

Progressive null-tracking for volumetric rendering

Zackary Misso
zackary.t.misso.gr@dartmouth.edu
Dartmouth College
USA

Yining Karl Li
Brent Burley
Daniel Teece
karl.li@disneyanimation.com
brent.burley@disneyanimation.com
daniel.teece@disneyanimation.com
Walt Disney Animation Studios
USA

Wojciech Jarosz
wojciech.k.jarosz@dartmouth.edu
Dartmouth College
USA

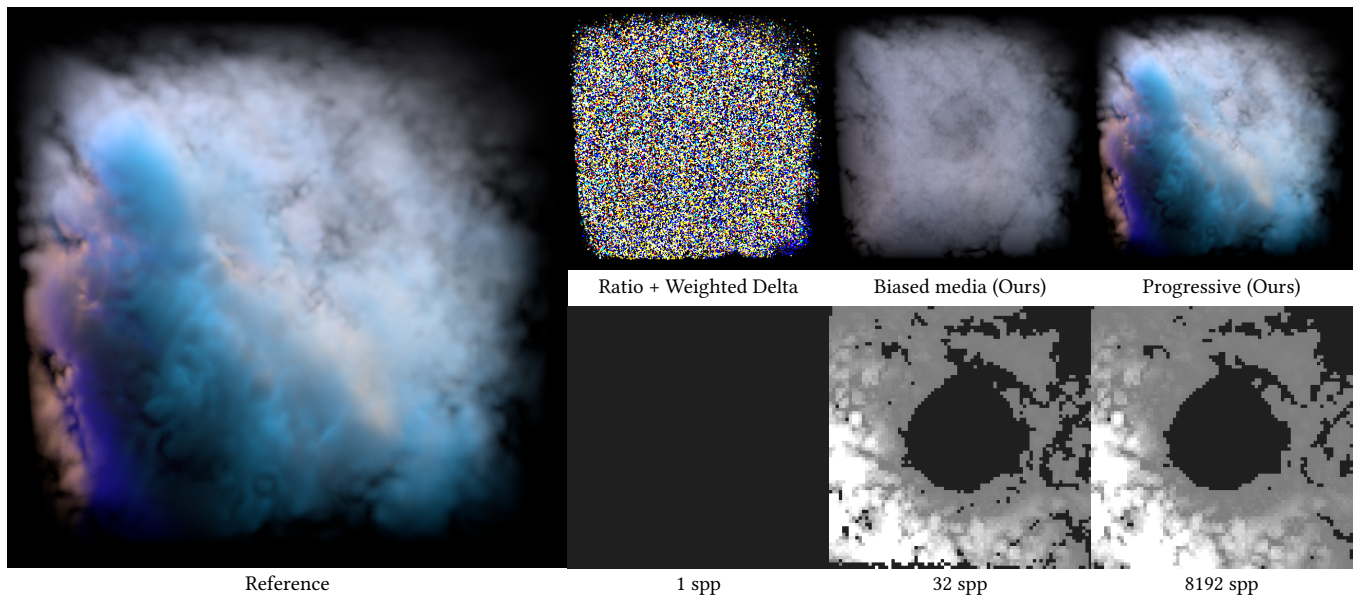


Figure 1: Most existing unbiased null-scattering methods for heterogeneous participating media require knowledge of a maximum density (majorant) to perform well. Unfortunately, bounding majorants are difficult to guarantee in production, and existing methods like ratio tracking and weighted delta tracking (top, left) suffer from extreme variance if the “majorant” ($\mu_t = 0.01$) significantly underestimates the maximum density of the medium ($\mu_t \approx 3.0$). Starting with the same poor estimate for a majorant ($\mu_t = 0.01$), we propose to instead clamp the medium density to the chosen majorant. This allows fast, low-variance rendering, but of a modified (biased) medium (top, center). We then show how to progressively update the majorant estimates (bottom row) to rapidly reduce this bias and ensure that the running average (top right) across multiple pixel samples converges to the correct result in the limit.

ABSTRACT

Null-collision approaches for estimating transmittance and sampling free-flight distances are the current state-of-the-art for unbiased rendering of general heterogeneous participating media. However, null-collision approaches have a strict requirement for

specifying a tightly bounding total extinction in order to remain both robust and performant; in practice this requirement restricts the use of null-collision techniques to only participating media where the density of the medium at every possible point in space is known a-priori. In production rendering, a common case is a medium in which density is defined by a black-box procedural function for which a bounding extinction cannot be determined beforehand. Typically in this case, a bounding extinction must be approximated by using an overly loose and therefore computationally inefficient conservative estimate. We present an analysis of how null-collision techniques degrade when a more aggressive initial guess for a bounding extinction underestimates the true maximum density and turns out to be *non-bounding*. We then build

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGGRAPH '23 Conference Proceedings, August 6–10, 2023, Los Angeles, CA, USA
© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0159-7/23/08...\$15.00
<https://doi.org/10.1145/3588432.3591557>

upon this analysis to arrive at two new techniques: first, a practical, efficient, consistent progressive algorithm that allows us to robustly adapt null-collision techniques for use with procedural media with unknown bounding extinctions, and second, a new importance sampling technique that improves ratio-tracking based on zero-variance sampling.

CCS CONCEPTS

• **Computing methodologies** → **Rendering; Ray tracing.**

KEYWORDS

progressive estimation, null-tracking, volumetric rendering

ACM Reference Format:

Zackary Misso, Yining Karl Li, Brent Burley, Daniel Teece, and Wojciech Jarosz. 2023. Progressive null-tracking for volumetric rendering. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Proceedings (SIGGRAPH '23 Conference Proceedings)*, August 6–10, 2023, Los Angeles, CA, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3588432.3591557>

1 INTRODUCTION

Null-scattering based methods for rendering participating media have become commonplace within the current generation of production renderers [Burley et al. 2018; Gamito 2018; Novák et al. 2018]. These methods take a *heterogeneous* medium—defined via its spatially varying absorption coefficient $\mu_a(x)$ and scattering coefficient $\mu_s(x)$ —and inject additional fictitious, or null, $\mu_n(x)$ density to create a (locally) *homogeneous* medium. In practice, the amount of fictitious density at any location x is specified *implicitly*,

$$\mu_t = \mu_n(x) + \underbrace{\mu_a(x) + \mu_s(x)}_{\mu_r(x)} \quad (1)$$

by taking the difference between the user-defined combined extinction¹ μ_t and the real density $\mu_r(x) = \mu_a(x) + \mu_s(x)$. By introducing a constant combined density μ_t , null-scattering allows using the same analytic sampling methods as in homogeneous media.

