



Delft University of Technology

Jackknife Transmittance and MIS Weight Estimation

Peters, Christoph

DOI

[10.1145/3763273](https://doi.org/10.1145/3763273)

Publication date

2025

Document Version

Final published version

Published in

ACM Transactions on Graphics

Citation (APA)

Peters, C. (2025). Jackknife Transmittance and MIS Weight Estimation. *ACM Transactions on Graphics*, 44(6), Article 204. <https://doi.org/10.1145/3763273>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



 Latest updates: <https://dl.acm.org/doi/10.1145/3763273>

RESEARCH-ARTICLE

Jackknife Transmittance and MIS Weight Estimation

CHRISTOPH PETERS, Delft University of Technology, Delft, Zuid-Holland, Netherlands

Open Access Support provided by:

Delft University of Technology



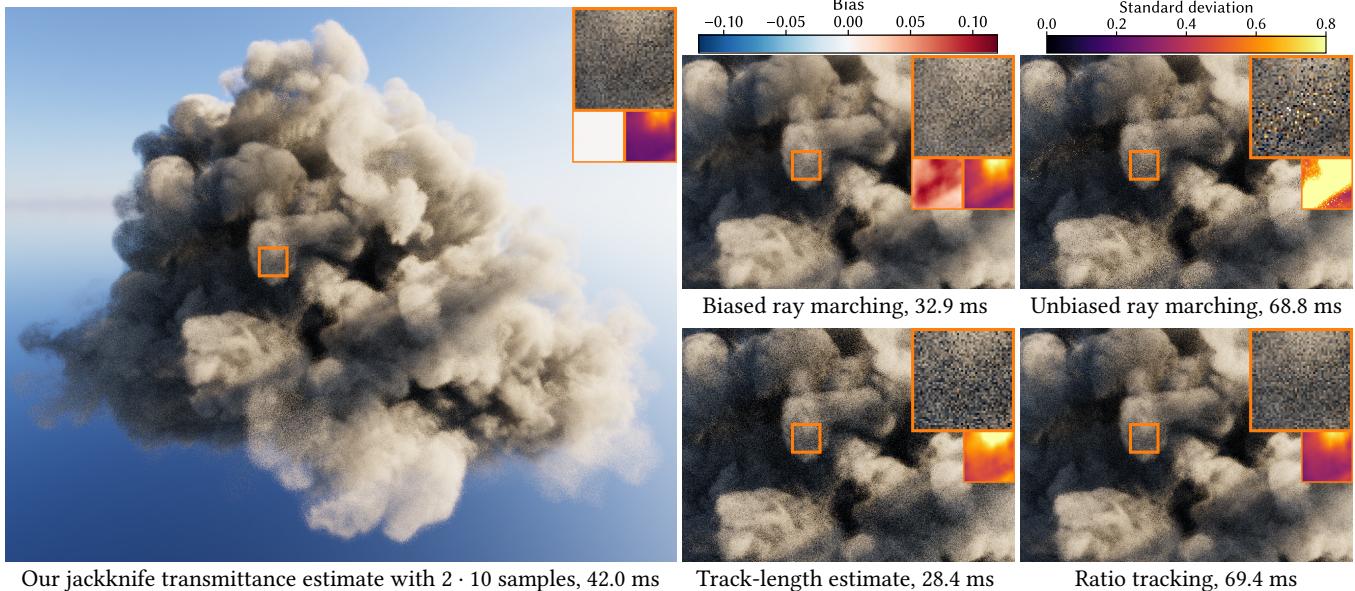
PDF Download
3763273.pdf
18 December 2025
Total Citations: 0
Total Downloads: 96

Published: 01 December 2025
Accepted: 09 August 2025
Received: 23 May 2025

[Citation in BibTeX format](#)

Jackknife Transmittance and MIS Weight Estimation

CHRISTOPH PETERS, Delft University of Technology, Netherlands



Our jackknife transmittance estimate with $2 \cdot 10$ samples, 42.0 ms

Track-length estimate, 28.4 ms

Ratio tracking, 69.4 ms

Fig. 1. A rendering of the Intel cloud under direct illumination from a light probe at 32 samples per pixel (spp). While biased ray marching [Kettunen et al. 2021] overestimates transmittance visibly, the bias of our jackknife transmittance estimator is negligible. Additionally, it is less prone to fireflies and faster than unbiased ray marching [Kettunen et al. 2021] and has lower variance than the track-length estimate and ratio tracking.

A core operation in Monte Carlo volume rendering is transmittance estimation: Given a segment along a ray, the goal is to estimate the fraction of light that will pass through this segment without encountering absorption or out-scattering. A naive approach is to estimate optical depth τ using unbiased ray marching and to then use $\exp(-\tau)$ as transmittance estimate. However, this strategy systematically overestimates transmittance due to Jensen's inequality. On the other hand, existing unbiased transmittance estimators either suffer from high variance or have a cost governed by random decisions, which makes them less suitable for SIMD architectures. We propose a biased transmittance estimator with significantly reduced bias compared to the naive approach and a deterministic and low cost. We observe that ray marching with stratified jittered sampling results in estimates of optical depth that are nearly normal-distributed. We then apply the unique minimum variance unbiased (UMVU) estimator of $\exp(-\tau)$ based on two such estimates (using two different sets of random numbers). Bias only arises from violations of the assumption of normal-distributed inputs. We further reduce bias and variance using a variance-aware importance sampling scheme. The underlying theory can be used to estimate any analytic function of optical depth. We use this generalization to estimate multiple importance sampling (MIS) weights and introduce two integrators: Unbiased MIS with biased MIS weights and a more efficient but biased combination of MIS and transmittance estimation.

Author's Contact Information: Christoph Peters, c.j.p.peters@tudelft.nl, Delft University of Technology, Delft, Netherlands.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 1557-7368/2025/12-ART204

<https://doi.org/10.1145/3763273>

CCS Concepts: • Computing methodologies → Visibility; Ray tracing; • Mathematics of computing → Probabilistic inference problems.

Additional Key Words and Phrases: transmittance estimation, volume rendering, path tracing, multiple importance sampling, MIS, statistics, UMVU, normal-distributed optical depth

ACM Reference Format:

Christoph Peters. 2025. Jackknife Transmittance and MIS Weight Estimation. *ACM Trans. Graph.* 44, 6, Article 204 (December 2025), 16 pages. <https://doi.org/10.1145/3763273>

1 Introduction

Physically-based volume rendering for participating media such as fog, steam or smoke is challenging. In the presence of such media, radiance is no longer constant along a ray between two surface interactions. Instead, we have to estimate integrals along the ray to account for absorption, out-scattering, in-scattering and emission of light. In many aspects, these effects can be handled similarly to surface rendering, e.g. by using a unidirectional path tracer. A major difference is that distance sampling and transmittance estimation take the place of ray tracing [Novák et al. 2018]. Distance sampling determines the distance to the next path vertex, whereas transmittance estimation computes the fraction of light that passes through the volume along a given ray segment. For heterogeneous (i.e. spatially varying) media, these two operations contribute substantially

to the overall cost of volume rendering since they are the main way in which the renderer interacts with the volume data.

Our main contribution in this work is an efficient and GPU-friendly solution for transmittance estimation. Let the extinction along a ray segment be given by $\mu(t)$ where $t \in [0, t_{\max}] \subset \mathbb{R}$ is the ray parameter. Then the transmittance is

$$T := T(t_{\max}) := \exp(-\tau), \text{ where } \tau := \tau(t_{\max}) := \int_0^{t_{\max}} \mu(t) dt. \quad (1)$$

The integral τ is called optical depth. The naive approach to estimate T is to compute an unbiased Monte Carlo estimate X of τ , e.g. by using ray marching with uniform jittered sampling [Kettunen et al. 2021; Pauly et al. 2000] and to use $\exp(-X)$ as transmittance estimate. However, this strategy systematically overestimates the transmittance, because by Jensen's inequality the expected value satisfies

$$E(\exp(-X)) \geq \exp(E(-X)) = \exp(-\tau) = T.$$

There are three major classes of unbiased transmittance estimators: Regular tracking [Amanatides and Woo 1987; Szirmay-Kalos et al. 2010] incurs a high cost to compute τ exactly. Null-collision methods [Coleman 1968; Novák et al. 2014] stochastically traverse a homogenized medium and estimate transmittance based on the difference to the original medium. Unbiased ray marching [Georgiev et al. 2019; Kettunen et al. 2021] combines a stochastically chosen and unbounded number of independent optical depth estimates to estimate the power series $\exp(-\tau) = \sum_{j=0}^{\infty} \frac{(-\tau)^j}{j!}$. Regular tracking has a predictable but high cost. The other two approaches are prone to divergent execution on GPUs since the amount of work that is conducted is random. Additionally, null-collision methods have relatively high variance.

At the same time, real-time volume rendering on GPUs is becoming viable, increasing the demand for efficient and GPU-friendly algorithms [Hofmann et al. 2021, 2023; Schneider 2023]. In this context, a small amount of bias is acceptable. However, the estimate $\exp(-X)$ seems suboptimal. The fact that its bias is always positive hints at the possibility for a better estimate. The bias grows with the variance of X . If we could somehow get a good estimate of the bias, we could compensate for it to form a better transmittance estimate.

To accomplish this, we rely on a single strong assumption: We assume that estimates of optical depth are normal-distributed, i.e. distributed with Gaussian density (Sec. 3.1). Under this assumption the unique minimum variance unbiased (UMVU) estimate of $\exp(-\tau)$ is known. It has been derived using the so-called generalized jackknife [Gray et al. 1973] (Sec. 3.2). We use it with two unbiased, independent and identically distributed (i.i.d.) estimates of optical depth X_0, X_1 , i.e. we estimate the optical depth τ twice in the same way, but using different random numbers. As a result of this choice, we obtain the following stunningly simple formula for our novel jackknife transmittance estimate (Sec. 3.3):

$$T \approx \cos\left(\frac{X_0 - X_1}{2}\right) \exp\left(-\frac{X_0 + X_1}{2}\right).$$

Under the assumption of normal-distributed optical depth estimates, this estimate is unbiased. In practice, we expect this assumption to be violated, which results in bias. However, we can keep this

bias small by using stratified jittered sampling instead of uniform jittered sampling [Pauly et al. 2000] (Sec. 4.1). This way, the optical depth estimate is a sum of independent random variables and the central limit theorem [Knight 1999, p. 145] implies that it approaches a normal distribution as the sample count grows. To make this reasoning more reliable in the presence of sparse volumes, we propose a variance-aware [Pantaleoni and Heitz 2017] importance sampling scheme (Sec. 4.2) and we keep the sample count fixed (Sec. 4.3).

The underlying theory [Gray et al. 1973] of our estimate is not restricted to estimating $\exp(-\tau)$. It can estimate $g(\tau)$ for any analytic function $g : \mathbb{C} \rightarrow \mathbb{C}$ (Sec. 5.1). We use this capability to implement multiple importance sampling (MIS) for distance sampling techniques, e.g. to combine free-flight distance sampling with equiangular sampling [Kulla and Fajardo 2012]. A well-known problem in doing so is that the probability density of free-flight sampling is $T(t)\mu(t)$ and is thus not known exactly [Miller et al. 2019; Novák et al. 2018]. We address this problem in two ways (Sec. 5.2): Either we estimate MIS weights in a biased fashion, which still results in an unbiased MIS estimate [Veach and Guibas 1995], or we estimate the product of MIS weight, reciprocal density and transmittance directly, which reduces the overall cost but introduces slight bias.

Our transmittance estimate has considerably lower bias than existing biased alternatives [Kettunen et al. 2021], lower variance than the track-length estimate or ratio tracking [Novák et al. 2014] and a lower cost than unbiased ray marching [Kettunen et al. 2021] (Sec. 6.1). Our MIS weight estimation offers considerably reduced variance compared to a technique based on ratio tracking [Miller et al. 2019] at equal sample count (Sec. 6.2). The source code of our GPU implementation is available on the project webpage.

2 Background and Related Work

The goal of volume rendering is to solve the radiative transfer equation. Let $\mathbf{x}(t) \in \mathbb{R}^3$ be the point at ray parameter t along the ray considered in Eq. 1 and let $\omega \in \mathbb{S}^2$ be the ray direction, where $\mathbb{S}^2 \subset \mathbb{R}^3$ is the unit sphere. Let ρ denote the phase function of the volume. Disregarding volume emission for simplicity, the incoming radiance at $\mathbf{x}(0)$ satisfies the following integral equation [Chandrasekhar 1950; Novák et al. 2018]:

$$L_i(\mathbf{x}(0), \omega) = \int_0^{\infty} T(t)\mu(t) \int_{\mathbb{S}^2} \rho(\mathbf{x}(t), -\omega, \omega_i) L_i(\mathbf{x}(t), \omega_i) d\omega_i dt \quad (2)$$

It is a transmittance-weighted integral over in-scattered radiance at each point along the ray. For rays that hit a surface, this equation is coupled with the usual rendering equation, including emission.

A simple unidirectional path tracer proceeds in the following steps to estimate this radiance [Novák et al. 2018]: First it samples a distance t proportional to $T(t)\mu(t)$ such that these factors cancel with the density. Then it samples a direction ω_i proportional to the phase function [Jendersie and d'Eon 2023; Witt 1977] and estimates the incoming radiance $L_i(\mathbf{x}(t), \omega_i)$ recursively. As with surface rendering, next-event estimation is a useful strategy to sample directions towards light sources at each vertex of the generated path. In this case, the transmittance from the path vertex to the light source must be estimated explicitly [Kettunen et al. 2021; Novák

et al. 2014]. Compared to surface rendering, distance sampling plays a similar role as closest-hit ray tracing and transmittance estimation replaces any-hit ray tracing for shadow rays. There are many alternative strategies for path construction [Novák et al. 2018], but they are orthogonal to our work.

