

Volume Path Guiding Based on Zero-Variance Random Walk Theory

SEBASTIAN HERHOLZ, University of Tübingen, Germany
YANGYANG ZHAO, McGill University, Montreal, Canada
OSKAR ELEK, Charles University, Prague, Czech Republic
DEREK NOWROUZEZHAI, McGill University, Montreal, Canada
HENDRIK P. A. LENSCH, University of Tübingen, Germany
JAROSLAV KŘIVÁNEK, Charles University, Prague, Czech Republic




Fig. 1. Path tracing of an optically dense medium (30 minutes), showing both a complete ‘beauty’ render ($S+V$), and a render with only volumetric transport (V). Building on an approximate adjoint solution of the incident and in-scattered radiance, our zero variance-based path construction forms near-optimal decisions for guided collision distance sampling, directional sampling, Russian roulette and path splitting. As such, our sampling methodology leads to significantly faster convergence compared to an unguided path tracer with standard transmittance-based collision and phase function sampling.

The efficiency of Monte Carlo methods, commonly used to render participating media, is directly linked to the manner in which random sampling decisions are made during path construction. Notably, path construction is influenced by scattering direction and distance sampling, Russian roulette, and splitting strategies. We present a consistent suite of volumetric path construction techniques where all these sampling decisions are guided by a cached estimate of the *adjoint transport solution*. The proposed strategy is based on the theory of zero-variance path sampling schemes, accounting for the spatial and directional variation in volumetric transport. Our key technical contribution, enabling the use of this approach in the context of volume light transport, is a novel guiding strategy for sampling the particle *collision distance* proportionally to the product of transmittance and the

adjoint transport solution (e. g., in-scattered radiance). Furthermore, scattering directions are likewise sampled according to the product of the phase function and the incident radiance estimate. Combined with guided Russian roulette and splitting strategies tailored to volumes, we demonstrate about an order-of-magnitude error reduction compared to standard unidirectional methods. Consequently, our approach can render scenes otherwise intractable for such methods, while still retaining their simplicity (compared to, e. g., bidirectional methods).

CCS Concepts: • **Computing methodologies** → **Ray tracing**;

Additional Key Words and Phrases: Monte Carlo, path guiding, zero variance path sampling theory, importance sampling, volume light transport, participating media rendering

ACM Reference Format:

Sebastian Herholz, Yangyang Zhao, Oskar Elek, Derek Nowrouzezahrai, Hendrik P. A. Lensch, and Jaroslav Krivánek. 2018. Volume Path Guiding Based on Zero-Variance Random Walk Theory. *ACM Trans. Graph.* 0, 0, Article 0 (2018), 19 pages. <https://doi.org/0>

1 INTRODUCTION

Recent industrial adoption of physically-based rendering methodologies has sparked renewed interest in Monte Carlo (MC) methods for light transport simulation. This wider adoption has also exposed many limitations inherent to MC methods, including their slow convergence in scenes with participating media [Fong et al. 2017].

Authors' addresses: Sebastian Herholz, University of Tübingen, Germany, sebastian.herholz@uni-tuebingen.de; Yangyang Zhao, McGill University, Montreal, Canada; Oskar Elek, Charles University, Prague, Czech Republic; Derek Nowrouzezahrai, McGill University, Montreal, Canada; Hendrik P. A. Lensch, University of Tübingen, Germany; Jaroslav Krivánek, Charles University, Prague, Czech Republic.

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 ACM 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.

© 2018 Association for Computing Machinery.
0730-0301/2018/0-ART0 \$15.00
<https://doi.org/0>

This can be attributed to a historical focus in computer graphics on optimizing MC methods for surface transport, and the added complexities of volumetric transport [Christensen and Jarosz 2016; Pharr et al. 2016].