One of the main benefits of this approach in production—in contrast to, say, ray marching [Perlin and Hoffert 1989]—is how straightforward it is to create unbiased estimators [Georgiev et al. 2019], meaning it generates predictable results that are correct in expectation. These methods also benefit from a variety of acceleration techniques [Kutz et al. 2017; Novák et al. 2014; Szirmay-Kalos et al. 2011; Yue et al. 2010] which allow for efficiently traversing or analytically accounting for portions of a medium. More recently, the null-scattering path integral [Miller et al. 2019] has enabled combining null-scattering via multiple importance sampling (MIS) [Veach and Guibas 1995] with other techniques, such as equi-angular sampling [Kulla and Fajardo 2012], for increased robustness. This combination of factors has made null-scattering techniques well-suited for the needs of production volume rendering.

However, the need to directly specify the total density μ_t of the medium becomes a major challenge since, when μ_t is *non-bounding* (i.e. there is negative null-density $\mu_n(x) < 0$), most current methods

¹Prior work has often referred to μ_t as the *majorant* since early methods like delta tracking [Woodcock et al. 1965] required μ_t to bound $\mu_r(x)$ from above. We will instead use the terminology of Miller et al. [2019] by referring to μ_t as the *combined* or *total extinction* since it is the sum of the null, absorption, and scattering coefficients.

[Carter et al. 1972; Cramer 1978; Galtier et al. 2013; Georgiev et al. 2019; Jonsson et al. 2020; Novák et al. 2014] can perform exceptionally poorly (see Fig. 1). This can become prohibitive in production since volumetric assets are often the result of a series of physical simulations, procedural manipulations, or other artistic workflows where the final density becomes a black box that can only be point-evaluated. Any choice of μ_t based on point evaluations becomes a guess for a truly bounding combined extinction μ_t .

In production rendering, we have effectively two options to remain robust and unbiased when using null-scattering: We can specify a conservative value for μ_t , but this may result in prohibitively expensive renders if the value is excessively loose for some or all of the volume. Alternatively, we can bake all volumes into voxel density grids as a pre-process. This is what we do in our current system, since it allows computing tight values for μ_t , but this has a startup cost, is memory intensive, and reduces the visual fidelity of all media. Clearly, neither option is ideal.

With an eye towards our next-generation production volume renderer, we aim to develop a technique resilient to non-bounding, otherwise “incorrect”, guesses for μ_t . While some specialized techniques do exist for side-stepping negative null coefficients [Carter et al. 1972; Cramer 1978; Galtier et al. 2016; Szirmay-Kalos et al. 2017], we aim to create a generalized framework with which we can easily adapt all existing null-scattering techniques that do not currently handle negative null coefficients well. To accomplish this, we first analyze why existing null-scattering methods exhibit such poor variance behavior in the presence of negative null densities (Sec. 3). With these insights, we propose (Sec. 4) a simple way to avoid this increased variance by clamping the real medium density μ_r to never exceed the specified total extinction μ_t . This approach in essence trades variance for bias, since we reduce variance but render a modified (“biased”) medium. We then introduce a consistent, progressive formulation (akin to progressive photon mapping [Hachisuka et al. 2008]) to eliminate this bias across multiple passes (Sec. 4.1) and a practical, scene-adaptive approach for updating the clamping parameter to keep render times efficient (Sec. 4.2). Inspired by our analysis, we additionally introduce (Sec. 5) an approximate zero-variance sampling scheme to specifically improve ratio-tracking [Novák et al. 2014]. We provide a longer review of how our contributions relate to prior work in Sec. 6, and include comparisons to existing approaches in Sec. 7. We provide our full implementation [online](#) [Misso et al. 2023].

2 BACKGROUND

To render heterogeneous participating media in the null-scattering framework [Kutz et al. 2017; Miller et al. 2019], the radiance arriving at a point x from direction ω can be written as

$$L(x, \omega) = \overline{\text{Tr}}(x, y) L^s(y, \omega) + \int_x^y p(t) L^m(t, \omega) dt. \quad (2)$$

This involves the radiance, $L^s(y, \omega)$, leaving the nearest surface intersection y along the ray [Immel et al. 1986; Kajiyama 1986], weighted by the combined transmittance $\overline{\text{Tr}}(x, y) = e^{-\|x-y\|\mu_t}$ which expresses the *probability* of encountering no real or null collisions when traveling from x to y . We additionally need to integrate the medium radiance L^m from all points between x and y along the ray, weighted by the combined free-flight distribution $p(t) = \mu_t e^{-t\mu_t}$,

which expresses the probability *density* of a collision occurring with a real or null particle at distance t . In a slight abuse of notation, we allow t to refer to both a distance, and a position at that distance when it is clear from context.

The medium radiance,

$$L^m(x, \omega) = \frac{\mu_a(x)}{\mu_t} L_e^m(x, \omega) + \frac{\mu_n(x)}{\mu_t} L(x, \omega) + \frac{\mu_s(x)}{\mu_t} \int_{S^2} \rho_m(\omega, x, -\omega') L(x, \omega') d\omega', \quad (3)$$

sums the contributions from emission $L_e^m(x, \omega)$, null-scattering in the forward direction, and integrated in-scattered radiance. We refer to the quantities $\alpha_*(x) = \mu_*(x)/\mu_t$ as the absorption, null, and scattering albedos for $\mu_* = \mu_a$, $\mu_* = \mu_n$, and $\mu_* = \mu_s$, respectively. The phase function ρ_m is the medium's directional distribution dictating how rays can scatter given a location and incoming direction.

Monte Carlo rendering estimates the above integral equations using Monte Carlo sampling. In practice, we often have a choice of different sampling PDFs that excel at sampling certain types of light paths or at sampling certain parts of the integrand. For instance, we can dramatically reduce variance compared to naive unidirectional path tracing by tracing shadow rays towards light sources [Shirley et al. 1996], and combine different strategies using MIS [Veach and Guibas 1995]. The formulation above readily supports directional MIS [Kutz et al. 2017] by considering different directional PDFs for estimating the in-scattered radiance integral in Eq. (3). Also, by inserting Eq. (3) into Eq. (2) and recursively considering all forward null events, we can obtain the Volterra integral formulation for the *real* transmittance [Georgiev et al. 2019],

$$\text{Tr}(x, y) = \overline{\text{Tr}}(x, y) + \int_x^y \alpha_n(t) p(t) \text{Tr}(t, y) dt, \quad (4)$$

which is necessary to compute when tracing shadow rays. The null-scattering path integral [Miller et al. 2019] then allows us to perform MIS between different strategies for evaluating transmittance.

For a more thorough overview of null scattering theory for volumetric rendering, we refer the reader to Novák et al. [2018]'s excellent state-of-the-art report.

3 ROOT CAUSE OF DEGENERATE BEHAVIOUR

To improve the robustness of null-scattering techniques, we must first analyze how existing techniques degrade when used with non-bounding estimates of μ_t . We analyze the root cause of increased variance when estimating transmittance using Eq. (4), and then similarly discuss the free-flight sampling case.