Through more powerful graphics hardware and recent work on neural denoising [Hofmann et al. 2021, 2023], GPU-accelerated real-time volume rendering is becoming viable. The underlying volume data can be specified procedurally and/or using voxel grids [Schneider 2023]. OpenVDB is a widely used data structure for voxel data, which uses a shallow hierarchy of sparse grids [Museth 2013]. NanoVDB is a stripped down version designed for access to volume data from GPUs [Museth 2021].

In the following, we discuss the two main operations through which a volume renderer interacts with the volume data in more detail: Distance sampling (Sec. 2.1) and transmittance estimation (Sec. 2.2). Then we discuss the challenges of using MIS [Veach and Guibas 1995] for distance sampling and existing solutions (Sec. 2.3).

2.1 Distance Sampling

Free-flight distance sampling uses the density $T(t)\mu(t)$ and is the most common strategy [Novák et al. 2018]. It provides perfect importance sampling for these factors and sidesteps the need for transmittance estimation since the transmittance cancels with the density. If $\xi \in [0, 1]$ is sampled uniformly, the distance t where $T(t) = 1 - \xi$ has the desired density.

Regular tracking computes this distance exactly, but at a high cost. It assumes a piecewise homogeneous volume, e.g. a voxel grid with nearest-neighbor interpolation. Using a digital differential analyzer (DDA) [Amanatides and Woo 1987] it steps along the ray one voxel at a time up to the point where $T(t) = 1 - \xi$ [Szirmay-Kalos et al. 2010]. A biased alternative is ray marching with fixed step size.

Delta tracking [Coleman 1968; Novák et al. 2018] uses an upper bound $\hat{\mu} \geq \mu(t)$ for the extinction. It then samples a tentative distance according to a homogenized medium with extinction $\hat{\mu}$ in closed form. With probability $\frac{\mu(t)}{\hat{\mu}}$ this distance is returned. Otherwise, it is treated as null-scattering event, meaning that the distance sampling procedure starts over from this distance. Null-scattering events can be interpreted as collisions with perfectly forward-scattering particles. A constant upper bound $\hat{\mu}$ may be too large for this method to be efficient. Szirmay-Kalos et al. [2010] instead prepare a low-resolution super-voxel grid, where each super-voxel stores a local maximum extinction and use regular tracking for this low-resolution grid. A k -d-tree with adaptive subdivision provides even better bounds [Yue et al. 2011].

Free-flight distance sampling becomes ineffective when the in-scattered radiance has high dynamic range. That happens for example when there are point lights in the volume, since their geometry terms approach infinity. To counteract that, equiangular sampling uses these geometry terms as density [Kulla and Fajardo 2012]. It can be generalized to area lights by first sampling a point on them.

2.2 Transmittance Estimation

With the sampling density of equiangular sampling, the transmittance $T(t)$ does not cancel in the Monte Carlo estimate, so it must

be estimated explicitly. The same holds true when next-event estimation is used to be able to deal with small or distant light sources.

Regular tracking as described above [Novák et al. 2018; Szirmay-Kalos et al. 2010] computes the exact optical depth τ along a ray and thus it can also compute the exact transmittance $\exp(-\tau)$, but at a high cost. As explained in the introduction, a biased alternative is to use a form of ray marching to obtain an unbiased estimate X of τ and to use $\exp(-X)$ as naive transmittance estimate.

Any free-flight distance sampling method can be turned into a transmittance estimator since it produces a sample $t > t_{\max}$ with probability $T(t_{\max})$. That gives rise to the track-length estimator, which is 1 if $t > t_{\max}$ and 0 otherwise. This estimate is unbiased, and when combined with delta tracking, its overhead is relatively low. However, it has the worst possible variance among unbiased estimates that return a transmittance in $[0, 1]$, namely $(1 - T)T$.

To reduce the variance, ratio tracking [Novák et al. 2014] replaces the random termination of delta tracking by accumulation of continuation probabilities. If tentative distances $t'_1, \dots, t'_k \in [0, t_{\max}]$ were sampled, the transmittance estimate is $\prod_{j=1}^{k-1} 1 - \frac{\mu(t'_j)}{\hat{\mu}(t'_j)}$. Residual ratio tracking [Novák et al. 2014] additionally stores a low-resolution extinction estimate for which the transmittance is computed exactly and only uses ratio tracking to estimate transmittance for the residual. These techniques force the traversal to continue longer than the track-length estimator, thus increasing the cost. Ratio tracking has lower variance than the track-length estimate when extinction bounds are inaccurate, but the same variance for perfect bounds.

Georgiev et al. [2019] develop the theory for a new class of transmittance estimators. The product of $j \in \mathbb{N}$ unbiased i.i.d. estimates of optical depth is an unbiased estimate of τ^j . That gives rise to an unbiased estimate of the power series $\exp(-\tau) = \sum_{j=0}^{\infty} \frac{(-\tau)^j}{j!}$ using a random and unbounded number of estimates. The unbiased ray marching transmittance estimator [Kettunen et al. 2021] is a practical realization of this theory. All possible permutations of estimates are treated equally to lower the variance (so-called U-statistics) and one estimate serves as expansion point for the power series. Additionally, aggressive Russian roulette limits the number of used estimates to just one most of the time, since more estimates give diminishing returns. Misso et al. [2022] introduce a similar method for pink-noise transmittance. They also propose an unbiased transmittance estimator that applies Russian roulette to a telescoping sum, which starts with a ray marching estimate and then corrects that using estimates with doubling sample counts.

On GPUs and other SIMD architectures, this Russian roulette causes thread divergence: Each thread in a SIMD group pays the same cost as the worst thread. Besides, it is tricky to allocate space for an unbounded number of estimates [Kettunen et al. 2021]. To estimate optical depth, Kettunen et al. use an adaptive sample count and importance sampling based on mean values in a super-voxel grid. This way, they also obtain an improved version of the naive biased transmittance estimate $\exp(-X)$.

2.3 Multiple Importance Sampling

MIS [Veach and Guibas 1995] is essential to combine sampling strategies in a way that retains their strengths. It may, for example, be

used to combine a distance sampling strategy with known probability density function $p_0(t)$ such as equiangular sampling with free-flight distance sampling, where $p_1(t) = T(t)\mu(t)$. When t_0, t_1 are sampled proportional to $p_0(t), p_1(t)$, respectively, we get the following MIS estimate:

$$\int_0^{t_{\max}} f(t) dt \approx \sum_{j=0}^1 w_j(t_j) \frac{f(t_j)}{p_j(t_j)}, \text{ where } w_j(t) := \frac{p_j(t)}{p_0(t) + p_1(t)}.$$

Here we have used the balance heuristic to define the MIS weights $w_j(t)$. However, it is important to note, that the MIS estimate is unbiased for any choice as long as $w_0(t) + w_1(t) = 1$ [Veach and Guibas 1995].

A well-known problem in this context is that costly regular tracking is the only way to compute the density $p_1(t) = T(t)\mu(t)$ exactly. When we merely have an unbiased estimate of $p_1(t)$, obtained e.g. using ratio tracking, the estimate of the MIS weights will still be biased since division is a non-linear operation. Additionally, it will inherit variance from the transmittance estimate. On this subject, Novak et al. [2018] write: “Since MIS only needs weights that sum up to one to be correct, it is possible to instead use crude approximations of the free-flight PDF. This may however deteriorate the variance of the combined estimator, so one may prefer to use deterministic transmittance estimators in scenarios where an algorithm depends heavily on MIS.”

Miller et al. [2019] dismiss biased MIS weights entirely and instead propose to work with a path space that treats null-scattering events as actual path vertices in a perfectly forward scattering medium. For ray segments that were not already generated by delta tracking, they use ratio tracking to produce null-scattering events. With this reinterpretation, all sampling densities are known exactly and MIS is applicable. However, the resulting MIS weights are different from the ones defined above and vary based on the random sample of null-scattering events. If ratio tracking has high variance in the rendered volume, the MIS estimates inherit this variance.

3 Our Transmittance Estimator

We now develop our novel transmittance estimator. To achieve a low and predictable computational cost, we sacrifice unbiasedness by making the assumption that optical depth estimates are normal-distributed (Sec. 3.1). Under this assumption, the optimal estimate of $\exp(-\tau)$ has been derived by Gray et al. [1973] (Sec. 3.2). We specialize their results to arrive at our jackknife transmittance estimator (Sec. 3.3). Suitable estimators of optical depth, which come close to being normal-distributed, are developed in Sec. 4.

3.1 Our Approach

We have discussed shortcomings of existing unbiased transmittance estimators in Sec. 2.2. Regular tracking is the only option with a deterministic cost, but this cost is the worst case. For delta tracking, ratio tracking [Novák et al. 2014] and unbiased ray marching [Kettunen et al. 2021], the number of samples of $\mu(t)$ is random and unbounded. This randomness incurs incoherent execution on SIMD architectures, which increases the overall cost. We conjecture that this is an inherent consequence of the lack of assumptions on the

extinction profile $\mu(t)$ and the fact that the exponential is a transcendental function. The interpretation of the problem as integral in an infinite-dimensional domain supports this conjecture [Georgiev et al. 2019].

If this is true, we have to make stronger assumptions to arrive at a technique with bounded cost. We seek a model of the distribution of optical depth estimates that is simple enough to compute our transmittance estimate efficiently without any additional assumptions. A normal distribution satisfies this requirement as we will see shortly. At the same time, we want our model to be reasonably close to the ground truth distribution. The reason why the normal distribution is of such great importance in statistics is the central limit theorem: Under relatively weak assumptions, a sum of independent random variables approaches a normal distribution as the number of random variables grows [Knight 1999, p. 145]. In Sec. 4.1, we demonstrate that this applies to a ray marching estimate of optical depth, as long as we jitter each sample independently. Thus, the normal distribution is a pragmatic and natural choice. It is not a perfect model, and therefore our estimate remains biased, but it is good enough to achieve much lower bias than with the naive estimate $\exp(-X)$.

3.2 Unique Minimum Variance Unbiased Estimation

To construct a better transmittance estimate than $\exp(-X)$, we need more than one estimate of optical depth. Thus, we compute multiple i.i.d. estimates of optical depth, i.e. we evaluate the same estimate multiple times using different random numbers. Based on our approach, we assume that each estimate is normal-distributed with unknown mean τ and unknown standard deviation. We then seek an unbiased estimate of the transmittance $T = \exp(-\tau)$. Additionally, we want this estimate to have minimal variance. In short, we want the UMVU estimate. UMVU estimates are well-studied in statistics. They are known for many different statistics and different distributions of X [Voinov and Nikulin 1993]. We found the solution for our specific case in the literature [Gray et al. 1973, Example 4]:

Theorem 1. *Let $m \in \mathbb{N}$ with $m \geq 2$ and let X_0, \dots, X_{m-1} be i.i.d. samples of a normal distribution with unknown mean $\tau \in \mathbb{R}$ and unknown standard deviation $\sigma \geq 0$. Consider the sample mean, biased sample variance and biased sample standard deviation defined by*

$$\bar{X} := \frac{1}{m} \sum_{j=0}^{m-1} X_j, \quad S^2 := \frac{1}{m} \sum_{j=0}^{m-1} X_j^2 - \bar{X}^2, \quad S := \sqrt{S^2}.$$

Let

$$K := \Gamma\left(\frac{m-1}{2}\right) \left(\frac{2}{S}\right)^{\frac{m-3}{2}} J_{\frac{m-3}{2}}(S) \exp(-\bar{X}), \quad (3)$$

where Γ denotes the gamma function and $J_{\frac{m-3}{2}}$ denotes the Bessel function of order $\frac{m-3}{2}$ [Akhmedova and Akhmedov 2019, p. 44]. Then K is the UMVU estimate of $\exp(-\tau)$, i.e. $E(K) = \exp(-\tau)$ and if another function $L(X_0, \dots, X_{m-1})$ satisfies $E(L(X_0, \dots, X_{m-1})) = \exp(-\tau)$, its variance cannot be lower:

$$V(K) \leq V(L(X_0, \dots, X_{m-1})).$$

As is, Eq. 3 requires us to evaluate the Bessel function, but we will eliminate the need for that shortly. Other than that, this looks like a straight-forward solution to our problem with all the properties that we wanted. The naive estimate $\exp(-\bar{X})$ is a factor in the estimate

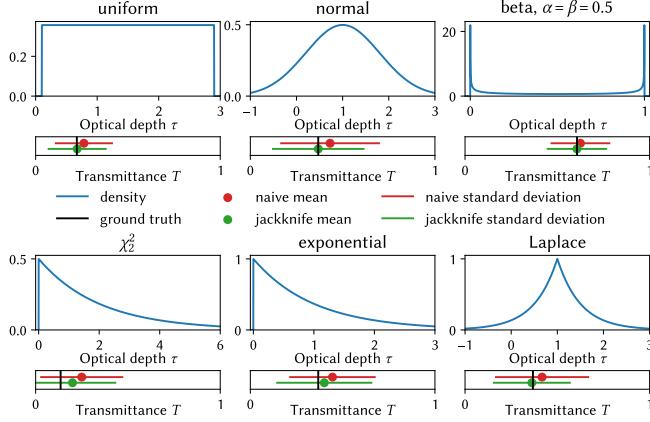


Fig. 2. A comparison of the jackknife estimate against the naive estimate $\exp(-\bar{X})$ using $m = 2$ i.i.d. optical depth estimates. We use various hypothetical distributions of optical depth estimates and compare the two transmittance estimates in terms of their mean and standard deviation. The jackknife estimate (Eq. 3) always reduces the bias.