Volumetric phenomena contribute significantly to realistic images. As such, recent work has focused on bridging the gap between surface- and volume-transport simulation [Jarosz 2013], including efficient importance sampling methods for unbiased [Georgiev et al. 2013; Kulla and Fajardo 2012; Kutz et al. 2017; Novák et al. 2014; Simon et al. 2017; Szirmay-Kalos et al. 2017] and biased [Křivánek et al. 2014; Novák et al. 2012b] rendering. The theoretical overlap between light and particle transport has also promoted knowledge transfer between the graphics and neutron transport communities [Dwivedi 1982; Galtier et al. 2013; Hoogenboom 2008; McCormick and Kuscer 1973; Spanier and Gelbard 1969]. Ultimately, however, these methods focus on either optimally sampling a *subset* of the transport equation terms, or they rely on restrictive assumptions about the underlying medium, such as homogeneity or scattering isotropy – that is, if they support volumetric media at all. Either way, important lighting features are left untreated during stochastic sampling, leading to estimators with high variance.

We leverage the key insight that, in order to improve convergence, MC volumetric light transport must optimally account for *every* stochastic decision: all spatial, directional and path-length sampling decisions should be treated in a consistent manner. Our main contribution is a *unified methodology* for close-to-optimal importance sampling of *volumetric transport paths*. Zero-variance path sampling theory [Hoogenboom 2008; Křivánek and d'Eon 2014; Meng et al. 2016] serves as the theoretical framework here, providing ground rules for building a *globally optimal*, joint path-space sampling distribution using but a set of *locally optimal* sampling decisions used during incremental path construction (Section 3.2). This process is referred to as *transport path guiding*.

Our contributions lie in the volume-specific aspects of path guiding, to complement the existing works with focus on surfaces (especially the ones compatible with the zero-variance sampling theory [Herholz et al. 2016; Vorba et al. 2014; Vorba and Křivánek 2016]). Specifically, we propose unbiased methods to efficiently and accurately conduct all the volumetric sampling decisions:

- **Our key technical contribution is a novel technique for the collision distance and scatter/no-scatter decisions performed in proportion to the full product of both transmittance and in-scattered radiance.** This proposed formulation avoids constructing the sampling distribution upfront and instead operates incrementally along the particle flight path, yielding an efficient solution, especially for optically thick media.
- **We sample scattering directions according to the product of the phase function and incident radiance,** each of which is represented using von Mises-Fisher (vMF) mixtures. Crucially, these allow for efficient evaluation, product and convolution operations. Previous path guiding work on surfaces is adapted for this purpose [Herholz et al. 2016; Vorba et al. 2014].

- Finally, **our path termination/splitting decisions locally adjust the density of traced paths according to their estimated contributions,** i. e., based on the expected throughput-weighted radiance. Again, we extend previous work focused on surface transport [Vorba and Křivánek 2016] to volumes.

Since a practical implementation requires knowledge of the *adjoint transport solution* (e. g., the steady-state spatio-directional radiance distribution), we cache *approximate* adjoints during a particle tracing pre-pass. Using parametric vMF mixture models for our cached representation enables closed-form evaluation of all required adjoint quantities. We then follow the framework of zero-variance path sampling to unify our sampling procedures, transferring and extending recent concepts on path guiding for surfaces [Dahm and Keller 2017; Herholz et al. 2016; Müller et al. 2017; Vorba et al. 2014; Vorba and Křivánek 2016] to full volumetric transport.

Compared to the state of the art, the proposed guided sampling techniques (Section 6) have two critical advantages. First, they sample proportionally to the *full integrand* of the volume transport equation, across respective local sub-domains (Section 3). Second, all local sampling decisions are consistently guided by the same underlying adjoint solution (Section 5). Consequently, our method can efficiently simulate a number of volumetric phenomena deemed difficult or costly for current rendering techniques, comprising both low-order (e. g., light shafts, volumetric caustics) and high-order (e. g., dense high-albedo media) transport effects (Section 7). In addition to our practical results, we include numerical analysis of our sampling techniques (in isolation and joint combinations), validating their variance-reduction ability across diverse scene configurations.

