y}, \omega')$ and phase function $\rho(\mathbf{y}, \omega, \omega')$ over solid angle, where \mathbf{y} is a randomly sampled location within the probe sphere's radius of influence and ω' and ω are directions from \mathbf{y} to a randomly sampled location on the light source and to the camera origin respectively (Figure 3a). During the rendering stage, nearby probes are queried along the

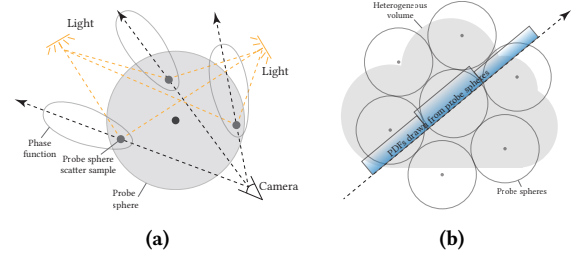


Figure 3: Scattering sample weight computation (a) and distance sampling CDF built on-the-fly from cached PDFs (b).

ray and their scattering sample weight s are used to construct a piecewise-linear CDF for distance sampling (Figure 3b).

We currently only employ this sampling strategy for camera rays and direct illumination to focus on improving visually prominent single-scattering effects while keeping overall performance overhead low. For scenarios such as high-order scattering in optically thick volumes, we combine our scattering-based distance samples with conventional transmittance-based samples using multiple importance sampling [Miller et al. 2019]. To implement the aforementioned probe spheres data structure, we use the cache points system described by Burley et al. [2018], originally designed to learn light-selection information. Compared to equi-angular sampling [Kulla and Fajardo 2012], our approach does better at handling highly anisotropic volumes since we take the phase-function term into account. Also, our approach does not need to sample a light vertex before sampling the distance because the scattering sample weight s contains global knowledge of direct illumination.

3 Conclusions

We proposed two practical sampling techniques that effectively reduce distance-sampling variance for cases that null-collision based volume rendering systems traditionally struggle with. Having a single estimator that produces robust results for various volume configurations and illumination scenarios has encouraged our artists to achieve the desired look through more straightforward in-render solutions instead of relying on approximations through compositing.

References

- Brent Burley, David Adler, Matt Jen-Yuan Chiang, Hank Driskill, Ralf Habel, Patrick Kelly, Peter Kutz, Yining Karl Li, and Dan Teece. 2018. The Design and Evolution of Disney's Hyperion Renderer. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 37, 3, Article 33 (Aug. 2018). <https://doi.org/10.1145/3072959.3073665>
- Christopher Kulla and Marcos Fajardo. 2012. Importance Sampling Techniques for Path Tracing in Participating Media. *CGF (Proc. of Eurographics Symposium on Rendering)* 31, 4 (June 2012), 1519–1528. <https://doi.org/10.1111/j.1467-8659.2012.03148.x>
- Peter Kutz, Ralf Habel, Yining Karl Li, and Jan Novák. 2017. Spectral and Decomposition Tracking for Rendering Heterogeneous Volumes. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 36, 4, Article 111 (July 2017). <https://doi.org/10.1145/3072959.3073665>
- Bailey Miller, Iliyan Georgiev, and Wojciech Jarosz. 2019. A Null-Scattering Path Integral Formulation of Light Transport. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 38, 4, Article 44 (July 2019). <https://doi.org/10.1145/3306346.3323025>
- Jan Novák, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz. 2018. Monte Carlo Methods for Volumetric Light Transport Simulation. *Computer Graphics Forum (Proc. of Eurographics)* 37, 2 (May 2018), 551–576. <https://doi.org/10.1111/cgf.13383>
- Florian Simon, Johannes Hanika, Tobias Zirr, and Carsten Dachsbacher. 2017. Line Integration for Rendering Heterogeneous Emissive Volumes. *CGF (Proc. of Eurographics Symposium on Rendering)* 36, 4 (July 2017), 101–110. <https://doi.org/10.1111/cgf.13228>
- Ryusuke Villemin and Christophe Hery. 2013. Practical Illumination from Flames. *Journal of Computer Graphics Techniques* 2, 2 (Dec. 2013), 142–155.