K , but there are additional factors depending on the sample standard deviation S , which compensate for the bias of this estimate.

To derive Eq. 3, Gray et al. [1973] apply the “generalized jackknife” to a biased estimate, thus removing the bias, and then use so-called Rao-Blackwellization to arrive at the UMVU estimate. The jackknife in turn is called that because it is “a useful tool in a variety of situations” [Gray et al. 1973]. Therefore, we refer to the estimate in Eq. 3 as jackknife estimate.

If our assumption of normal-distributed estimates holds true, the jackknife estimate is unbiased. Fig. 2 demonstrates that it still achieves lower bias than $\exp(-\bar{X})$ for a variety of other distributions. In these experiments, the improvement is smallest for single-sided distributions. The Laplace distribution is remarkable in that the jackknife estimate slightly underestimates the transmittance. By Jensen’s inequality, the naive estimate can never do so. In terms of their standard deviation, both methods perform similarly.

3.3 Jackknife Transmittance Estimation

We now introduce our jackknife transmittance estimator. The main design decision that we still have to make to complete this part of our technique is what number of i.i.d. optical depth estimates m we should use. We have to use at least $m = 2$ estimates, because otherwise it is not possible to estimate the sample variance S^2 . On the other hand, we would like to use as few independent estimates as possible, for two reasons: If we use fewer estimates, we can afford more ray marching steps per estimate. Thus, the reasoning that these estimates, as sum of many independent samples, are normal-distributed is more sound and we expect less bias. Besides, \bar{X} converges at a rate of $\frac{1}{\sqrt{m}}$, whereas stratified jittered sampling with $N \in \mathbb{N}$ samples has a convergence rate closer to $\frac{1}{N}$ (Fig. 5). Therefore, it is preferable to invest the available sample budget into more ray marching steps. The data in Table 1 (and in the supplemental) support this reasoning: The choice $m = 2$ is optimal in terms of bias and standard deviation.

Table 1. Statistics about the accuracy of the estimate in Eq. 3 when estimating transmittance for the extinction profile $\mu(t)$ shown in Fig. 4a. Each estimate uses $mN = 24$ samples of $\mu(t)$ overall, but using the maximal number of samples for ray marching with $m = 2$ is optimal. The bias for $m = 2$ is below the margin of error with 10^8 trials. The row for $m = 1$ uses the naive estimate $\exp(-\bar{X})$, which has much higher bias but slightly smaller standard deviation (std.).

Estimate count m	Ray marching sample count N	Transmittance estimate		
		mean	bias	std.
2	12	0.13411	0.00000	0.0405
3	8	0.13425	0.00014	0.0440
4	6	0.13421	0.00010	0.0528
6	4	0.13433	0.00022	0.0490
8	3	0.13485	0.00074	0.0694
1	24	0.13797	0.00386	0.0331

This choice also has a third advantage in that it simplifies evaluation of Eq. 3. We know [Akhmedova and Akhmedov 2019, p. 49] that

$$J_{-\frac{1}{2}}(S) = \sqrt{\frac{2}{\pi S}} \cos(S), \quad \Gamma\left(\frac{1}{2}\right) = \sqrt{\pi}.$$

Then for $m = 2$

$$K = \sqrt{\pi} \left(\frac{2}{S}\right)^{-\frac{1}{2}} \sqrt{\frac{2}{\pi S}} \cos(S) \exp(-\bar{X}) = \cos(S) \exp(-\bar{X}).$$

Furthermore,

$$S = \sqrt{\frac{X_0^2 + X_1^2}{2} - \left(\frac{X_0 + X_1}{2}\right)^2} = \sqrt{\frac{X_0^2 - 2X_0X_1 + X_1^2}{4}} = \frac{|X_0 - X_1|}{2}.$$

With that, we have derived the core of our method. Given two unbiased i.i.d. estimates of optical depth X_0, X_1 , our jackknife transmittance estimate is

$$K = \cos(S) \exp(-\bar{X}), \text{ where } \bar{X} = \frac{X_0 + X_1}{2}, \quad S = \frac{|X_0 - X_1|}{2}. \quad (4)$$

The computational cost is negligible. It is quite surprising that the cosine shows up in this manner, but even without diving into the proof of Thm. 1, we can convince ourselves that it is plausible. The estimate $\exp(-\bar{X})$ systematically overestimates the transmittance. The factor $\cos(S) \in [-1, 1]$ reduces this estimate to compensate for this bias. For small sample standard deviation S , the correction factor remains close to 1. For large standard deviations, our transmittance estimate can become negative. That is an undesirable property, but we found these negative values to be rare unless the transmittance is close to zero (see the supplemental document) and clamping them away causes increased bias. Sec. 5.1 generalizes Eq. 4 and provides another explanation why $\cos(S)$ arises in this formula.

In Appendix A.2, we derive the variance of our estimate for $m = 2$ and normal-distributed X_0, X_1 with standard deviation $\sigma \geq 0$:

$$V(K) = \frac{1}{2}(\exp(\sigma^2) - 1) \exp(-2\tau).$$

The corresponding relative root mean square error (rRMSE) is

$$\frac{\sqrt{V(K)}}{T} = \sqrt{\frac{\exp(\sigma^2) - 1}{2}}.$$

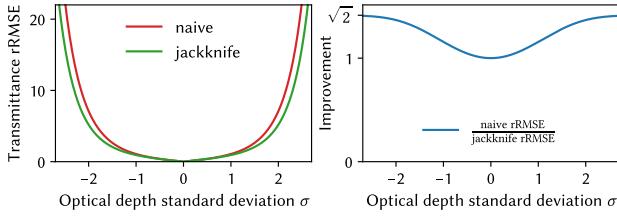


Fig. 3. Left: The rRMSE of our jackknife transmittance estimator and the naive estimate $\exp(-\bar{X})$ for normal-distributed optical depth estimates with standard deviation σ . Right: Their ratio.

In Appendix A.1, we derive the rRMSE of $\exp(-\bar{X})$ for the same two samples:

$$\sqrt{\exp(\sigma^2) - 2 \exp\left(\frac{\sigma^2}{4}\right) + 1}.$$

Fig. 3 shows plots of these two functions and their ratio. We observe that the error of our method is never worse and lower by a factor of nearly $\sqrt{2}$ for large σ . We note however, that this comparison is not entirely fair, since the naive estimate could also use $m = 1$ to benefit from the faster convergence rate of stratified jittered sampling.

4 Our Optical Depth Estimator

We now have to design an optical depth estimator that fits the requirements of our jackknife transmittance estimator well. First, we elaborate on our reasoning that stratified jittered sampling results in nearly normal-distributed estimates (Sec. 4.1). Since this reasoning can fail for sparse volumes, we use importance sampling with a variance-minimizing strategy [Pantaleoni and Heitz 2017] (Sec. 4.2). Finally, we describe non-trivial aspects of the implementation (Sec. 4.3).

4.1 Stratified Jittered versus Uniform Jittered Sampling

Ray marching with uniform jittered sampling [Pauly et al. 2000] is a well-established way to attain unbiased estimates of optical depth [Kettunen et al. 2021]. It depends on a single uniform random number $\xi \in [0, 1]$ and uses it to jitter $N \in \mathbb{N}$ equidistant samples. Sample $j \in \{0, \dots, N-1\}$ is placed at $t_j := \frac{j+\xi}{N} t_{\max}$. Stratified jittered sampling similarly ensures exactly one sample per stratum $[\frac{j}{N}, \frac{j+1}{N})$ but places each sample independently. Sample j is placed at $t_j := \frac{j+\xi_j}{N} t_{\max}$ where $\xi_0, \dots, \xi_{N-1} \in [0, 1]$ are uniform and independent. In both cases, the Monte Carlo estimate of the optical depth is $X = \frac{t_{\max}}{N} \sum_{j=0}^{N-1} \mu(t_j)$.

In Fig. 4, we show histograms of optical depth estimates obtained using many evaluations of both of these techniques. It clearly shows that the distribution of stratified jittered sampling quickly approaches a normal distribution. With ten samples only minor deviations from a normal distribution remain. On the other hand, the distribution of the estimates from uniform jittered sampling does not approach a normal distribution at all.

This behavior can be explained by consulting the central limit theorem (CLT) [Knight 1999]. The most common formulations of the CLT deal with a sum of i.i.d. random variables. That does not

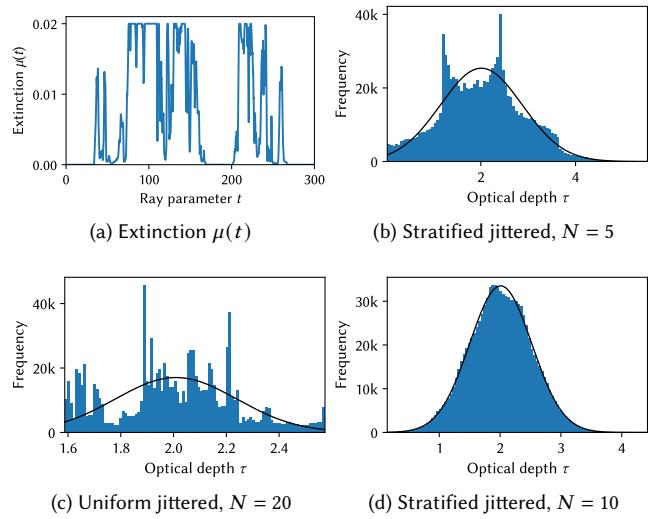


Fig. 4. Histograms of 10^6 estimates of optical depth for the shown extinction profile $\mu(t)$. With stratified jittered sampling, the distribution already approaches a normal distribution for $N = 5$ samples and is close to normal at $N = 10$. The black line shows a fitted normal density. Uniform jittered sampling shows no sign of approaching a normal distribution at $N = 20$.

match stratified jittered sampling where each term of the sum has a different distribution describing a different stratum. However, the Lyapunov CLT [Knight 1999, p. 145] is applicable. In Appendix B, we make the following claim precise and prove it: As long as we keep encountering samples with non-zero variance, the distribution of the optical depth estimate converges to a normal distribution as we continue marching along the ray.

Since our jackknife transmittance estimate depends on normal-distributed inputs to be unbiased, it is clear now that we should prefer stratified jittered sampling over uniform jittered sampling. Nonetheless, bias is not the only concern. We also want low variance. Fig. 5 compares convergence rates of three variants of stratified sampling as well as i.i.d. sampling, which chooses each $t_j \in [0, t_{\max}]$ uniformly. Like uniform jittered sampling, stratified jittered sampling converges faster than i.i.d. sampling at a rate close to $\frac{1}{N}$. For sample counts between 1 and 100, which is the regime we care about based on Fig. 4, the two techniques perform similarly. Once ray marching steps reach the scale of individual voxels, the techniques with equidistant samples begin to outpace stratified jittered sampling. That is probably because uniform jittered sampling inherits a $\frac{1}{N^2}$ convergence rate from the midpoint rule in this case. We provide graphs for other extinction profiles $\mu(t)$ as supplemental, which confirm these conclusions.

4.2 Variance-Aware Importance Sampling

Our reasoning using the CLT has a caveat: If a ray mostly travels through homogeneous regions of the volume, many samples will not contribute any variance to the estimate of optical depth. Usually, that would be desirable. In our case, however, it means that we wind up with an estimate of optical depth that is farther from being

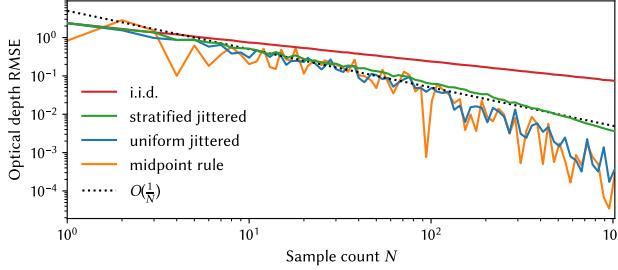


Fig. 5. A log-log-plot illustrating the convergence rate of various optical depth estimators for the extinction $\mu(t)$ shown in Fig. 4a. We compute the RMSE for each sample count N using 10^4 random realizations of each optical depth estimate. The deterministic midpoint rule is an exception. The stratified techniques are closer to linear convergence than to the $\frac{1}{\sqrt{N}}$ convergence of uniform i.i.d. samples in $[0, t_{\max}]$. Between 1 and 100 samples, stratified and uniform jittered sampling perform similarly. We provide examples with other extinction profiles $\mu(t)$ as supplemental material.

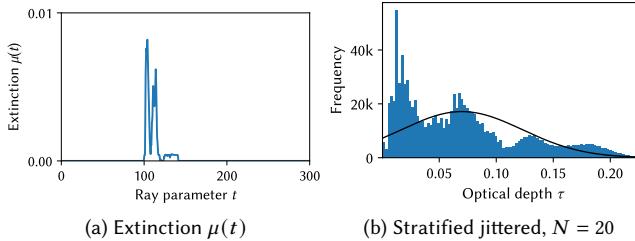


Fig. 6. A histogram like the ones in Fig. 4 but for a sparse extinction profile $\mu(t)$. In spite of using 20 samples, the distribution is still far from normal.

normal-distributed and thus it drives up bias in the transmittance estimate. Fig. 6 demonstrates that this is not a purely theoretical concern.

To remedy this problem, we use importance sampling, similar to Kettunen et al. [2021]. That is also an effective way to reduce variance. We compute a low-resolution super-voxel grid, where each super-voxel stores an importance for $R = r^3$ voxels (we use 16^3). Then the sampling density within each super-voxel is made to be proportional to its importance (Sec. 4.3).