2 RELATED WORK

We focus on efficient sampling for Monte Carlo solutions to transport problems. Interested readers can refer to broader surveys of volumetric media rendering in research [Novák et al. 2018] and production [Fong et al. 2017].

Local importance sampling. The efficiency of Monte Carlo methods depends largely on their ability to generate light transport paths with a distribution resembling the underlying transport equation as closely as possible. Since this is a difficult proposition, in practice the sampled paths are instead drawn from one or more of the transport equation’s factors (Section 3.1). Specialized schemes have thus been designed to importance-sample them: Rayleigh [Frisvad 2011] and Henyey-Greenstein [Witt 1977] phase functions can be sampled analytically and adaptive sampling can be conducted according to coarse visibility oracles [Belcour et al. 2014; Engelhardt and Dachsbacher 2010]. In optically thin media dominated by low-order transport effects, convergence can be significantly improved by importance-sampling single scattering from point [Kulla and Fajardo 2012] or environmental sources [von Radziewsky et al. 2017].

Extensions to bidirectional path connection schemes are useful for generalizing these gains to scenes with refractive interfaces [Koerner et al. 2016] or two- and three-bounce transport [Georgiev et al. 2013]. While the latter work shares the same objective as our method – that of *joint importance sampling* of transport paths – it is limited to importance sampling of the geometry terms and the phase function over (up to) three random events, ignoring transmittance.

Our method, on the other hand, can build *full paths* by importance sampling *all* terms of the transport equation.

A significant effort has also been devoted to sampling proportionally to the transmittance and density of the medium. Here, tabulated transmittance sampling [Amanatides and Woo 1987] introduces bias, whereas its unbiased counterpart – delta tracking [Raab et al. 2006; Woodcock et al. 1965] – becomes inefficient for media with significant optical density variation. One approach to alleviate this issue is to bound regions with similar density [Szirmay-Kalos et al. 2011; Yue et al. 2010]. Independently, a *weighted* tracking scheme instead of an *analog* one has likewise proven effective [Galtier et al. 2013; Morgan and Kotlyar 2015]. These techniques reduce the number of required density look-ups, and can be combined with control variates to further reduce variance [Novák et al. 2014]. Several new unbiased tracking methods have then been developed based on these formulations [Kutz et al. 2017; Szirmay-Kalos et al. 2017].

Regardless, all the methods above only address individual components of the global transport sampling problem (Section 3.1). This results in unpredictable residual variance in general scenes, i. e., where transport characteristics are not known beforehand. Our main goal is, therefore, a unified volumetric path sampling approach that supersedes these more specific techniques.

Transport path guiding. Many efficient bidirectional rendering algorithms leverage the reciprocity of light transport [Christensen 2003; Krivánek et al. 2014; Lafortune and Willems 1995; Veach 1997]. Similarly, dual (adjoint) transport solutions generated in a pre-pass can be used to guide a unidirectional solver [Hey and Purgathofer 2002; Jensen 1995; Lafortune and Willems 1995]. In the latter context, both continuous [Herholz et al. 2016; Vorba et al. 2014] and discrete [Dahm and Keller 2017; Müller et al. 2017] representations of the adjoint transport solution have recently been explored. Building atop a unidirectional solution is desirable since a vast majority of rendering in practice relies on such approaches.

We are motivated by the fact that the above methods exclusively treat surface light transport, whereas path guiding in volumetric media has remained largely unexplored. Nevertheless our approach is strongly influenced by these works. Notable exceptions of volumetric guiding in graphics is the sequential MC method of Pegoraro et al. [2008] or the significance cache-based method by Bashford-Rogers et al. [2012], who perform directional sampling according to a histogram or cosine lobe representation of the approximated adjoint transport solution. Although being a step forward, these methods are limited to near-isotropic phase functions, while not considering the effect of transmittance or path length on the ideal sampling distribution. In contrast, our method performs sampling according to the *full volume rendering equation* (Section 3.1), without imposed restrictions on the form or properties of the medium.