3.1 Transmittance estimation

Most tracking-based transmittance estimators [Georgiev et al. 2019; Novák et al. 2014; Woodcock et al. 1965] can be derived by applying Monte Carlo estimation directly to Eq. (4), giving

$$\langle \text{Tr}(x, y) \rangle = \frac{\overline{\text{Tr}}(x, y)}{P_{\text{Tr}}} + \frac{\alpha_n(t) p(t) \langle \text{Tr}(t, y) \rangle}{p_{\text{MC}}(t) P_{\text{rec}}}, \quad (5)$$

where P_{Tr} is the probability that the first term is evaluated, P_{rec} is the probability that the recursive second term is evaluated, $\langle \text{Tr}(x, y) \rangle$ denotes an estimate of the quantity $\text{Tr}(x, y)$, and we use $p_{\text{MC}}(t)$ to




Figure 2: A comparison of the behaviour of ratio tracking with bounding (top) and non-bounding (bottom) total extinctions. On the right we plot the ground truth transmittance (red), the average of 100 invocations of ratio tracking (black), and also plot the individual instances (thin blue).

denote the conditional probability density used to sample a distance t for the recursive integral. We include the subscript to be explicit that this is a distribution used for *Monte Carlo* and is generally distinct from the combined free-flight distribution $p(t)$ in the integral equation itself.

3.1.1 Ratio tracking. As a specific example for how negative null-density can impact estimators for Eq. (5), let us consider the ratio-tracking family of techniques [Novák et al. 2014]. In its most basic form, ratio tracking works by sampling a tentative distance proportional to the medium's combined free-flight distribution $p(t)$. It chooses to only evaluate the first term of Eq. (5) if a distance $t > y$ is sampled, and only evaluates the second term when $t \leq y$ is sampled. This is equivalent to setting $P_{\text{Tr}} = \overline{\text{Tr}}(x, y)$ and $P_{\text{rec}} = 1 - \overline{\text{Tr}}(x, y)$. The conditional probability density of sampling a distance is $p_{\text{MC}}(t) = p(t)/P_{\text{rec}}$, where P_{rec} ensures that the PDF is properly normalized within the integration interval from x to y .

Inserting these values into Eq. (5), ratio tracking can thus be written succinctly as,

$$\langle \text{Tr}(x, y) \rangle = \prod_{j=1}^{m-1} \alpha_n(x_j) = \prod_{j=1}^{m-1} \frac{\mu_n(x_j)}{\mu_t} = \prod_{j=1}^{m-1} \frac{\mu_t - \mu_r(x_j)}{\mu_t}, \quad (6)$$

where m is the total number of recursions in Eq. (5) or steps taken, and x_j is the j th ordered location along the ray [Novák et al. 2014]. The resulting running product is a nice monotonically decreasing, positive function when μ_t is bounding, i.e. $\mu_r(x) < \mu_t$, however when μ_t becomes non-bounding this behaviour deteriorates (Fig. 2).

When $\mu_t < \mu_r(x) < 2\mu_t$, the sign of the transmittance estimates (6) will begin to flip, but their absolute values will remain a monotonically decreasing function bounded within $[-1, 1]$ since each term in the product is between $[-1, 1]$. We refer to this case as slightly non-bounding since variance will increase, but may still be manageable. On the other hand, once μ_t is significantly non-bounding, $\mu_r(x) > 2\mu_t$, the absolute magnitude of each term in Eq. (6) will become greater than one, causing the running product, and hence the variance, to become unbounded as m increases (Fig. 3).

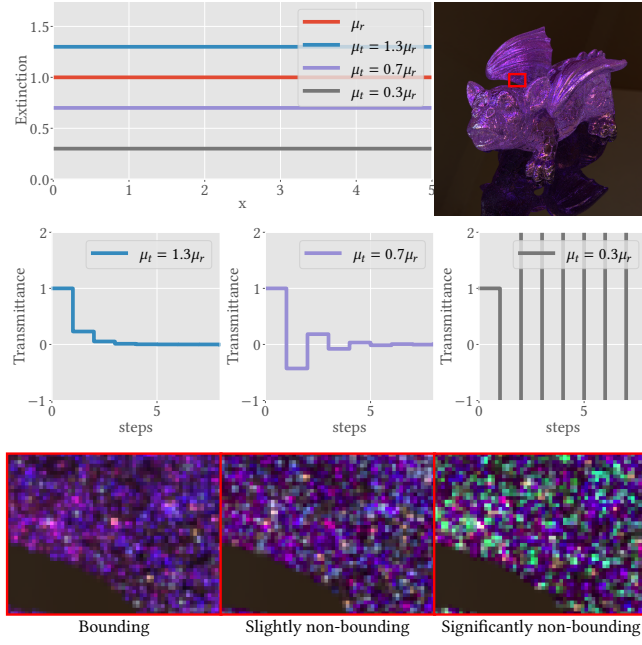


Figure 3: A comparison of the different behaviours exhibited by ratio tracking when μ_t is bounding, slightly non-bounding, and significantly non-bounding for a constant density absorbing medium (top row). We visualize the returned transmittance as a function of the number of exponential distances sampled (middle row) and visualize the difference in render quality for approximately equal spp renders (bottom row).

Variance analysis. Georgiev et al. [2019] previously derived an analytic expression for the variance of ratio tracking in the special case of a homogeneous medium where $\mu_r(t) = \mu_r$:

$$V[\langle \text{Tr}(x, y) \rangle] = e^{-2\mu_r d} \left(e^{\left(\frac{\mu_r d}{\mu_t}\right)^2} - 1 \right), \quad (7)$$

where $d = y - x$ is the distance between x and y . Georgiev et al. focused specifically on the behaviour of ratio tracking when the combined extinction μ_t is tightly versus loosely bounding, however, they did not analyze the non-bounding case. By rewriting Eq. (7) as

$$V[\langle \text{Tr}(x, y) \rangle] = e^{\left(\frac{\mu_r - 2\mu_t}{\mu_t}\right)\mu_r d} - e^{-2\mu_r d}, \quad (8)$$

we can analytically explain the different behaviours of ratio tracking under varying degrees of bounding extinctions (see Fig. 3).

The crucial part of Eq. (8) is the first exponential, since the choice of μ_t can influence the sign of the exponent and whether variance remains bounded. When the extinction is slightly non-bounding ($\mu_t < \mu_r < 2\mu_t$) the exponent remains negative, so variance is bounded and decreases as a function of distance. Once μ_t is significantly non-bounding ($\mu_r > 2\mu_t$) the exponent becomes positive and the variance unbounded. Even worse, the variance increases *exponentially* as a function of distance.




Figure 4: We visualize the degradation of weighted delta tracking [Cramer 1978; Galtier et al. 2013] for free-flight distance sampling in the presence of non-bounding extinctions for a scene containing a homogeneous medium while only rendering direct illumination. All renders are roughly equal extinction calls, and transmittance is evaluated analytically.