Let $\mu_0, \dots, \mu_{R-1} \geq 0$ be the extinction values of all the voxels that fall into the super-voxel in question. We have flattened the 3D index into a 1D index here for notational convenience. Kettunen et al. [2021] use the mean as importance, i.e. $\frac{1}{R} \sum_{j=0}^{R-1} \mu_j$. Although importance sampling proportional to the integrand is perfect, this strategy is known to give a suboptimal piecewise constant density. The optimal piecewise constant density within each super-voxel is proportional to the root mean square extinction along the ray segment within the super-voxel [Pantaleoni and Heitz 2017]. Since it is not practical to compute this optimal density for every ray, we approximate it by $\sqrt{\frac{1}{R} \sum_{j=0}^{R-1} \mu_j^2}$, similar to the mean used by Kettunen et al. [2021].

One problem remains: The volume may contain large homogeneous regions with a non-zero extinction. Ray marching samples in

these regions do not contribute to the variance either, so our importance sampling scheme should avoid them. To this end, we compute the minimal extinction $\mu_{\min} := \min\{\mu_j\}_{j=0}^{R-1}$ for each super-voxel. Then our importance sampling and the integration only account for the difference $\mu_j - \mu_{\min}$. The optical depth of the low-resolution minimum volume is computed exactly using regular tracking and added to the end result. This is a typical application of control variates in volume rendering [Novák et al. 2014]. Our final super-voxel importance is

$$P := \sqrt{\frac{1}{R} \sum_{j=0}^{R-1} (\mu_j - \mu_{\min})^2}. \quad (5)$$

In principle, we can also use control variates that exceed the minimum value, but we do not recommend it. With such control variates, the noisy estimate of optical depth \bar{X} can become negative. Most of the time this is not a problem, but in rare cases the exponential growth of $\exp(-\bar{X})$ will make an outlier sample much worse. Thus, control variates in excess of the minimum introduce fireflies and harm robustness. Without them, our transmittance estimates are guaranteed to remain in the interval $[-1, 1]$ (Eq. 4).

4.3 Implementation

Designing a good implementation of our optical depth estimator is slightly challenging. It has to support our importance sampling scheme in conjunction with stratified jittered sampling, whilst also producing $m = 2$ optical depth estimates at the same time. Alg. 1 presents pseudo code for our implementation. Until it reaches the end of the ray, each pass through the outer loop takes exactly one sample $\mu(t)$ of the high-resolution volume. On SIMD architectures, that ensures coherent execution of this expensive step. Prior to each sample, the inner loop advances traversal of the super-voxel grid until the cumulative distribution function (CDF) F exceeds the CDF value for the next sample $F_{\Delta I_j}$. Note that the used CDF is an unnormalized integral over super-voxel importance values.

An input to Alg. 1 is the CDF step size F_{Δ} , i.e. the integral of importance within each stratum. Since our importance values in Eq. 5 are essentially extinctions, this value has the same unit as optical depth. It is meaningful in an absolute and scene-independent sense and can be set manually. However, that has two drawbacks: First, a ray traveling through dense and high-frequency regions of the volume will take an unnecessarily high number of samples, especially when the final transmittance is close to zero. Second, rays that only traverse thin parts in the outer rim of a volume do not take enough samples to arrive at a normal-distributed estimate.

To address these problems, we advocate for a fixed sample count N : We run regular tracking through the super-voxels twice. The first traversal computes the unnormalized CDF F as shown in Alg. 2. Then we set the CDF step size to $F_{\Delta} := \frac{F}{N}$. This is also a good opportunity to determine a tighter ray segment $[t'_{\min}, t'_{\max}]$ for use in Alg. 1 since we enforce a small non-zero importance for all non-empty super-voxels. If $F = 0$, optical depth estimation returns 0. With this strategy, the cost of the transmittance estimator is entirely predetermined by the length of the ray segment and our evaluation

Algorithm 1 Our optical depth estimator.

Input: Ray $\mathbf{x}(0), \omega, t_{\max}$ and CDF step size F_Δ .
Output: Two i.i.d. optical depth estimates X_0, X_1 .

```

 $F := 0, \tau_c := 0$            ▷ Current CDF and control variate
 $X_0 := 0, X_1 := 0$            ▷ Optical depth estimates
 $I_0 := \text{rand}(), I_1 := \text{rand}()$  ▷ CDFs at next sample, divided by  $F_\Delta$ 
 $j := \arg \min_{k \in \{0,1\}} I_k$     ▷ Which estimate is next?
Initialize DDA super-voxel traversal for the ray  $\mathbf{x}(0) + t\omega$ .
 $t_n := 0, t_{n+1} := 0, \mu_{\min} := 0, P := 0$ 
while  $F_\Delta I_j < F$  or  $t_{n+1} < t_{\max}$ :    ▷ Loop over high-res. samples
  while  $F_\Delta I_j \geq F$  and  $t_{n+1} < t_{\max}$ : ▷ Loop over super-voxels
    DDA step to next super-voxel with ray segment  $[t_n, t_{n+1}] \subseteq [0, t_{\max}]$ , importance  $P$  and minimum extinction  $\mu_{\min}$ .
     $F := F + (t_{n+1} - t_n)P$            ▷ Update the CDF
     $\tau_c := \tau_c + (t_{n+1} - t_n)\mu_{\min}$  ▷ Update the control variate
    if  $F_\Delta I_j < F$ :                  ▷ If we are not at  $t_{\max}$  yet:
       $t := t_{n+1} - \frac{F - F_\Delta I_j}{P}$     ▷ Rewind to the sample location
       $X_j := X_j + \frac{\mu(t) - \mu_{\min}}{P}$     ▷ Accumulate the sample
       $I_j := \text{ceil}(I_j) + \text{rand}()$     ▷ Advance CDF for next sample
       $j := \arg \min_{k \in \{0,1\}} I_k$            ▷ Which estimate is next?
     $X_0 := F_\Delta X_0 + \tau_c, X_1 := F_\Delta X_1 + \tau_c$  ▷ Complete the estimates
  Return  $(X_0, X_1)$ .

```

Algorithm 2 Preparations for optical depth estimation.

Input: Ray $\mathbf{x}(0), \omega, t_{\max}$.
Output: The integral of importance F (such that $F_\Delta = \frac{F}{N}$) and a tighter ray segment $[t'_{\min}, t'_{\max}]$.

$F := 0, t'_{\min} := t_{\max}, t'_{\max} := 0$
Initialize DDA super-voxel traversal for the ray $\mathbf{x}(0) + t\omega$.
For each super-voxel, covering a ray segment $[t_n, t_{n+1}] \subseteq [0, t_{\max}]$:
 Read the super-voxel importance P .
 $F := F + (t_{n+1} - t_n)P$
 if $P \neq 0$: $t'_{\min} := \min\{t_n, t'_{\min}\}, t'_{\max} := \max\{t_{n+1}, t'_{\max}\}$
Return F, t'_{\min}, t'_{\max} .

in Sec. 6.1 shows that it usually gives less bias. The cost of low-resolution regular tracking is significant though. We could store importance values in Alg. 2, but on GPU we prefer to reread them.

Our implementation stores high-resolution extinction values as 16-bit floats. For all low-resolution grids, we use BC4 compression [Hofmann and Evans 2021] but found that it makes a relatively minor difference. During compression, we take special care of proper rounding: Densities are always rounded up, because rounding to zero causes bias. Minima are always rounded down to avoid negative estimates of optical depth.

Rather than using trilinear interpolation for the high-resolution volume, we jitter each sample location stochastically [Hofmann et al. 2021; Pharr et al. 2024]. When that is enabled, we use overlapping super-voxels. To compute the jitters, we use PCG3D. For stratified jittered sampling, we choose PCG as inexpensive random number

generator with reasonable quality [Jarzynski and Olano 2020] (`rand()` in Alg. 1).

5 Our MIS Weight Estimator

At this point, we have established that the assumption of normal-distributed optical depth estimates is a good approximation. A natural next question is whether there are more ways to utilize this insight and the work of Gray et al. [1973]. And indeed, it turns out that we can not only estimate $\exp(-\tau)$ but any analytic function of optical depth (Sec. 5.1). As an important application, we demonstrate how to estimate MIS weights for distance sampling (Sec. 5.2).

5.1 Jackknife Function Estimation

We now introduce the UMVU estimate of $g(\tau) \in \mathbb{R}$. We demand that g is an analytic function, i.e. a function that can be written as a power series $g(\tau) = \sum_{j=0}^{\infty} a_j \tau^j$ with coefficients $a_0, a_1, \dots \in \mathbb{R}$. The coefficients can be obtained as derivatives using a Taylor series: $a_j := \frac{g^{(j)}(0)}{j!}$ for all $j \in \mathbb{N}_0$. Then the UMVU estimate of $g(\tau)$ is known [Gray et al. 1973, Theorem 4]:

Theorem 2. Let $m, X_0, \dots, X_{m-1}, \bar{X}$ and S be defined as in Thm. 1. Let $g(\tau) = \sum_{j=0}^{\infty} a_j \tau^j \in \mathbb{R}$ be analytic on the real axis \mathbb{R} . Let

$$G := g(\bar{X}) + \sum_{j=1}^{\infty} \frac{\Gamma(\frac{m-1}{2})}{\Gamma(\frac{m-1}{2} + j)} \frac{(-1)^j g^{(2j)}(\bar{X})}{j!} \left(\frac{S^2}{4}\right)^j. \quad (6)$$

Then under appropriate regularity conditions, G is the UMVU estimator of $g(\tau)$.

The aforementioned “appropriate regularity conditions” are stated implicitly in the proof of this theorem [Gray et al. 1973]. Alternatively, Gray et al. recommend to verify directly that the estimator is unbiased, for any specific function of interest $g(\tau)$.

It is not immediately clear how we should evaluate the series in Eq. 6 algorithmically. However, we found that this problem goes away in the special case $m = 2$. In Appendix C, we use basic algebraic manipulations to show that for $m = 2$

$$G = \sum_{j=0}^{\infty} \frac{(-1)^j g^{(2j)}(\bar{X})}{(2j)!} S^{2j}.$$

We now combine the sample mean and standard deviation into a single complex number $\bar{X} + iS \in \mathbb{C}$, where $i := \sqrt{-1}$ is the imaginary unit. Next, we study the real part of $g(\bar{X} + iS)$ using a Taylor series expansion around \bar{X} :

$$\begin{aligned} \Re g(\bar{X} + iS) &= \Re \sum_{j=0}^{\infty} \frac{g^{(j)}(\bar{X})}{j!} (iS)^j \stackrel{(*)}{=} \sum_{j=0}^{\infty} \frac{g^{(2j)}(\bar{X})}{(2j)!} (iS)^{2j} \\ &= \sum_{j=0}^{\infty} \frac{(-1)^j g^{(2j)}(\bar{X})}{(2j)!} S^{2j} = G. \end{aligned} \quad (7)$$

This turns out to be the sought-after UMVU estimate G for $m = 2$! In the step marked $(*)$, we have exploited that $\Re(i^j) = 0$ for odd j .

This result is rather confounding. We have interpreted the sample standard deviation S of our optical depth estimate as imaginary component of a complex optical depth estimate. Then simply feeding this complex optical depth into g and taking the real part gives us

the desired UMVU estimate. Although we have a proof that this is indeed correct, we do not have a simple intuitive argument why. However, we can at least convince ourselves that it is consistent with our formula for the jackknife transmittance estimate:

$$\Re \exp(-(\bar{X} + iS)) = \Re(\exp(-iS)) \exp(-\bar{X}) = \cos(S) \exp(-\bar{X}).$$

5.2 Jackknife MIS Weight Estimation

We will now apply this kind of function estimation to compute MIS weights [Veach and Guibas 1995]. As explained in Sec. 2.3, MIS still gives unbiased results when the MIS weights themselves are biased, but prior work has dismissed this approach [Miller et al. 2019; Novák et al. 2018].

To keep our discussion simple, we focus on the MIS combination of two distance sampling strategies: One strategy (e.g. equiangular sampling [Kulla and Fajardo 2012]) uses a known density $p_0(t)$, the other is free-flight distance sampling with density $p_1(t) = T(t)\mu(t)$. Then the MIS weight for strategy $j \in \{0, 1\}$ with the balance heuristic is

$$w_j(t) := \frac{p_j(t)}{p_0(t) + p_1(t)} = \begin{cases} 1 - \frac{T(t)\mu(t)}{p_0(t) + T(t)\mu(t)} & \text{if } j = 0, \\ \frac{T(t)\mu(t)}{p_0(t) + T(t)\mu(t)} & \text{if } j = 1. \end{cases} \quad (8)$$

It is also useful to study the MIS estimate more holistically. Let t_0, t_1 denote the distances sampled from $p_0(t), p_1(t)$, respectively. Let $L_{s,0}, L_{s,1}$ denote the Monte Carlo estimates of in-scattered radiance for these distance samples. Then the MIS estimate for the incoming radiance described in Eq. 2 is

$$L_i(\mathbf{x}(0), \omega) \approx L_{i,\text{MIS}} := \sum_{j=0}^1 w_j(t_j) \frac{T(t_j)\mu(t_j)L_{s,j}}{p_j(t_j)}.$$

Expanding this slightly and canceling $p_1(t_1)$ with $T(t_1)\mu(t_1)$ in the second term, we arrive at

$$L_{i,\text{MIS}} = w_0(t_0) \frac{T(t_0)\mu(t_0)L_{s,0}}{p_0(t_0)} + w_1(t_1)L_{s,1}. \quad (9)$$

Alternatively, we can merge MIS weight estimation and transmittance estimation:

$$L_{i,\text{MIS}} = \sum_{j=0}^1 \frac{T(t_j)\mu(t_j)}{p_0(t_j) + T(t_j)\mu(t_j)} L_{s,j}. \quad (10)$$

No matter whether we use Eq. 8 or Eq. 10, we have to estimate

$$\frac{T(t)\mu(t)}{p_0(t) + T(t)\mu(t)} = \frac{u_0 \exp(-\tau(t)) + u_1}{v_0 \exp(-\tau(t)) + v_1} =: g(\tau(t)), \quad (11)$$

where $u_0 := v_0 := \mu(t)$, $u_1 := 0$ and $v_1 := p_0(t)$ do not depend on $\tau(t)$. Applying Eq. 7, we get the biased estimate

$$g(\tau) \approx \Re \frac{u_0 \exp(-(\bar{X} + iS)) + u_1}{v_0 \exp(-(\bar{X} + iS)) + v_1}. \quad (12)$$

In this estimate, \bar{X}, S are computed from two unbiased estimates of optical depth X_0, X_1 exactly as in Eq. 4. Appendix F provides formulas that avoid complex arithmetic.