Specialized methods. Volumetric analogues of caching methods based on photon density estimation [Bitterli and Jarosz 2017; Jarosz et al. 2011; Jensen and Christensen 1998; Novák et al. 2012a] or virtual light sources [Engelhardt et al. 2012; Novák et al. 2012b] often face similar sampling decisions as the purely path-based methods, and have consequently led to a suite of dedicated sampling techniques. These methods can be combined to improve the convergence of individual estimators [Krivánek et al. 2014]. Explicit cache

placement can also be guided by importance sampling according to illumination gradients [Jarosz et al. 2008; Ren et al. 2008].

Yet again, similarly to the local methods, these typically consider only the terms of the transport equation that can be readily sampled within the respective caching scheme. In contrast, we designed our caching to support all the necessary sampling decisions dictated by the theoretical foundations of volumetric transport (Section 3.2).

Zero-variance path sampling theory. Ultimately, we seek to obtain an optimal MC estimator in radiative transport. Such an estimator would sample according to all factors of the transport equation, since any ignored one results in variance increases. While that is the established idea of importance sampling, we can make an even stronger claim on the basis of the zero-variance random walk theory [Hoogenboom 2008; Kalos and Whitlock 2008]: that a perfect *global* MC sampling can be achieved if every *local* decision is made optimally, that is, in exact proportion to the respective integrand of the transport equation.

This proposition is, however, only hypothetical – infinite computation and memory is needed to perform this sampling, since the adjoint solution itself is nested in the transport equation. Exceptions comprise idealized scenarios with simplified closed-form adjoints, such as homogeneous infinite half-space transport [Dwivedi 1982; McCormick and Kuscer 1973]. Recent works in computer graphics have nevertheless leveraged these results to improve the convergence for sub-surface scattering simulation [Krivánek et al. 2014; Meng et al. 2016].

The conclusion that can be drawn from these works is that even an approximate adjoint solution can serve an MC approach that follows the zero-variance principles [Vorba and Krivánek 2016]. We build on this key observation and propose a unified sampling scheme with state-of-the-art convergence rates (Sections 4–6). Crucially, any residual variance of our method results only from imperfections in our approximate adjoint solution (as examined in Section 7).


3 BACKGROUND

We first briefly revisit the integral form of the volume rendering equation (VRE), which underlies the volumetric light transport problems. The MC estimator for the VRE is then described in Section 3.2, focusing on defining the target probability density functions (PDFs) for optimal (zero-variance) sampling.

3.1 Volume Rendering Equation

We rely on the standard parametrization for participating media [Pharr et al. 2016] via spherical phase function f , absorption coefficient σ_a , scattering coefficient σ_s and extinction coefficient (optical density) $\sigma_t = \sigma_a + \sigma_s$. All these quantities can be spatially varying. To simplify the exposition we will however not consider emissive media (see Simon et al. [2017] for a dedicated treatment of these).

The method presented in this paper provides a numerical solution to the integral form of the






Fig. 2. *Top*: An example path containing volume (blue) and surface (yellow) vertices. Note that, in the path tracing algorithm, the path carries visual importance, and thus is generated in the opposite direction of the flow of light. *Bottom*: The four considered zero-variance sampling decisions, mirroring the Equations 7–10. Optimal decisions are whether to scatter within or outside the medium, how far to travel until the next scattering event, how to choose the scattering direction and when to terminate the path.

volume rendering equation (VRE), expressing the (scalar, monochromatic) radiance L arriving at point \mathbf{x} from direction ω :

$$L(\mathbf{x}, \omega) = \underbrace{T(\mathbf{x}, \mathbf{x}_s)L_o(\mathbf{x}_s, \omega)}_{L_s(\mathbf{x}, \omega)} + \underbrace{\int_0^s T(\mathbf{x}, \mathbf{x}_d)\sigma_s(\mathbf{x}_d)L_i(\mathbf{x}_d, \omega) dd}_{L_m(\mathbf{x}, \omega)}, \quad (1)$$