3.2 Free-flight sampling

Evaluating transmittance to the next surface scattering event only accounts for the first term in Eq. (2). Estimating the remaining integral typically requires free-flight distance sampling routines, such as delta tracking [Woodcock et al. 1965], weighted delta-tracking [Cramer 1978; Galtier et al. 2013], or decomposition tracking [Kutz et al. 2017]. These techniques estimate Eq. (3) by assigning probabilities P_a, P_n, P_s to evaluate emission, null-scattering, and in-scattered radiance, respectively:

$$\langle L^m(x, \omega) \rangle = \frac{\mu_a(x)}{\mu_t} \frac{L_e^m(x, \omega)}{P_a} + \frac{\mu_n(x)}{\mu_t} \frac{L(x, \omega)}{P_n} + \frac{\mu_s(x)}{\mu_t} \frac{\rho_m(\omega, x, -\omega') L(x, \omega')}{p_{MC}(\omega') P_s}. \quad (9)$$

Unfortunately, when negative null-densities occur, free-flight sampling of this form suffers from a similar explosion in variance as we previously saw for transmittance estimation.

For instance, weighted delta tracking [Cramer 1978; Galtier et al. 2013] sets P_a, P_n , and P_s to

$$P_*(x) = \frac{\mu_*(x)}{\mu_r(x) + |\mu_n(x)|}. \quad (10)$$

Inserting these probabilities leads to a shared factor of $\frac{\mu_r(x) + |\mu_n(x)|}{\mu_t}$ multiplying each of the three branches in Eq. (9). When μ_t is non-bounding, this factor is greater than 1 and the weight (and variance) of the estimator grows exponentially with distance (see Fig. 4) in a similar manner to ratio tracking.

4 OUR METHOD

While slightly non-bounding extinctions might be usable in a production setting they result in increased variance, and the potential for unbounded variance due to significantly non-bounding extinctions is prohibitive. Instead of modifying individual estimators to be more resilient to non-bounding extinctions, the core idea behind our method is to first bias the medium itself to avoid non-bounding extinctions and then use a progressive technique to remove the bias over time.

For a chosen extinction μ_t , we simply do not allow negative null densities. Given the relationship in Eq. (1), we accomplish this by defining our modified medium with a clamped real extinction

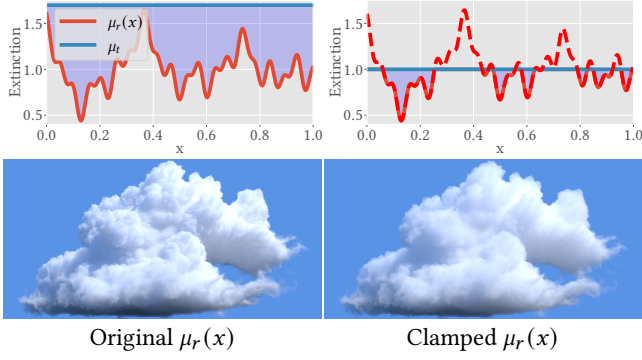


Figure 5: We visualize the idea of biasing a medium through clamping its real density to avoid $\mu_n < 0$. Note that clamping $\mu_r(x)$ to μ_n allows for the use of null-tracking but results in lost detail.




Figure 6: We visually illustrate an example of how taking the average of n renders while linearly increasing the total extinction can result in converging images. For visualization purposes, each iteration is rendered to near convergence instead of just one pixel sample.

coefficient $\mu_r^+(x) = \min(\mu_r(x), \mu_t)$. The corresponding clamped null density $\mu_n^+(x) = \mu_t - \mu_r^+(x)$, is then always non-negative.

We can now apply any unbiased tracking-based estimator on this modified medium while avoiding the degenerate behavior we discussed in Sec. 3. However, since we’ve modified the medium, we will get an unbiased estimate, but of the wrong answer (see Fig. 5). In essence, this trades variance for bias.

4.1 Progressive Extinctions

We then formulate a progressive and consistent estimator (akin to progressive photon mapping [Hachisuka and Jensen 2009; Hachisuka et al. 2008; Knaus and Zwicker 2011]), that ensures the bias will go to zero. Let us denote with $I(\mu_t)$ the expected value of rendering a medium clamped to a particular choice of μ_t . We form our progressive estimator by averaging n images, where the images are rendered using a sequence $\{\mu_t^{(k)}\} := \{\mu_t^{(1)}, \mu_t^{(2)}, \dots, \mu_t^{(n)}\}$ of non-decreasing total extinctions:

$$\langle I \rangle = \frac{1}{n} \sum_{k=1}^n \langle I(\mu_t^{(k)}) \rangle, \quad (11)$$

where $\mu_t^{(k)}$ is the total extinction used to render image k .

As long as $\mu_t^{(k)}$ bounds $\mu_r(x)$ after some finite iteration j , the bias will disappear in the limit of n since an infinite number of unbiased iterations will overpower any bias accumulated from the first j iterations. While we provide a more rigorous proof in the supplemental, this means Eq. (11) is consistent, eliminating bias in the limit,

$$I = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \langle I(\mu_t^{(k)}) \rangle. \quad (12)$$

A naive choice that satisfies these conditions could be to linearly increase the total extinction in each iteration:

$$\mu_t^{(k+1)} = c\mu_t^{(k)}, \quad (13)$$

for any constant $c > 1$ which we visualize in Fig. 6. Unfortunately, this is impractical since the cost of rendering the image increases as $\mu_t^{(k)}$ increases.

4.2 Adaptive Approach

For practical applications we would like a progressive sequence $\{\mu_t^{(k)}\}$ which reaches a bounding μ_t and then stops increasing μ_t to maintain low render times. Luckily, the individual images we average in our progressive formulation (11) become unbiased once we reach $\mu_t^{(k)} \geq \mu_r$, which affords us considerable flexibility to adaptively update $\mu_t^{(k)}$ from one iteration to the next.

We choose to update μ_t only when encountering non-bounded real densities in the previous rendering pass:

$$\mu_t^{(k+1)} = \max_{\forall x^{(k)}} (\mu_t^{(k)}, \mu_r(x^{(k)}) + \epsilon), \quad (14)$$

where $x^{(k)}$ is any medium lookup performed in pass k , and ϵ is a small user-defined constant. As long as there is a non-zero measure of points within an ϵ of the function’s maximum, Eq. (14) guarantees that a bounding extinction is expected to be found in finite time. In other words, all density functions where we have a non-zero chance of randomly sampling a position whose density is close to (but not necessarily equal to) the peak density are supported. This accounts for all densities typically used in production, like tri-linearly or tri-cubically interpolated and layered grids, and most procedural manipulations.

Our entire clamping and progressive approach can be practically described as replacing all density evaluations in a renderer with Algorithm 1 which returns a potentially clamped density, and an updated total extinction for use in the next render pass.