To evaluate our unbiased MIS estimate of $L_i(\mathbf{x}(0), \omega)$, we produce distance samples t_0, t_1 , invoke Algs. 2 and 1 and Eq. 12 twice to estimate $w_0(t_0), w_1(t_1)$ and finally use an unbiased transmittance estimator with independent random numbers for $T(t_0)$. Then using these estimates in Eq. 9 gives an unbiased estimate (see Appendix D).

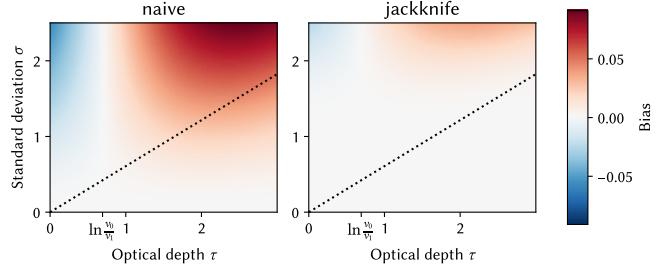


Fig. 7. We estimate $g(\tau)$ as defined in Eq. 11 for $(u_0, u_1, v_0, v_1) = (1, 0, 1, 1/2)$ and analyze the bias. For this experiment, the optical depth estimates are normal-distributed with varying mean τ and standard deviation σ . The dashed line marks $\sigma = 1.645\tau$, where the normal distribution has 5% of its mass across negative values. Our jackknife MIS weight estimate has low bias below this line and generally much less bias than the naive estimate $g(\bar{X})$. Both estimates have little bias for $\tau = \ln \frac{v_0}{v_1}$, which aligns with the location of poles of $g(\tau)$ (Appendix E). Our supplemental material provides results for other choices of u_0, u_1, v_0, v_1 , which support these conclusions.

Compared to the method of Miller et al. [2019], this procedure has more steps: Miller et al. reuse intermediate results from delta tracking and only need one invocation of ratio tracking for $T(t_0)$ on top of that. However, their MIS weights are defined differently. In our method, we can avoid the cost of estimating $T(t_0)$ separately, by using Eq. 10 instead of Eq. 9, but that introduces bias. Either way, we can also use the naive estimate of MIS weights $g(\bar{X})$ instead of our jackknife estimate. The unbiased estimate of radiance remains unbiased then.

It should be noted that $g(\tau)$ is not actually analytic, due to the division. Appendix E explains that the convergence radius of the power series $g(\tau) = \sum_{j=0}^{\infty} a_j \tau^j$ is $|\ln \frac{v_0}{v_1} + i\pi|$. Since analytic functions must have infinite convergence radius, Thm. 2 does not apply. Again, this is not a purely theoretical problem: Fig. 7 shows that our jackknife MIS weight estimator is indeed biased, even when the optical depth estimates are exactly normal-distributed. However, this bias is much weaker than for the naive estimate $g(\bar{X})$. In addition, the bias is small for combinations of optical depth τ and standard deviation σ that we care about: Our optical depth estimates are non-negative by design, but for $\sigma \geq 1.645\tau$ the normal distribution has more than 5% of its mass across negative values. In this case, our optical depth estimates cannot be close to being normal-distributed. Thus, the relevant parts of Fig. 7 are the parts below the dashed line, where our estimator has negligible bias.

Our derivation above only covers the simplest case, namely the balance heuristic for two strategies with one sample each. However, the procedure is easy to generalize, e.g. to the power heuristic: Derive MIS weights as function of optical depth $g(\tau)$ and evaluate $\Re g(\bar{X} + iS)$.

6 Results

We now evaluate our techniques in comparison to related work. We start with transmittance estimation (Sec. 6.1), proceed to MIS (Sec. 6.2) and finally discuss shared limitations (Sec. 6.3). To this end,

we use four volumes: Bunny cloud ($576 \times 571 \times 437$), Intel cloud ($625 \times 349 \times 566$), Disney cloud ($993 \times 675 \times 1224$) and explosion ($200 \times 271 \times 229$). Our supplemental material provides an interactive viewer with full sets of results for seven volumes.

Our implementation runs in a Vulkan fragment shader. All techniques use BC4-compressed super-voxel grids with one super-voxel per 16^3 voxels. For unbiased ray marching [Kettunen et al. 2021], we limit the maximal degree of the power series to eight, to avoid the need for dynamic memory allocation within a shader, we do not use endpoint matching and we use a variant of Alg. 1 for mean-based importance sampling, including tight ray segments from Alg. 2. All other parameters are chosen as proposed by Kettunen et al. [2021], although we sometimes double the sample count for higher quality. All transmittance estimators use stochastic texture filtering instead of trilinear interpolation. The reported timings refer to frames of resolution 1920×1080 rendered on an NVIDIA RTX 5070 Ti with GPU and memory clocks locked to 2452 MHz and 13801 MHz, respectively.

6.1 Transmittance Estimation

To evaluate transmittance estimators, we compute, display and analyze the transmittance for primary rays (with the exception of Figs. 1 and 12). That is not how they are typically used in a renderer (Sec. 2) but makes it easier to assess the bias and variance.

Biased Techniques. In Fig. 8, we compare different biased techniques: The biased ray marching technique of Kettunen et al. [2021] with doubled sample count and super-voxel size (compared to their evaluation) and three variants of our method. The supplemental additionally includes results for ray marching with equidistant steps throughout the volume bounding box, but techniques with importance sampling consistently perform better.

Among these techniques, our jackknife transmittance estimator with $N = 10$ samples is a clear winner (Fig. 8e). For three of the four volumes, it has the lowest bias. As expected (Sec. 3.3), the naive estimate with the same overall sample count (Fig. 8c) has slightly lower standard deviation but roughly three times more bias. Tying the ray marching sample count to the overall importance along the ray instead of keeping it fixed (Fig. 8d) is only beneficial for the explosion. On all scenes, biased ray marching (Fig. 8b) has greater bias than our method with fixed sample count. The frame times of all techniques are quite similar. The insets focus on challenging regions, but Fig. 10a provides full-size bias images for our technique.

The explosion is a failure case of most ray marching techniques, with clearly visible bias due to poor importance sampling. We discuss it in more detail under limitations (Sec. 6.3).

Variance-Aware Importance Sampling. As an ablation, Fig. 9 compares our variance-aware importance sampling with the minimum as control variate (Sec. 4.2) to importance sampling proportional to the super-voxel mean without a control variate. We observe a clear reduction of bias in all cases and a slight reduction of the standard deviation, which confirms that our approach is beneficial.

Unbiased Techniques. Fig. 10 compares our technique to unbiased transmittance estimators using the same scenes as in Fig. 8. In this

comparison, the timings and standard deviations have greater variation. Timings of our method (Fig. 10f) and the track-length estimate are similar and much lower than for ratio tracking and unbiased ray marching. However, the track-length estimate (Fig. 10b) has the worst possible standard deviation for a transmittance estimate that produces estimates in $[0, 1]$: It is $\sqrt{(1 - T)T}$. Our transmittance estimator achieves much lower standard deviation.

Ratio tracking still has much greater standard deviation and its frame times are 2 to 3 times greater than for our method (Fig. 10c). The standard deviation of unbiased ray marching [Kettunen et al. 2021] depends heavily on the quality of the importance sampling (Fig. 10d). In this regard, the explosion is a serious failure case (more so than for our method) with extreme fireflies. Doubling the sample count (Fig. 10e) mitigates this problem, but makes the frame times be inferior to ratio tracking.

Sample Count. Fig. 11 compares various techniques in terms of the number of samples in the high-resolution volume. Samples in the BC4-compressed super-voxel grid are not counted, but are predictable for all methods except the track-length estimate: Per traversed super-voxel, ratio tracking queries one value, unbiased ray marching two values and ours three values. Note that these reads are extremely cache coherent: Even for the Disney cloud, a single super-voxel grid takes only 102 KiB of memory.

The number of samples for the track-length estimate is extremely low. Though the timings in Fig. 10a hardly reflect that, presumably due to thread divergence. Ratio tracking and unbiased ray marching take most samples in regions of low transmittance, where they might not contribute meaningfully to the end result. Our method is configured to take exactly $mN = 20$ samples when it encounters a non-empty super-voxel and zero otherwise.

6.2 Multiple Importance Sampling

As explained in Sec. 2.3, there is little related work on MIS for distance samplers. To the best of our knowledge, the naive estimate

$$g(\bar{X}) = \frac{\mu(t) \exp(-\bar{X})}{p_0(t) + \mu(t) \exp(-\bar{X})}$$

was mentioned [Novák et al. 2018] but never systematically evaluated in a peer-reviewed publication. We do so, using our optical depth estimates. One can argue about whether or not that constitutes a contribution, but our jackknife MIS weight estimate $\mathfrak{Ag}(\bar{X} + iS)$ certainly does.

Other than that, the MIS of Miller et al. [2019] is the most important related work. Miller et al. describe the measurement contribution function and sampling density for complete transport paths with null-scattering vertices. In our experiments, we focus on direct illumination from point lights with equiangular sampling [Kulla and Fajardo 2012] as sampling strategy $p_0(t)$. We found that for MIS weights with the balance heuristic, many factors cancel. What remains of their MIS weight estimate is

$$w_1(t) \approx \frac{\mu(t) T_{rt}(t)}{p_0(t) + \mu(t) T_{rt}(t)}, \quad \text{where } T_{rt}(t) := \prod_{j=1}^{k-1} 1 - \frac{\mu(t'_j)}{\hat{\mu}(t'_j)}$$

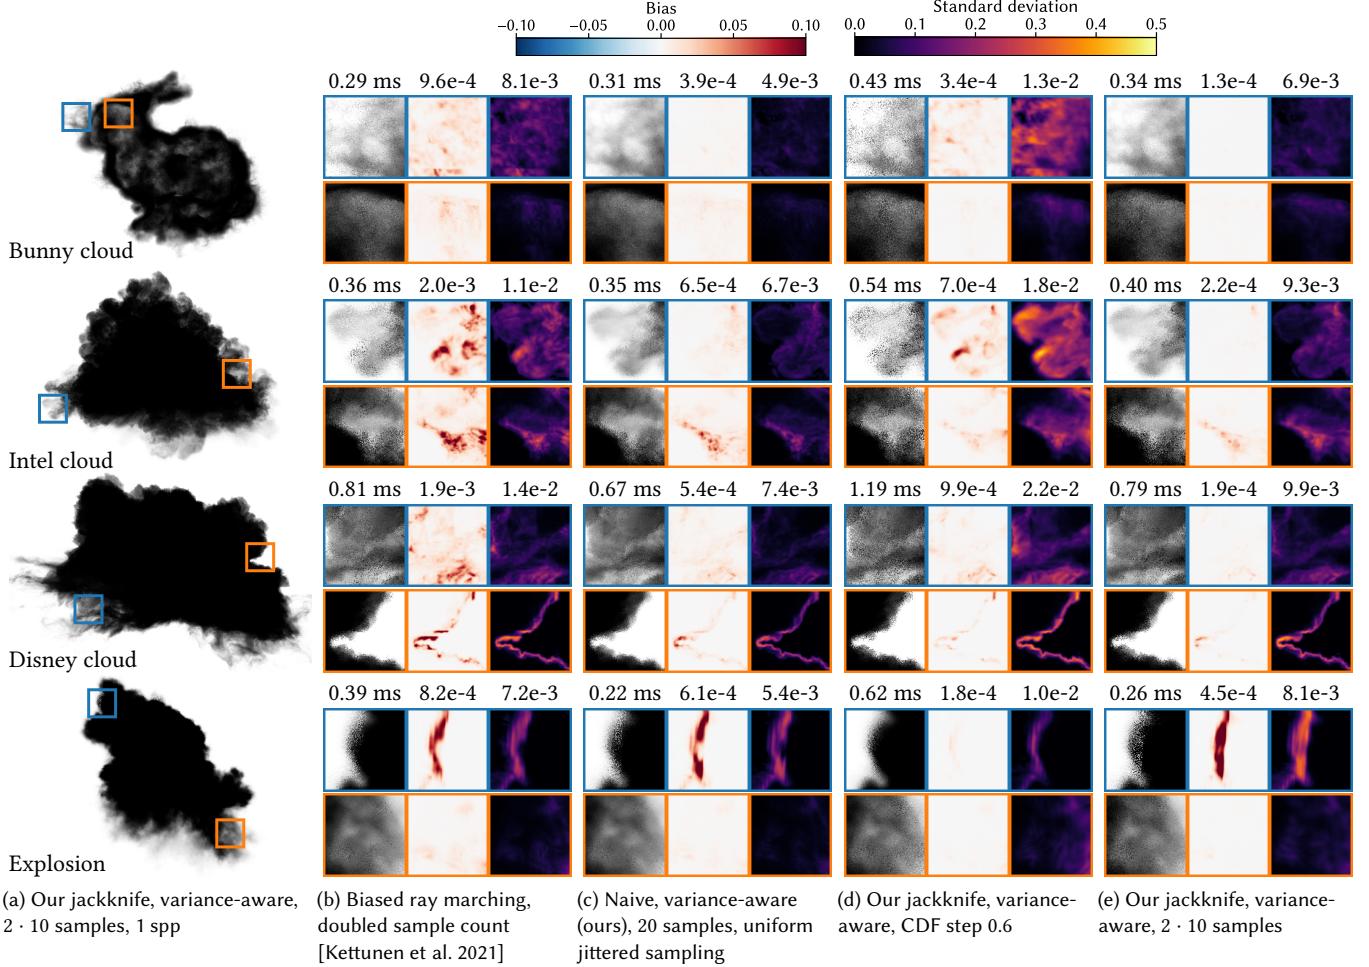


Fig. 8. Results of biased transmittance estimators. We show 1 sample per pixel (spp) full-size images for our main technique (a) and magnified insets for other techniques (b-d). Additional insets show the bias and standard deviation (computed from 2^{18} spp). For each full-size image, we report the total frame time at 1 spp as well as the average bias and standard deviation across all pixels. Our main technique achieves remarkably low bias and a good standard deviation (e).

is the transmittance estimate of ratio tracking (Sec. 2.2). For free-flight sampling, $T_{\text{rt}}(t_1)$ reuses the null-scattering events of delta tracking.