Algorithm 1: EvaluateDensity($\mu_t^{(k)}, o, \omega, t$)

Input : $\mu_t^{(k)}$, ray origin o , ray direction ω , distance t

Output: $\mu_r^+, \mu_t^{(k+1)}$

- 1 $\mu_r^+ \leftarrow \min(\mu_r(o + t\omega), \mu_t^{(k)})$
 - 2 $\mu_t^{(k+1)} \leftarrow \max(\mu_t^{(k)}, \mu_r(o + t\omega) + \epsilon)$
 - 3 **return** $\mu_r^+, \mu_t^{(k+1)}$
-




Figure 7: An illustration of the difference between our adaptive importance sampling and ratio-tracking’s traditional importance sampling. The black vertical line represents the location of the current tracking in the medium (left graph). The right graph shows that our technique will sample the next exponential distance proportional to $\mu_n(t)$ (blue) instead of $\mu_t(t)$ (red).

5 IMPORTANCE SAMPLING EXTINCTIONS

Most notable null-tracking techniques sample distances from a distribution $p_{MC}(t)$ which is directly proportional to the combined free-flight distribution of the underlying medium. However, this is not the most optimal choice. In a similar vein to d’Eon and Novák [2021]; Křivánek and d’Eon [2014], we formulate a theoretically zero-variance estimator for evaluating transmittance with ratio tracking. While the zero-variance estimator is not achievable, we use it to derive an approximate but practical importance sampling strategy to directly improve ratio tracking.

Zero variance analysis. Assuming we have a bounding extinction $\mu_t \geq \mu_r$, and the real density is homogeneous, we can formulate a zero variance estimator for ratio tracking by sampling distances proportional to,

$$p_{MC}(t) = \mu_n \hat{\text{Tr}}(x, t), \text{ where } \hat{\text{Tr}}(x, t) = e^{-\mu_n t} = \frac{\overline{\text{Tr}}(x, t)}{\text{Tr}(x, t)}. \quad (15)$$

To confirm this, we can plug Eq. (15) into Eq. (4),

$$\langle \text{Tr}(x, y) \rangle = \frac{\overline{\text{Tr}}(x, y)}{P_{\text{Tr}}} + \frac{\alpha_n(t) p(t) \langle \text{Tr}(t, y) \rangle}{\mu_n \hat{\text{Tr}}(x, t) P_{\text{rec}}}. \quad (16)$$

By noting that ratio tracking obtains P_{rec} and P_{Tr} by integrating and renormalizing $p_{MC}(t)$, this equation simplifies to

$$\langle \text{Tr}(x, y) \rangle = \text{Tr}(x, y) + \text{Tr}(x, t) \langle \text{Tr}(t, y) \rangle, \quad (17)$$

which always returns the exact transmittance (note that ratio tracking evaluates only the first or second term on the right-hand side).

This result says that instead of sampling distances t proportional to the total free-flight distribution (as ratio tracking normally does), we would get zero variance if we sampled according to the “null-free-flight distribution” in Eq. (15). Unfortunately, this is not possible in practice since $\mu_n = \mu_n(x)$ is an unknown, spatially varying function.

However, this exercise reveals an important insight. Ratio tracking can be improved by *decoupling* the medium’s combined free-flight distribution from the distribution used for sampling. While this may seem obvious, to our knowledge, this is the first time this idea has been proposed for any tracking based technique.

An importance sampled approach. Motivated by our zero variance analysis of ratio tracking and progressive technique for adapting μ_t during renders (Sec. 4.2), we propose an importance-sampled form of ratio tracking Algorithm 2 which chooses $p_{MC}(t)$ for the next

Algorithm 2: AdaptiveRatioTracking($\mu_t^{(k)}, o, \omega, d$)

Input : $\mu_t^{(k)}, o, \omega$, medium length d
Output : T : transmittance, $\mu_t^{(k+1)}$: next iteration majorant

- 1 $T \leftarrow 1$
- 2 $t \leftarrow 0$
- 3 $\mu_{MC} \leftarrow \mu_t^{(k)}$
- 4 **while** $t < d$ and $\mu_{MC} > 0$ **do**
- 5 $x \leftarrow \min(-\ln(1-u)/\mu_{MC}, d-t)$
- 6 $T \leftarrow T * \exp((\mu_{MC} - \mu_t^{(k)}) * x)$
- 7 **if** $x \neq d - t$ **then**
- 8 $t \leftarrow t + x$
- 9 $\mu_r^+, \mu_t^{(k+1)} = \text{EvaluateDensity}(\mu_t^{(k)}, o, \omega, t)$
- 10 $T \leftarrow T * (\mu_t^{(k)} - \mu_r^+) / \mu_{MC}$
- 11 $\mu_{MC} \leftarrow \mu_t^{(k)} - \mu_r^+$
- 12 **return** $T, \mu_t^{(k+1)}$

recursion of Eq. (4) by setting it proportional to Eq. (15) where we use $\mu_{MC} = \mu_n^+(x)$ from the current recursion as a proxy for the true μ_n . Compared to traditional ratio tracking, our approach favors taking larger free-flight distances which we illustrate in Fig. 7.

6 RELATION TO PRIOR WORK

In this section, we review how prior work has gone about tackling the problem of unknown total extinctions and progressive estimation as well as how they differ from the contributions we have introduced in this paper.

Handling unknown total extinctions. Previous efforts at handling heterogeneous media with unknown total extinction values focused on side-stepping the problem of potential negative null coefficients entirely. Szirmay-Kalos et al. [2017] and Szirmay-Kalos et al. [2018] demonstrate methods to bypass negative null coefficients by manipulating the medium itself to be analytically sampleable while maintaining the same expected radiance as the original medium. Weighted tracking [Carter et al. 1972; Cramer 1978] allows the total extinction estimate to occasionally underestimate the true extinction, but as we saw in Sec. 3, reweighting the samples to remain unbiased leads to severe variance for significantly non-bounding extinctions. Galtier et al. [2013, 2016] allow for non-bounding extinctions by allowing negative Monte Carlo weights, though at the cost of lowering the convergence rate. More recent estimators, such as Taylor series-based unbiased ray marching [Kettunen et al. 2021], debiased raymarching using an unbiased telescoping series formulation [Misso et al. 2022], and the Taylor series based pseries-cumulative and pseries-cdf estimators [Georgiev et al. 2019], trade off variance for computation time. Of these methods, debiased raymarching does not depend on null collision theory and both pseries-cumulative and unbiased raymarching can effectively bypass the need for directly specifying a total extinction. These estimators work well for low-albedo media where transmittance is the dominant source of image variance. However, for high-albedo media with multiple scattering, this tradeoff is often less worthwhile

since overall image variance tends to be dominated by other factors; here inexpensive estimators with manageable variance tend to considerably outperform expensive, but lower variance, estimators since affording more samples per pixel helps reduce other variance sources; our work focuses on these inexpensive estimators.

Progressive estimation. Our progressive medium clamping is conceptually similar to both progressive photon mapping [Hachisuka et al. 2010; Hachisuka and Jensen 2009; Hachisuka et al. 2008; Jarosz et al. 2011; Knaus and Zwicker 2011; Novák et al. 2012] and progressive many-light methods [Dachsbacher et al. 2014; Davidovič et al. 2012; Keller 1997; Walter et al. 2005]. A key difference is that bias exists in every iteration of these methods—as blurring (photon mapping) or loss of short-range and glossy transport (many-lights)—which necessitates careful error analysis to ensure a converging radius or clamping reduction formula. In contrast, our approach’s individual passes become unbiased once we reach a bounding extinction, which significantly simplifies our needed analysis and readily supports the scene-adaptive and spatially-adaptive update formulas proposed in Sec. 4.2. Misso et al. [2022] introduced a debiasing framework and showed how generalized progressive formulations can be derived directly from it. They also introduced progressive finite differences for evaluating derivatives as well as unbiased progressive photon mapping. Since individual iterations become unbiased after a finite number of steps, further debiasing is not needed.

7 RESULTS AND IMPLEMENTATION

While our goal is to eventually incorporate our adaptive techniques into our production renderer, we implemented our current prototype within PBRT [Pharr et al. 2016] for easier testing and comparison. Though we previously denoted μ_t as a scene-wide homogeneous constant for notational simplicity, in practice we store and update spatially varying bounding extinctions μ_t on a coarse-resolution grid as in prior work [Novák et al. 2014; Szirmay-Kalos et al. 2011]. We use an 80^3 grid fit to the medium extents where each grid cell, or super-voxel, stores its own estimate of μ_t for the portion of the medium it spans. In essence, μ_t becomes a piecewise function for any path through the medium. When rendering pass k we update the total extinction $\mu_t^{(k+1)}$ for the next pass independently for each super-voxel.

For all of our comparisons we choose to initialize the extinctions for all super-voxels to be near zero ($\mu_t^{(1)} = .01$) when employing our progressive approach. This corresponds to the near-worst case scenario for prior methods and we choose to do so for the purpose of conveying the robustness of our technique. In Fig. 1 we visualize just how poorly prior work performs for the same worst-case scenario by comparing our progressive technique against ratio tracking + weighted delta tracking in a scene containing a fully procedural medium. While prior methods exhibit variance which can never converge in any practical amount of time, our technique effectively removes this variance through biasing the medium, and then discovers bounding μ_t , eventually removing the bias.

In Fig. 10 we instead give ratio tracking the benefit of knowing the true bounds a-priori, and despite using a near-worst case initialization, our technique very quickly discovers bounding values for μ_t (top graph), and reports a nearly equivalent mean squared error




Figure 8: We compare the work-normalized variance (cost \times variance) of traditional ratio-tracking [Novák et al. 2014] against our importance sampling strategy (right) across a variety of canonical extinction functions (left).

to ratio tracking (bottom graph). The convergence of our technique greatly benefits from the fact that every path generated during a render can contribute to discovering bounding μ_t for all super-voxels. By the time the first few pixel samples have been evaluated, enough bounding extinctions have populated grid that any bias is both visually and numerically negligible. Our progressive approach maintains the same good performance of prior techniques while successfully making them resilient to non-bounding extinctions.

In Fig. 8 we compare classical ratio tracking with and without our importance sampling strategy (Sec. 5) for a variety of canonical extinction functions. For each one, ratio tracking with our adaptive sampling strategy results in a more efficient estimation technique in terms of work-normalized variance.

In Fig. 9 we compare our technique to the state of the art in transmittance estimation in a scene containing highly scattering media. Most of the more recent transmittance estimators such as unbiased raymarching [Kettunen et al. 2021], debiased raymarching [Misso et al. 2022], and pseries-cdf [Georgiev et al. 2019] trade lower variance transmittance estimates for higher costs and are also at least partially resilient to non-bounding extinctions. However, in highly scattering scenes such as this one which are very common in production, the majority of the noise is due to other sources than transmittance estimation so utilizing a cheaper technique like ratio tracking is currently more viable in production rendering. Our progressive approach with near-worst case initializations for μ_t still maintains similar performance to ratio tracking with perfect knowledge of bounding μ_t .

8 CONCLUSION

Our ultimate goal in this work was to alleviate one of the most glaring drawbacks of using null-scattering based techniques in production: the need to directly specify bounding extinctions μ_t . Existing methods suffer from high variance in the presence of non-bounding extinctions and high-albedo media, both of which are occurring with increasingly frequency in our productions. We introduced a novel progressive technique that makes existing null-tracking based methods resilient to non-bounding extinctions, while maintaining their computational efficiency. While our approach introduces bias, it is consistent, converging to the correct result in the limit. In doing so, we have alleviated the most notable caveat of using null-tracking in production.

We have also introduced a novel adaptive sampling routine for ratio tracking based on zero-variance estimation which we hope can motivate future improvements to other null-tracking techniques.

8.1 Future work

For future work, we plan on exploring different strategies for initializing, clamping, and updating bounding extinctions μ_t . Currently, we conservatively clamp densities to ensure the total extinction is bounding, but our analysis in Sec. 3 implies that allowing some negative null-density (as long as the medium remains only slightly non-bounding) may further reduce bias while still ensuring bounded variance.

Additionally, we plan on expanding our method to support residual trackers, which would simply involve storing additional local information other than μ_t . We also hope to apply our importance sampling idea to other null-tracking techniques.

Finally, while our current prototype is implemented in PBRT, we intend on incorporating this technique into our production renderer. Our current implementation utilizes a naive multithreading scheme for updating the bounding extinctions: separate threads write updated bounding extinctions into the super-voxel grid independently with no synchronization. The nature of our method happens to be resilient to the kind of race conditions that can result, and our implementation eventually produces bounding extinctions once consistency has been reached. At intermediate sample counts, however, the state of the super-voxel grid will be non-deterministic. A more sophisticated multithreading scheme is needed if fully deterministic intermediate behavior is desired in a production environment.

ACKNOWLEDGMENTS

The cloud model in Fig. 9 is from Walt Disney Animation Studios. This work was generously supported by NSF awards 1844538.