Fig. 12 shows our results. Since it is not obvious whether the MIS of Miller et al. [2019] is compatible with stochastic texture filtering, this figure uses nearest-neighbor interpolation. In this test case, we observe typical shortcomings of free-flight sampling and equiangular sampling: Free-flight sampling undersamples bright regions close to the point lights and deep inside the volume (green inset, Fig. 12a). Equiangular sampling undersamples the front of the volume (red inset, Fig. 12b). The improvement over pure free-flight sampling with the MIS of Miller et al. is relatively small (Fig. 12c). It inherits variance from ratio tracking and the Disney cloud is a worst case in this regard: Its interior has a constant extinction, where ratio tracking is no better than the track-length estimate.

The remaining three techniques, which all use our optical-depth estimator with $N = 10$, all succeed in combining the strengths of

free-flight and equiangular sampling. Their differences are largest in the purple inset, but even there they are hardly visible. The high quality of the naive estimate $g(\bar{X})$ with $m = 1$ (Fig. 12d) is surprising, given the dismissive statements about this approach in prior work [Miller et al. 2019; Novák et al. 2018] (Sec. 2.3). Apparently, our optical depth estimates make this approach viable. Fig. 7 clearly shows that our jackknife MIS weight estimator (Fig. 12f) produces considerably less bias. Though, bias in MIS weights does not cause bias in radiance estimates (Appendix D).

The biased version of our MIS produces nearly identical results (Fig. 12e), so it is justifiable as a way to avoid the cost of an additional transmittance estimation, especially when the transmittance estimate would be biased anyway. Fig. 13 shows that the combination of biased MIS (Eq. 10) with the naive estimate $g(\bar{X})$ gives visible bias, while our jackknife MIS weight estimation does not. The timings for our MIS variants are comparatively high, but clearly justified by the significant reduction in variance (Fig. 12).

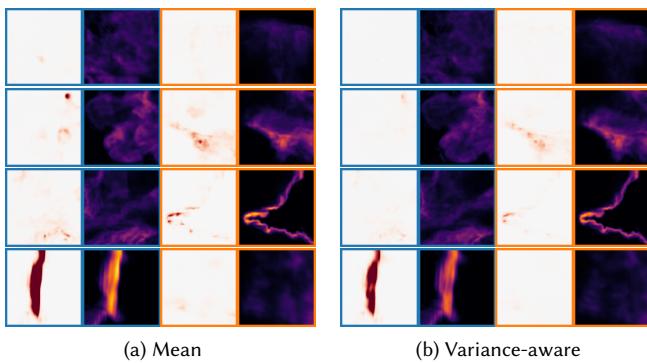


Fig. 9. An ablation study comparing our variance-aware importance sampling with the minimum as control variate (Sec. 4.2) to importance sampling proportional to the super-voxel mean. Both results use $2 \cdot 10$ samples for our jackknife transmittance estimate. The insets in (b) are the same as in Fig. 8e and the same colorbars apply. Variance-aware sampling reduces the bias considerably and the standard deviation slightly.

6.3 Limitations

The accuracy of our method is wholly dependent on the quality of the importance sampling for the optical depth estimates. As illustrated in Fig. 6, a sparse extinction signal may result in optical depth estimates that are far from being normal-distributed. Most of the time, our variance-aware importance sampling (Sec. 4.2) mitigates this problem, but sometimes it fails. The blue inset of the explosion (Fig. 10a) shows such a failure case and Fig. 14 provides further analysis: Since the ray traverses non-empty super-voxels without actually interacting with the volume, too few samples remain for the parts of the ray where the extinction is non-zero. For $N = 10$ samples, we can say with certainty that eight of them will make zero contribution and the ninth will make a small contribution most of the time. In this case, the assumption of normal-distributed optical depth estimates fails and the result is clear bias.

It should be noted that our method still gives results in $[-1, 1]$ with moderate standard deviation, whereas unbiased ray marching produces extremely high variance with transmittance estimates far outside the range $[-1, 1]$ (Fig. 10d). Null-scattering methods can handle this situation well, by simply taking smaller steps, which is reflected by their longer timings for the explosion in comparison to our method (Figs. 10b, 10c and 10f). The most pragmatic mitigation for the problem is to use smaller super-voxels. With our chosen super-voxel size of 16^3 , the super-voxel grid for the (fairly low-resolution) explosion contains only $13 \times 17 \times 15$ super-voxels. At a super-voxel size of 12^3 , the problems are mostly gone. For use cases in production rendering with much larger volumes, we anticipate that larger ray marching sample counts N would be used and that also mitigates these problems: Every good sample brings the estimate closer to a normal distribution. In general, the choice of the sample count N after the first pass through the volume (Alg. 2) is not set in stone and can be changed depending on the requirements of a renderer.

7 Conclusions

The bias of our biased transmittance estimator is so low that it is a serious contender for unbiased techniques. At the same time, its variance and run time cost are low and predictable. Failure cases arise when the importance sampling works poorly. Future work could address this problem using more sophisticated data structures, e.g. k -d trees with adaptive subdivision [Yue et al. 2011] or hierarchies of sparse grids [Museth 2021]. Our MIS weight estimation also works as desired and supports unbiased rendering.

All of this is accomplished using stunningly simple formulas, which date back five decades [Gray et al. 1973]. Beyond our specific applications, the concept of UMVU estimates holds great promise for Monte Carlo rendering and graphics in general: Our reasoning with the central limit theorem is quite broadly applicable and UMVU estimates are known for many statistics [Voinov and Nikulin 1993]. For example, we can use this method to estimate the reciprocal of a quantity for which an unbiased estimate is available. The estimate from Eq. 7 for $g(\tau) := \tau^{-1}$ is

$$\Re \left(\frac{X_0 + X_1}{2} + i \frac{X_0 - X_1}{2} \right)^{-1} = \frac{2(X_0 + X_1)}{(X_0 + X_1)^2 + (X_0 - X_1)^2} = \frac{X_0 + X_1}{X_0^2 + X_1^2}.$$

In the supplemental document, we show the bias and standard deviation of this estimate, which look promising. We would have shown the same figure for the naive estimate but neither the mean nor the standard deviation of \tilde{X}^{-1} converge [Robert 1991].

Acknowledgments

I would like to thank Sebastian Herholz, Tobias Zirr and Johannes Hanika for pointing out that MIS with biased MIS weights is unbiased, Sebastian Herholz, Weizhen Huang and all reviewers for their feedback on my manuscript and Mariia Soroka for encouraging submission to SIGGRAPH Asia instead of a smaller venue and for help in acquiring a rare technical report. I also thank Intel and Disney for their cloud datasets. The other volumes are OpenVDB samples. This publication is part of the project Aim@epilepsy with project number KIC1.ST03.21.016 of the research programme KIC - MISSIE 2021 which is (partly) financed by the Dutch Research Council (NWO).

References

Valeriya Akhmedova and Emil T. Akhmedov. 2019. *Selected Special Functions for Fundamental Physics*. doi:10.1007/978-3-030-35089-5

John Amanatides and Andrew Woo. 1987. A Fast Voxel Traversal Algorithm for Ray Tracing. In *EG 1987-Technical Papers*. Eurographics Association. doi:10.2312/egtp.19871000

Subrahmanyam Chandrasekhar. 1950. *Radiative Transfer*. Oxford University Press.

W. A. Coleman. 1968. Mathematical Verification of a Certain Monte Carlo Sampling Technique and Applications of the Technique to Radiation Transport Problems. *Nuclear Science and Engineering* 32, 1 (1968). doi:10.13182/NSE68-1

Iliyan Georgiev, Zackary Misso, Toshiya Hachisuka, Derek Nowrouzezahrai, Jaroslav Krivánek, and Wojciech Jarosz. 2019. Integral formulations of volumetric transmittance. *ACM Trans. Graph.* 38, 6 (2019). doi:10.1145/3355089.3356559

H. L. Gray, T. A. Watkins, and W. R. Schucany. 1973. On the jackknife statistic and its relation to UMVU estimators in the normal case. *Communications in Statistics* 2, 4 (1973). doi:10.1080/03610927308827077

Nikolai Hofmann and Alex Evans. 2021. *Ray Tracing Gems II*. Apress, Chapter Efficient Unbiased Volume Path Tracing on the GPU. doi:10.1007/978-1-4842-7185-8

Nikolai Hofmann, Jon Hasselgren, Petrik Clarberg, and Jacob Munkberg. 2021. Interactive Path Tracing and Reconstruction of Sparse Volumes. *Proc. ACM Comput. Graph. Interact. Tech.* 4, 1, Article 5 (2021). doi:10.1145/3451256

Nikolai Hofmann, Jon Hasselgren, and Jacob Munkberg. 2023. Joint Neural Denoising of Surfaces and Volumes. *Proc. ACM Comput. Graph. Interact. Tech.* 6, 1 (2023). doi:10.1145/3585497

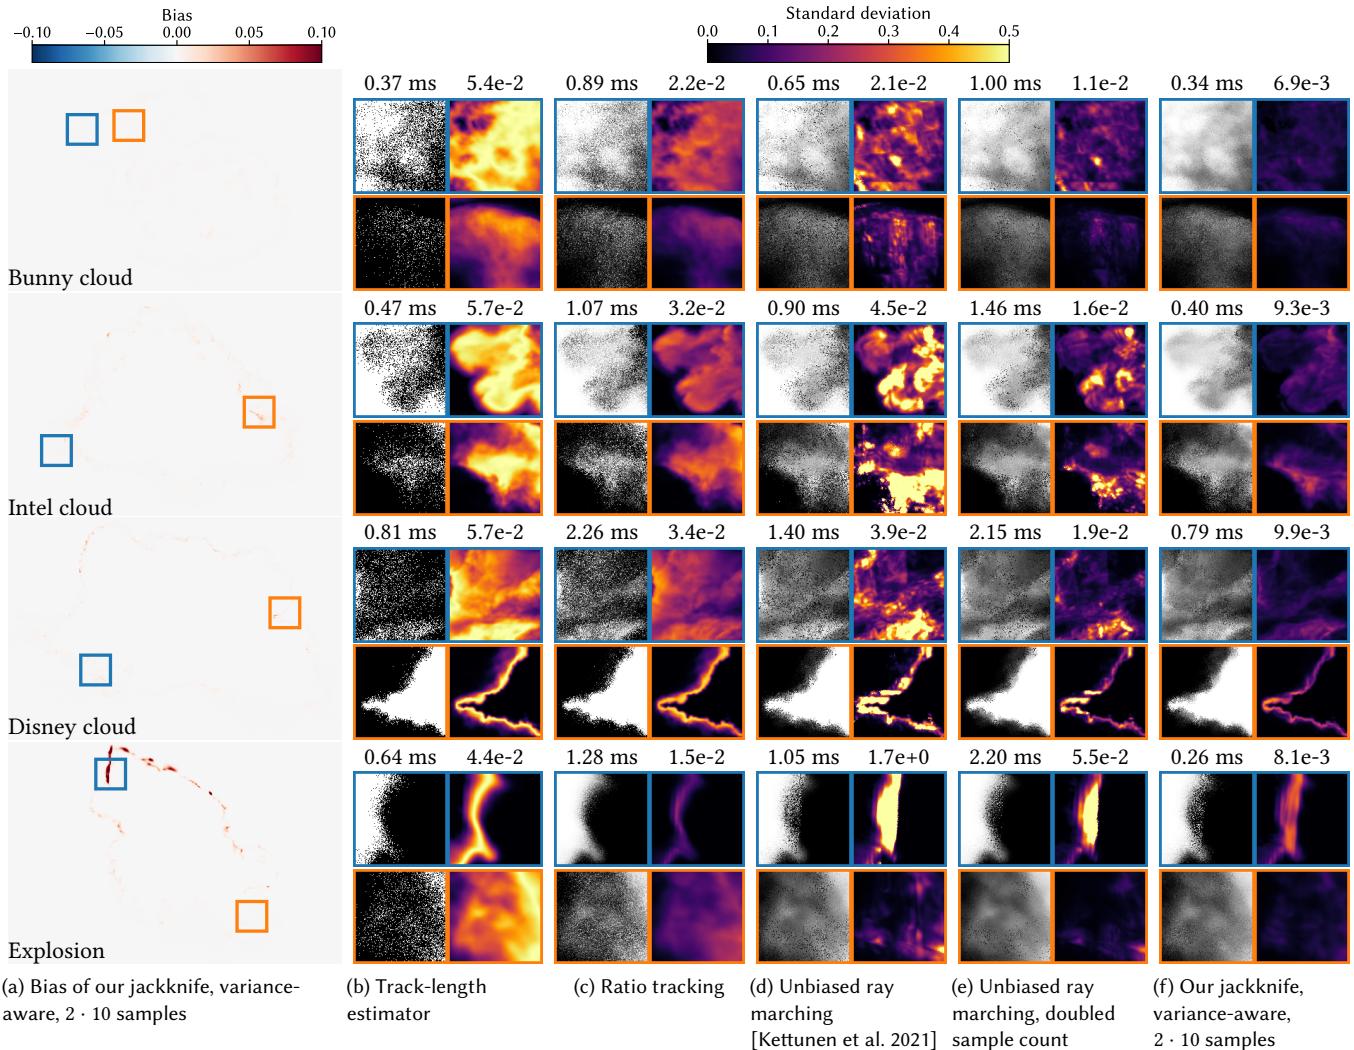


Fig. 10. A comparison of several unbiased transmittance estimators to our biased method. Per full-size image we report the frame time and average standard deviation. Our method is about as expensive as the track-length estimator, but with much lower standard deviation. Ratio tracking and unbiased ray marching fall somewhere in between in terms of standard deviation, whilst having much higher frame times.