REFERENCES

Brent Burley, David Adler, Matt Jen-Yuan Chiang, Hank Driskill, Ralf Habel, Patrick Kelly, Peter Kutz, Yining Karl Li, and Daniel Teece. 2018. The Design and Evolution of Disney's Hyperion Renderer. *ACM Transactions on Graphics* 37, 3 (July 2018), 33:1–33:22. <https://doi.org/10/gfjm8w>

L. L. Carter, E. D. Cashwell, and W. M. Taylor. 1972. Monte Carlo Sampling with Continuously Varying Cross Sections along Flight Paths. *Nuclear Science and Engineering* 48, 4 (Aug. 1972), 403–411. <https://doi.org/10/gfzngg>

S. N. Cramer. 1978. Application of the Fictitious Scattering Radiation Transport Model for Deep-Penetration Monte Carlo Calculations. *Nuclear Science and Engineering* 65, 2 (1978), 237–253. <https://doi.org/10/gfzq74>

Carsten Dachsbacher, Jaroslav Krivánek, Miloš Hašan, Adam Arbree, Bruce Walter, and Jan Novák. 2014. Scalable Realistic Rendering with Many-Light Methods. *Computer*

Graphics Forum 33, 1 (Feb. 2014), 88–104. <https://doi.org/10/f5twgw>

Tomáš Davidovič, Iliyan Georgiev, and Philipp Slusallek. 2012. Progressive Lightcuts for GPU. In *ACM SIGGRAPH Talks*. ACM Press, 1. <https://doi.org/10/gfzq7w>

Eugene d'Eon and Jan Novák. 2021. Zero-Variance Transmittance Estimation. In *Eurographics Symposium on Rendering - DL-only Track*. Eurographics Association. <https://doi.org/10/grqvdk>

M. Galtier, S. Blanco, C. Caliot, C. Coustet, J. Dauchet, M. El Hafi, V. Eymet, R. Fournier, J. Gautrais, A. Khuong, B. Piaud, and G. Terrée. 2013. Integral Formulation of Null-Collision Monte Carlo Algorithms. *Journal of Quantitative Spectroscopy and Radiative Transfer* 125 (Aug. 2013), 57–68. <https://doi.org/10/f446pg>

Mathieu Galtier, Stéphane Blanco, Jérémie Dauchet, Mouna El Hafi, Vincent Eymet, Richard Fournier, Maxime Roger, Christophe Spiesser, and Guillaume Terrée. 2016. Radiative Transfer and Spectroscopic Databases: A Line-Sampling Monte Carlo Approach. *Journal of Quantitative Spectroscopy and Radiative Transfer* 172 (March 2016), 83–97. <https://doi.org/10/gfzq7j>

Manuel Gamito. 2018. Path Tracing the Framerstorian Way. In *Path Tracing in Production (ACM SIGGRAPH Courses)*, 52–61. <https://doi.org/10/gfz2ck>

Iliyan Georgiev, Zackary Misso, Toshiya Hachisuka, Derek Nowrouzezahrai, Jaroslav Krivánek, and Wojciech Jarosz. 2019. Integral Formulations of Volumetric Transmittance. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 38, 6 (Nov. 2019), 154:1–154:17. <https://doi.org/10/dffn>

Toshiya Hachisuka, Wojciech Jarosz, and Henrik Wann Jensen. 2010. A Progressive Error Estimation Framework for Photon Density Estimation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 29, 6 (Dec. 2010), 144:1–144:12. <https://doi.org/10/dmrtfg>

Toshiya Hachisuka and Henrik Wann Jensen. 2009. Stochastic Progressive Photon Mapping. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 28, 5 (Dec. 2009), 130:1–130:8. <https://doi.org/10/d8xxn3>

Toshiya Hachisuka, Shinji Ogaki, and Henrik Wann Jensen. 2008. Progressive Photon Mapping. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 27, 5 (Dec. 2008), 130:1–130:8. <https://doi.org/10/cn8h39>

David S. Immel, Michael F. Cohen, and Donald P. Greenberg. 1986. A Radiosity Method for Non-Diffuse Environments. *Computer Graphics (Proceedings of SIGGRAPH)* 20, 4 (Aug. 1986), 133–142. <https://doi.org/10/dmjm9t>

Wojciech Jarosz, Derek Nowrouzezahrai, Robert Thomas, Peter-Pike Sloan, and Matthias Zwicker. 2011. Progressive Photon Beams. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 30, 6 (Dec. 2011), 181:1–181:12. <https://doi.org/10/fn5xzz>

Daniel Jonsson, Joel Kronander, Jonas Unger, Thomas B. Schon, and Magnus Wrenninge. 2020. Direct Transmittance Estimation in Heterogeneous Participating Media Using Approximated Taylor Expansions. *IEEE Transactions on Visualization and Computer Graphics* (2020), 1–1. <https://doi.org/10/gj32rc>

James T. Kajiya. 1986. The Rendering Equation. *Computer Graphics (Proceedings of SIGGRAPH)* 20, 4 (Aug. 1986), 143–150. <https://doi.org/10/cvf53j>

Alexander Keller. 1997. Instant Radiosity. In *Annual Conference Series (Proceedings of SIGGRAPH)*. ACM Press, 49–56. <https://doi.org/10/fqch2z>

Markus Kettunen, Eugene d'Eon, Jacopo Pantaleoni, and Jan Novák. 2021. An Unbiased Ray-Marching Transmittance Estimator. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 40, 4 (July 2021), 137:1–137:20. <https://doi.org/10/gqjn79> arXiv:2102.10294 [cs.GR]

Claude Knaus and Matthias Zwicker. 2011. Progressive Photon Mapping: A Probabilistic Approach. *ACM Transactions on Graphics* 30, 3 (May 2011), 25:1–25:13. <https://doi.org/10/bcw2ph>

Jaroslav Krivánek and Eugene d'Eon. 2014. A Zero-Variance-Based Sampling Scheme for Monte Carlo Subsurface Scattering. In *ACM SIGGRAPH Talks*. ACM Press, 66:1–66:1. <https://doi.org/10/gfzq7n>

Christopher Kulla and Marcos Fajardo. 2012. Importance Sampling Techniques for Path Tracing in Participating Media. *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 31, 4 (June 2012), 1519–1528. <https://doi.org/10/f35f4k>

Peter Kutz, Ralf Habel, Yining Karl Li, and Jan Novák. 2017. Spectral and Decomposition Tracking for Rendering Heterogeneous Volumes. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 36, 4 (July 2017), 111:1–111:16. <https://doi.org/10/gbxjxg>

Bailey Miller, Iliyan Georgiev, and Wojciech Jarosz. 2019. A Null-Scattering Path Integral Formulation of Light Transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 38, 4 (July 2019), 44:1–44:13. <https://doi.org/10/gf6rzb>

Zackary Misso, Benedikt Bitterli, Iliyan Georgiev, and Wojciech Jarosz. 2022. Unbiased and Consistent Rendering Using Biased Estimators. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 41, 4 (July 2022), 48:1–48:13. <https://doi.org/10/gqjn66>

Zackary Misso, Yining Karl Li, Brent Burley, Daniel Teece, and Wojciech Jarosz. 2023. *Progressive null-tracking for volumetric rendering supplemental code and data*. <https://doi.org/10.5281/zenodo.7893499>

Jan Novák, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz. 2018. Monte Carlo Methods for Volumetric Light Transport Simulation. *Computer Graphics Forum (Proceedings of Eurographics State of the Art Reports)* 37, 2 (May 2018), 551–576. <https://doi.org/10/gd2jqj>

Jan Novák, Derek Nowrouzezahrai, Carsten Dachsbacher, and Wojciech Jarosz. 2012. Progressive Virtual Beam Lights. *Computer Graphics Forum (Proceedings of the*

- Eurographics Symposium on Rendering* 31, 4 (June 2012), 1407–1413. <https://doi.org/10/gfzndw>
- Jan Novák, Andrew Selle, and Wojciech Jarosz. 2014. Residual Ratio Tracking for Estimating Attenuation in Participating Media. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 33, 6 (Nov. 2014), 179:1–179:11. <https://doi.org/10/f6r2nq>
- Ken H. Perlin and Eric M. Hoffert. 1989. Hypertexture. *Computer Graphics (Proceedings of SIGGRAPH)* 23, 3 (July 1989), 253–262. <https://doi.org/10/fdmsxd>
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically Based Rendering: From Theory to Implementation* (3rd ed.). Morgan Kaufmann, Cambridge, MA.
- Peter Shirley, Changyaw Wang, and Kurt Zimmerman. 1996. Monte Carlo Techniques for Direct Lighting Calculations. *ACM Transactions on Graphics* 15, 1 (Jan. 1996), 1–36. <https://doi.org/10/ddgbgg>
- László Szirmay-Kalos, Iliyan Georgiev, Milán Magdics, Balázs Molnár, and Dávid Légrády. 2017. Unbiased Light Transport Estimators for Inhomogeneous Participating Media. *Computer Graphics Forum (Proceedings of Eurographics)* 36, 2 (2017), 9–19. <https://doi.org/10/gbm5tr>
- László Szirmay-Kalos, Milán Magdics, and Mateu Sbert. 2018. Multiple Scattering in Inhomogeneous Participating Media Using Rao-Blackwellization and Control Variates. *Computer Graphics Forum (Proceedings of Eurographics)* 37, 2 (2018), 63–74. <https://doi.org/10.1111/cgf.13342>
- László Szirmay-Kalos, Balázs Tóth, and Milán Magdics. 2011. Free Path Sampling in High Resolution Inhomogeneous Participating Media. *Computer Graphics Forum* 30, 1 (March 2011), 85–97. <https://doi.org/10/bpkdzm>
- Eric Veach and Leonidas J. Guibas. 1995. Optimally Combining Sampling Techniques for Monte Carlo Rendering. In *Annual Conference Series (Proceedings of SIGGRAPH)*, Vol. 29. ACM Press, 419–428. <https://doi.org/10/d7b6n4>
- Bruce Walter, Sebastian Fernandez, Adam Arbree, Kavita Bala, Michael Donikian, and Donald P Greenberg. 2005. Lightcuts: A Scalable Approach to Illumination. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 24, 3 (Aug. 2005), 1098–1107. <https://doi.org/10/dhp5d3>
- E. R. Woodcock, T. Murphy, P. J. Hemmings, and T. C. Longworth. 1965. Techniques Used in the GEM Code for Monte Carlo Neutronics Calculations in Reactors and Other Systems of Complex Geometry. In *Applications of Computing Methods to Reactor Problems*. Argonne National Laboratory.
- Yonghao Yue, Kei Iwasaki, Bing-Yu Chen, Yoshinori Dobashi, and Tomoyuki Nishita. 2010. Unbiased, Adaptive Stochastic Sampling for Rendering Inhomogeneous Participating Media. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 29, 6 (Dec. 2010), 177:1–177:8. <https://doi.org/10/bp73xb>




Figure 9: An equal extinction lookups comparison of our progressive method initialized with near-zero total extinction versus many of the state-of-the-art transmittance estimators given bounding extinctions for a highly scattering medium. All methods employ weighted-delta tracking for free-flight sampling. MSE values are computed as the average MSE across the entire image.

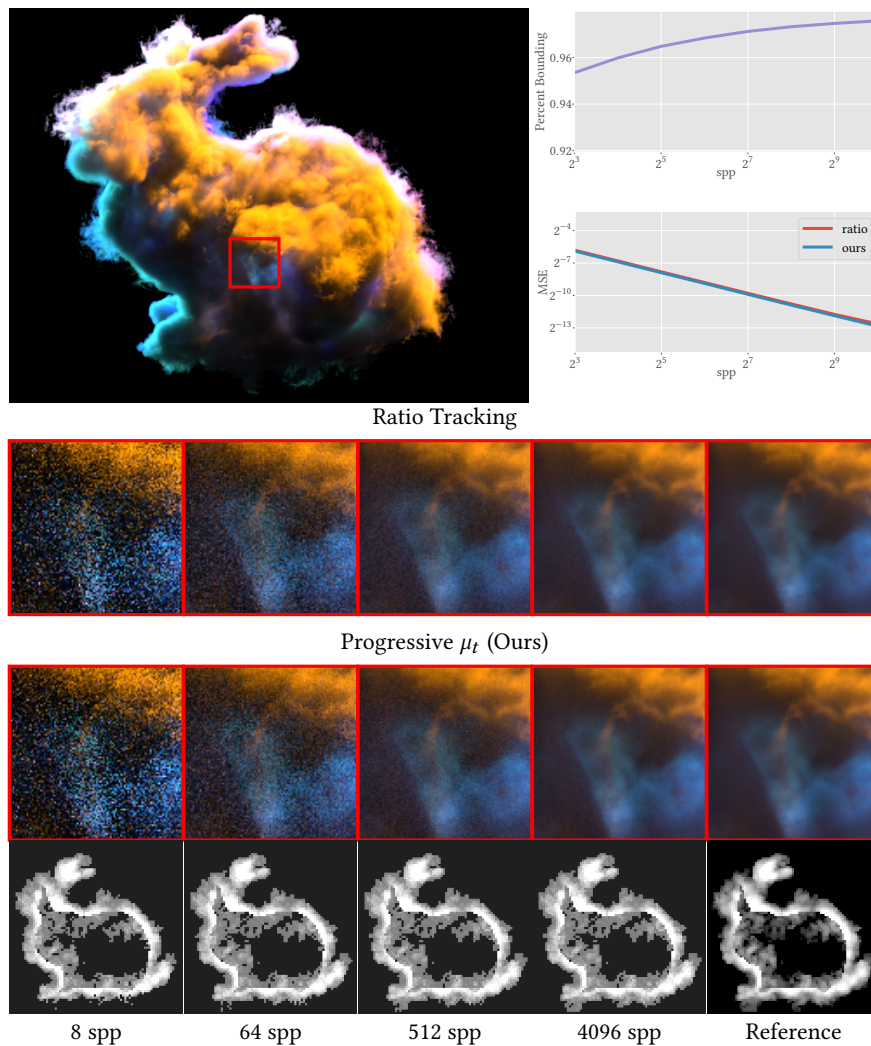


Figure 10: We compare the performance of our progressive estimator to ratio-tracking while also showing the convergence of our super-voxel grid. We visualize a slice of the super-voxel grid (bottom row) throughout a render. We also compare the convergence of the MSE of our method to ratio tracking (bottom graph) and convey the convergence of our super-voxel grid by plotting the percentage of voxels with bounding μ_t 's (top graph).