Mark Jarzynski and Marc Olano. 2020. Hash Functions for GPU Rendering. *Journal of Computer Graphics Techniques* 9, 3 (2020). [https://www.jcgt.org/published/0009/03/02/](https://www.jcgt.org/published/0009/03/)

Johannes Jendersie and Eugene d'Eon. 2023. An Approximate Mie Scattering Function for Fog and Cloud Rendering. In *SIGGRAPH 2023 Talks*. ACM. doi:10.1145/3587421.3595409

Markus Kettunen, Eugene D'Eon, Jacopo Pantaleoni, and Jan Novák. 2021. An unbiased ray-marching transmittance estimator. *ACM Trans. Graph.* 40, 4 (2021). doi:10.1145/3450626.3459937

Keith Knight. 1999. *Mathematical Statistics*. Chapman & Hall/CRC. doi:10.1201/9780367805319

Christopher Kulla and Marcos Fajardo. 2012. Importance Sampling Techniques for Path Tracing in Participating Media. *Computer Graphics Forum* 31, 4 (2012). doi:10.1111/j.1467-8659.2012.03148.x

Bailey Miller, Iliyan Georgiev, and Wojciech Jarosz. 2019. A null-scattering path integral formulation of light transport. *ACM Trans. Graph.* 38, 4, Article 44 (2019). doi:10.1145/3306346.3323025

Zackary Missos, Benedikt Bitterli, Iliyan Georgiev, and Wojciech Jarosz. 2022. Unbiased and consistent rendering using biased estimators. *ACM Trans. Graph.* 41, 4, Article 48 (2022). doi:10.1145/3528223.3530160

Ken Museth. 2013. VDB: High-resolution sparse volumes with dynamic topology. *ACM Trans. Graph.* 32, 3 (2013). doi:10.1145/2487228.2487235

Ken Museth. 2021. NanoVDB: A GPU-Friendly and Portable VDB Data Structure For Real-Time Rendering And Simulation. In *ACM SIGGRAPH 2021 Talks*. doi:10.1145/3450623.3464653

Jan Novák, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz. 2018. Monte Carlo Methods for Volumetric Light Transport Simulation. *Computer Graphics Forum* 37, 2 (2018). doi:10.1111/cgf.13383

Jan Novák, Andrew Selle, and Wojciech Jarosz. 2014. Residual ratio tracking for estimating attenuation in participating media. *ACM Trans. Graph.* 33, 6 (2014). doi:10.1145/2661229.2661292

Jacopo Pantaleoni and Eric Heitz. 2017. Notes on optimal approximations for importance sampling. (2017). doi:10.48550/arxiv.1707.08358

Mark Pauly, Thomas Kollig, and Alexander Keller. 2000. Metropolis Light Transport for Participating Media. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000*. Springer, 11–22.

Matt Pharr, Bartłomiej Wrónski, Marco Salví, and Marcos Fajardo. 2024. Filtering After Shading With Stochastic Texture Filtering. *Proc. ACM Comput. Graph. Interact. Tech.* 7, 1 (2024). doi:10.1145/3651293

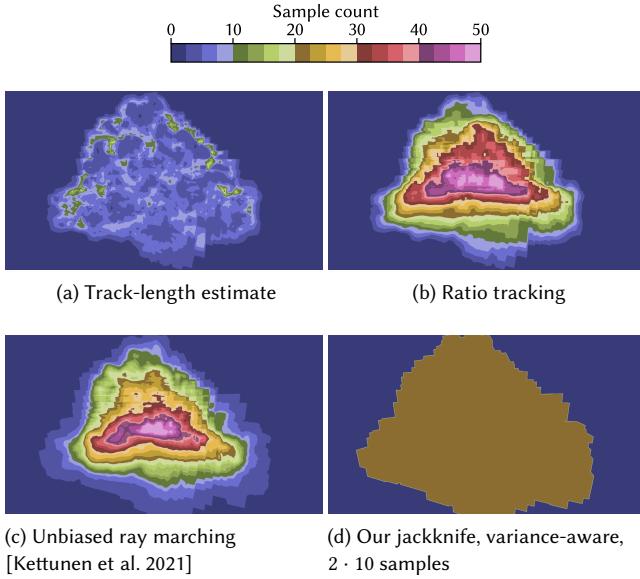


Fig. 11. A comparison of transmittance estimators in terms of the average number of high-resolution volume samples per pixel. Samples in the BC4-compressed super-voxel grid are not counted. The sample count of our method is fixed to 20 for all rays that hit non-empty super-voxels.

Christian Robert. 1991. Generalized inverse normal distributions. *Statistics & Probability Letters* 11, 1 (1991), 37–41. doi:10.1016/0167-7152(91)90174-P

Andrew Schneider. 2023. Nubis³: Methods (and madness) to model and render immersive real-time voxel-based clouds. In *ACM SIGGRAPH 2023 Courses*. Article 23. doi:10.1145/3587423.3607877 Advances in Real-Time Rendering (Part I).

László Szirmay-Kalos, Balázs Tóth, Milán Magdics, and Balázs Csébfalvi. 2010. Efficient Free Path Sampling in Inhomogeneous Media. In *Eurographics 2010 - Posters*, Anders Hast and Ivan Viola (Eds.). doi:10.2312/egp.20101020

Eric Veach and Leonidas J. Guibas. 1995. Optimally combining sampling techniques for Monte Carlo rendering. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*. doi:10.1145/218380.218498

Vassily G. Voinov and Mikhail S. Nikulin. 1993. *Unbiased Estimators and Their Applications. Volume 1: Univariate Case*. Kluwer Academic Publishers. doi:10.1007/978-94-011-1970-2

Adolf N. Witt. 1977. Multiple scattering in reflection nebulae. I. A Monte Carlo approach. *Astrophysical Journal Supplement Series* 35 (1977). doi:10.1086/190463

Yonghao Yue, Kei Iwasaki, Bing-Yu Chen, Yoshinori Dobashi, and Tomoyuki Nishita. 2011. Toward Optimal Space Partitioning for Unbiased, Adaptive Free Path Sampling of Inhomogeneous Participating Media. *Computer Graphics Forum* 30, 7 (2011). doi:10.1111/j.1467-8659.2011.02049.x

A Bias, Variance and RMSE

We now derive formulas for the bias, variance and RMSE of the naive transmittance estimator $\exp(-X)$ and our jackknife transmittance estimator under the assumption of normal-distributed estimates of optical depth. In doing so, we repeatedly make use of the moment-generating function of a normal-distributed random variable X with mean τ and standard deviation σ for an $s \in \mathbb{C}$ [Knight 1999, p. 44]:

$$E(\exp(sX)) = \exp\left(s\tau + \frac{s^2\sigma^2}{2}\right). \quad (13)$$

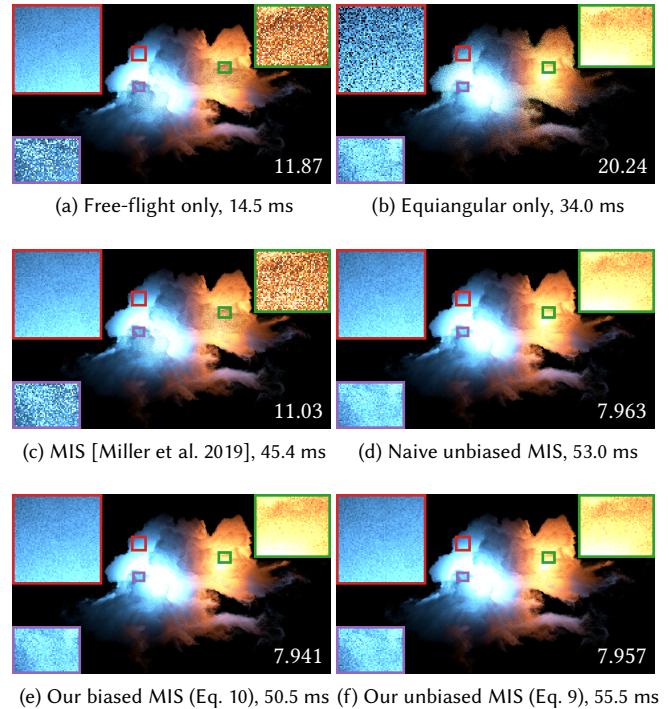


Fig. 12. The Disney cloud, directly lit by two point lights inside the cloud. All images use 16 spp, but for MIS that amounts to 32 distance samples. The timings are frame times for 16 spp. Additionally, we report the symmetric mean absolute percentage error (SMAPE) with $\epsilon := 0.01$ added to the mean linear RGB value in the denominator. Transmittance estimation for the ray to the chosen point light uses our jackknife transmittance estimator with $N = 10$ in all cases. The improvement of the MIS of Miller et al. [2019] over free-flight sampling is moderate. All other variants of MIS successfully combine the strengths of both strategies with minor differences.

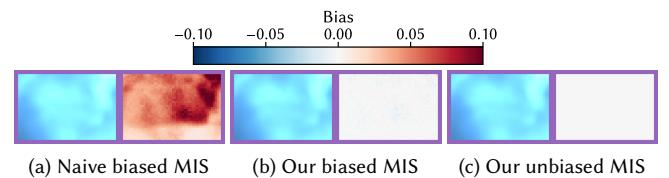


Fig. 13. We render the purple inset from Fig. 12 at 2^{18} spp using three variants of MIS. The images show the rendering itself and the bias in the luminance. The bias with our jackknife MIS weight estimation is not visible, whereas the result with the naive estimate is visibly too bright.

A.1 Naive Transmittance Estimation

The ground-truth transmittance is $T = \exp(-\tau)$. According to Eq. 13, the first two moments of the transmittance estimate $L := \exp(-X)$ for a normal-distributed optical depth estimate X are

$$E(L) = E(\exp((-1)X)) = \exp\left(-\tau + \frac{\sigma^2}{2}\right) = \exp\left(\frac{\sigma^2}{2}\right)T,$$

$$E(L^2) = E(\exp((-2)X)) = \exp(-2\tau + 2\sigma^2) = \exp(2\sigma^2)T^2.$$

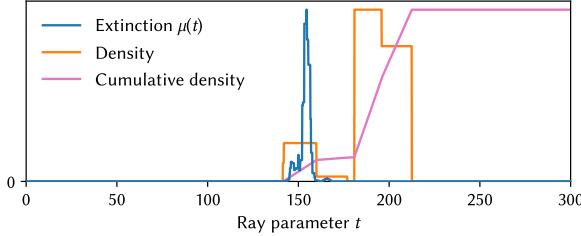


Fig. 14. An analysis of the source of the problems for a pixel in the middle of the blue inset of the explosion (Fig. 10a). Since the ray traverses several non-empty super-voxels without encountering parts of the volume, 86% of all samples contribute nothing to the estimate of optical depth. Note that this figure does not use stochastic trilinear interpolation or BC4 compression and that BC4 compression aggravates the problem further.

Thus, we get the

$$\begin{aligned} \text{bias: } E(L) - T &= \left(\exp\left(\frac{\sigma^2}{2}\right) - 1 \right) T, \\ \text{variance: } E(L^2) - (E(L))^2 &= (\exp(2\sigma^2) - \exp(\sigma^2))T^2, \\ \text{RMSE: } E((L - T)^2) &= E(L^2 - 2LT + T^2) \\ &= \left(\exp(2\sigma^2) - 2\exp\left(\frac{\sigma^2}{2}\right) + 1 \right) T^2. \end{aligned}$$

In Sec. 3.3, we use the mean \bar{X} of $m = 2$ estimates instead of a single sample X . That halves the variance σ^2 .

A.2 Jackknife Transmittance Estimation

Since the sign of S is irrelevant for the value of $\cos(S)$, we use $S := \frac{X_0 - \bar{X}_1}{2}$ for the following derivation. Since the matrix $\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ is orthogonal, the random variables S and \bar{X} are independent and normal-distributed [Knight 1999, p. 84]. Furthermore, \bar{X} has mean τ and variance $\frac{\sigma^2}{2}$ and S has mean 0 and variance $\frac{\sigma^2}{2}$. First, we prove unbiasedness using $\cos(S) = \Re \exp(iS)$ and Eq. 13:

$$\begin{aligned} E(K) &= \Re E(\exp(iS))E(\exp(-\bar{X})) \\ &= \Re \exp\left(\frac{i^2\sigma^2}{4}\right) \exp\left(-\tau + \frac{\sigma^2}{4}\right) = \exp(-\tau) = T. \end{aligned}$$

We already knew that this estimate is unbiased [Gray et al. 1973] (for normal-distributed optical depth estimates), but now we have a concise proof. Computing the second moment works similarly, but exploits $\cos^2(S) = 1/2(1 + \cos(2S))$:

$$\begin{aligned} E(K^2) &= E(\cos^2(S))E(\exp(-2\bar{X})) \\ &= 1/2E(1 + \Re \exp(2iS))\exp(-2\tau + \sigma^2) \\ &= 1/2(1 + \Re \exp(-\sigma^2))\exp(\sigma^2)T^2 \\ &= 1/2(\exp(\sigma^2) + 1)T^2. \end{aligned}$$

The variance is

$$V(K) = E(K^2) - T^2 = 1/2(\exp(\sigma^2) - 1)T^2.$$

Since the estimate is unbiased, the RMSE matches the standard deviation.

B Applying the Central Limit Theorem

The Lyapunov CLT is essential to our reasoning, since it explains why our estimates of optical depth with stratified jittered sampling are approximately normal-distributed. We will show that as a ray travels through a volume with enough variation, the sum in the ray marching estimate approaches a normal distribution. Arguably, it would be preferable to show that the optical depth estimate for a fixed ray segment approaches a normal distribution as the sample count N increases. That is however not viable, since each new sample count gives us a new set of strata with new distributions of extinction values. By considering the sum along a ray with a fixed step size, we avoid this issue. Our precise statement is as follows:

Proposition 3. Consider an extinction profile $0 \leq \mu(t) \leq \hat{\mu}$ with ray parameter $t \geq 0$. Let $\xi_0, \xi_1, \dots \in [0, 1]$ be uniform and i.i.d.. Let $t_\Delta > 0$ be the used ray marching step size such that sample $j \in \mathbb{N}_0$ is taken at $t_j := (j + \xi_j)t_\Delta$. Let $Y_j := \mu(t_j)$ denote the corresponding extinction values. The stratified jittered sampling estimate of $\tau(nt_\Delta)$ is $t_\Delta \sum_{j=0}^{n-1} Y_j$. Assume that

$$\lim_{n \rightarrow \infty} \frac{n^{\frac{2}{3}}}{\sum_{j=0}^{n-1} V(Y_j)} = 0, \quad (14)$$

i.e. the sum of variances grows faster than $n^{\frac{2}{3}}$. Then for $n \rightarrow \infty$, the following normalized version of the optical depth estimate converges in distribution to a standard normal distribution:

$$\frac{(t_\Delta \sum_{j=0}^{n-1} Y_j) - \int_0^{nt_\Delta} \mu(t) dt}{t_\Delta \sqrt{\sum_{j=0}^{n-1} V(Y_j)}}. \quad (15)$$

PROOF. We consider a scaled zero-mean version of Y_j , along with its variance and the third moment of its absolute value:

$$\begin{aligned} Z_j &:= t_\Delta Y_j - E(t_\Delta Y_j) = t_\Delta Y_j - \int_{jt_\Delta}^{(j+1)t_\Delta} \mu(t) dt, \\ \sigma_j^2 &:= V(Z_j) = t_\Delta^2 V(Y_j), \\ \gamma_j &:= E(|Z_j|^3) \leq E((t_\Delta \hat{\mu})^3) = (t_\Delta \hat{\mu})^3. \end{aligned}$$

Raising Eq. 14 to the $\frac{3}{2}$ -th power and multiplying by $(t_\Delta \hat{\mu})^3$ gives

$$\lim_{n \rightarrow \infty} \frac{\sum_{j=0}^{n-1} (t_\Delta \hat{\mu})^3}{\left(\sum_{j=0}^{n-1} V(Y_j) \right)^{\frac{3}{2}}} = 0 \quad \Rightarrow \quad \lim_{n \rightarrow \infty} \frac{\sum_{j=0}^{n-1} \gamma_j}{\left(\sum_{j=0}^{n-1} \sigma_j^2 \right)^{\frac{3}{2}}} = 0.$$

This is exactly the prerequisite to be able to apply the Lyapunov CLT [Knight 1999, p. 145], which then tells us that the following random variable converges in distribution to a standard normal distribution:

$$\frac{\sum_{j=0}^{n-1} Z_j}{\sqrt{\sum_{j=0}^{n-1} \sigma_j^2}} = \frac{(t_\Delta \sum_{j=0}^{n-1} Y_j) - \int_0^{nt_\Delta} \mu(t) dt}{t_\Delta \sqrt{\sum_{j=0}^{n-1} V(Y_j)}}.$$

□

Taken literally, this proof requires an infinite volume with never-ending variation. More pragmatically, it tells us that the ray marching estimate approaches a normal distribution with each new sample that has non-zero variance. Fig. 4 underpins this finite interpretation empirically and Sec. 4.2 ensures that most samples will have non-zero variance.

C Jackknife Function Estimation for Two Samples

Our goal is to prove that for $m = 2$

$$\frac{\Gamma\left(\frac{m-1}{2}\right)}{\Gamma\left(\frac{m-1}{2} + j\right)} \frac{(-1)^j g^{(2j)}(\bar{X})}{j!} \left(\frac{S^2}{4}\right)^j = \frac{(-1)^j g^{(2j)}(\bar{X})}{(2j)!} S^{2j}.$$

For all $j \in \mathbb{N}$, the gamma function satisfies [Akhmedova and Akhmedov 2019, pp. 7, 3]

$$2^{2j-1} \Gamma(j) \Gamma\left(\frac{1}{2} + j\right) = \sqrt{\pi} \Gamma(2j), \quad \Gamma(j) = (j-1)!.$$

Thus,

$$\begin{aligned} \frac{\Gamma\left(\frac{1}{2}\right)}{\Gamma\left(\frac{1}{2} + j\right)} &= \frac{2^{2j-1} \Gamma(j) \Gamma\left(\frac{1}{2}\right)}{2^{2j-1} \Gamma(j) \Gamma\left(\frac{1}{2} + j\right)} \\ &= \frac{2^{2j-1} (j-1)! \sqrt{\pi}}{\sqrt{\pi} \Gamma(2j)} = \frac{4^j (j-1)!}{2(2j-1)!}. \end{aligned}$$

Now

$$\begin{aligned} \frac{\Gamma\left(\frac{m-1}{2}\right)}{\Gamma\left(\frac{m-1}{2} + j\right)} \frac{(-1)^j g^{(2j)}(\bar{X})}{j!} \left(\frac{S^2}{4}\right)^j \\ = \frac{1}{2(2j-1)!} \frac{(-1)^j g^{(2j)}(\bar{X})}{j} S^{2j} = \frac{(-1)^j g^{(2j)}(\bar{X})}{(2j)!} S^{2j}. \end{aligned}$$

D Unbiasedness of MIS

As explained in Sec. 2.3, MIS is unbiased for any choice of MIS weights as long as $w_0(t) + w_1(t) = 1$ [Veach and Guibas 1995]. However, this result pertains to deterministic MIS weights, whereas our MIS weights are random variables. Additionally, we have canceled $p_1(t_1)$ with $T(t_1) \mu(t_1)$ in Eq. 9. We now prove that our variants of MIS based on Eq. 9 are unbiased nonetheless. We are interested in conditional expectations of random variables w.r.t. $t_j = t \in \mathbb{R}$ where $j \in \{0, 1\}$. We denote the random variable for the estimate of the MIS weight $w_j(t)$ by $\tilde{w}_j(t)$. Similarly, we denote the unbiased estimates of transmittance and inscattered radiance at t by $\tilde{T}(t)$ and $\tilde{L}_s(t)$, respectively. The estimates $\tilde{w}_j(t)$, $\tilde{T}(t)$, $\tilde{L}_s(t)$ all consume different, independent random numbers. Therefore, $E(\tilde{w}_j(t) \tilde{T}(t) \tilde{L}_s(t)) = E(\tilde{w}_j(t)) T(t) E(\tilde{L}_s(t))$. The values of $\mu(t)$, $p_0(t)$, $p_1(t)$ are deterministic. Then the expected value of Eq. 9 is:

$$\begin{aligned} E(L_{i,\text{MIS}}) &= E\left(\tilde{w}_0(t_0) \frac{\tilde{T}(t_0) \mu(t_0) \tilde{L}_s(t_0)}{p_0(t_0)} + \tilde{w}_1(t_1) \tilde{L}_s(t_1)\right) \\ &= \int_0^\infty E\left(\tilde{w}_0(t) \frac{\tilde{T}(t) \mu(t) \tilde{L}_s(t)}{p_0(t)}\right) p_0(t) + E(\tilde{w}_1(t) \tilde{L}_s(t)) p_1(t) dt \\ &= \int_0^\infty E(\tilde{w}_0(t)) T(t) \mu(t) E(\tilde{L}_s(t)) + E(\tilde{w}_1(t)) E(\tilde{L}_s(t)) T(t) \mu(t) dt \\ &= \int_0^\infty (E(\tilde{w}_0(t)) + E(\tilde{w}_1(t))) E(\tilde{L}_s(t)) T(t) \mu(t) dt \end{aligned}$$

Therefore, this radiance estimate will be unbiased if $E(\tilde{w}_0(t)) + E(\tilde{w}_1(t)) = 1$ for all $t \geq 0$. A trivial way to satisfy this requirement is to define $\tilde{w}_0(t) := 1 - \tilde{w}_1(t)$ as suggested by Eq. 8 (using different random numbers for evaluation). Though, applying our jackknife

estimate for $\tilde{w}_0(t)$, $\tilde{w}_1(t)$ more directly also works:

$$\begin{aligned} E(\tilde{w}_1(t)) &= E\left(\Re \frac{\mu(t) \exp(-(\bar{X} + iS))}{p_0(t) + \mu(t) \exp(-(\bar{X} + iS))}\right) \\ &= E\left(\Re \frac{p_0(t) + \mu(t) \exp(-(\bar{X} + iS)) - p_0(t)}{p_0(t) + \mu(t) \exp(-(\bar{X} + iS))}\right) \\ &= 1 - E\left(\Re \frac{p_0(t)}{p_0(t) + \mu(t) \exp(-(\bar{X} + iS))}\right) = 1 - E(\tilde{w}_0(t)) \end{aligned}$$

Replacing \bar{X} by X_0 and S by 0 in this derivation shows that the naive MIS weight estimate also satisfies this requirement for unbiasedness. Of course, the MIS weight estimates also contribute to the variance but our results in Sec. 6.2 validate our approach in this regard.

E Poles and Convergence Radius

The function $g(\tau) = \frac{u_0 \exp(-\tau) + u_1}{v_0 \exp(-\tau) + v_1}$ is not well-defined for all $\tau \in \mathbb{C}$. Since we know $v_0, v_1 \geq 0$, the denominator $v_0 \exp(-\tau) + v_1$ cannot vanish for real τ . However, there are poles in the complex plane:

$$\begin{aligned} v_0 \exp(-\tau) + v_1 &= 0 \\ \Leftrightarrow \exp(-\Re \tau) \exp(-i\Im \tau) &= \exp(-\tau) = -\frac{v_1}{v_0} \\ \Leftrightarrow \exp(-\Re \tau) &= \frac{v_1}{v_0} \quad \wedge \quad \exp(-i\Im \tau) = -1 \\ \Leftrightarrow \Re \tau &= -\ln \frac{v_1}{v_0} = \ln \frac{v_0}{v_1} \quad \wedge \quad \exists k \in \mathbb{Z} : \Im \tau = \pi + 2\pi k. \end{aligned}$$

The convergence radius of the power-series representation $g(\tau) = \sum_{j=0}^\infty a_j \tau^j$ will end exactly at one of these poles, i.e. it is $\left| \ln \frac{v_0}{v_1} + i\pi \right|$.

F Evaluating MIS Weights

Without using complex arithmetic explicitly, our goal is to evaluate

$$\Re \frac{U}{V} = \frac{\Re U \Re V + \Im U \Im V}{(\Re V)^2 + (\Im V)^2}, \quad \text{where} \quad \begin{aligned} U &:= u_0 \exp(-(\bar{X} + iS)) + u_1, \\ V &:= v_0 \exp(-(\bar{X} + iS)) + v_1, \end{aligned}$$

for real $u_0, u_1, v_0, v_1, \bar{X}, S$. We note

$$\begin{aligned} \Re \exp(-(\bar{X} + iS)) &= \Re \exp(-iS) \exp(-\bar{X}) = \cos(S) \exp(-\bar{X}), \\ \Im \exp(-(\bar{X} + iS)) &= \Im \exp(-iS) \exp(-\bar{X}) = \sin(-S) \exp(-\bar{X}). \end{aligned}$$

Thus,

$$\begin{aligned} \Re U &= u_0 \cos(S) \exp(-\bar{X}) + u_1, & \Im U &= u_0 \sin(-S) \exp(-\bar{X}), \\ \Re V &= v_0 \cos(S) \exp(-\bar{X}) + v_1, & \Im V &= v_0 \sin(-S) \exp(-\bar{X}). \end{aligned}$$