where $\mathbf{x}_s = \mathbf{x} - s\omega$ and $\mathbf{x}_d = \mathbf{x} - d\omega$. In the presence of participating media the evaluation is typically split into two parts: the attenuated surface contribution L_s arriving from the closest surface intersection at \mathbf{x}_s , and the volume contribution L_m , scattered along and into ω within the medium. The outgoing surface radiance L_o is the sum of the emitted radiance L_e and the reflected radiance L_r [Kajiya 1986]. The transmittance function T defines the loss of energy due to photon-medium collisions (i. e., volumetric absorption and out-scattering, expressed by the optical thickness τ) along an optical path between two points:

$$T(\mathbf{x}_1, \mathbf{x}_2) = e^{-\tau(\mathbf{x}_1, \mathbf{x}_2)}, \quad \tau(\mathbf{x}_1, \mathbf{x}_2) = \int_0^{\|\mathbf{x}_1 - \mathbf{x}_2\|} \sigma_t(\mathbf{x}_d) dd. \quad (2)$$

Finally, the medium radiance contribution L_m gathers the incoming radiance L_i across a sphere \mathcal{S} :

$$L_i(\mathbf{x}_d, \omega) = \int_{\mathcal{S}} f(\mathbf{x}_d, \omega, \omega')L(\mathbf{x}_d, \omega') d\omega', \quad (3)$$

which is weighted by σ_s in Equation 1 and attenuated by the transmittance along the optical path to \mathbf{x} .

The VRE (Equation 1) has a structure similar to the standard rendering equation [Kajiya 1986]; in fact it reduces to it if $\sigma_a = \sigma_s = 0$. It however bears additional significant complexity due to the integration within the medium. In a stochastic MC solver such as path tracing, this leads to additional linear, spherical, and discrete sampling decisions covered in the following section.

3.2 Zero-variance Path Sampling Theory

We solve the VRE using MC estimation with optimized importance sampling according to the theory of zero-variance path sampling.

Monte-Carlo estimator. Solving the VRE (Equation 1) using MC estimation leads to a procedure where a set of N visual importance particles are traced until being stochastically terminated. Along its

trajectory, each particle collects radiance at emitting positions, or by doing next event estimation [Pharr et al. 2016], yielding one or several *light transport paths*. Since Equation 1 assumes non-emissive media, emission can only occur on a light source surface. To simplify the exposition a step further, we focus on the so called ‘last-event’ estimator [Hoogenboom 2008]: the particle only collects radiance at an emitting position and is then immediately terminated.

The contribution of each particle’s path r_i is given by the cumulative particle weight a and the emitted radiance L_e at the end of the path:

$$I = \mathbb{E} \left[\frac{1}{N} \sum_i^N a(r_i) \cdot L_e(\mathbf{x}_{i,M}, \omega_{i,M-1}) \right]. \quad (4)$$

Along its path each of the N particles undergoes $M - 1$ scattering events (on surface or in volume, see Figure 2). At each event and its coupled propagation step the particle performs a state change, i. e., its position and orientation changes from (\mathbf{x}_j, ω_j) to $(\mathbf{x}_{j+1}, \omega_{j+1})$. The positions \mathbf{x}_0 and \mathbf{x}_M define the start and the end of the path. The full particle weight is then the product of the weights for each individual state change j :

$$a(r) = \prod_{j=1}^{M-1} a_j(r). \quad (5)$$

The individual state change weight is given by the throughput of the path segment divided by the PDF of performing the corresponding state change. If for the sake of brevity we only consider volumetric interactions, the weight of a segment is defined as:

$$a_j(r) = \frac{1}{q_{j-1}} \frac{T(\mathbf{x}_{j-1}, \mathbf{x}_j)\sigma_s(\mathbf{x}_j)f(\mathbf{x}_j, \omega_{j-1}, \omega_j)}{p(\mathbf{x}_j, \omega_j|\mathbf{x}_{j-1}, \omega_{j-1})}, \quad (6)$$

with q being the probability of surviving the termination at the respective event and p being the PDF for generating the next path segment, both defined below.

Zero-variance path sampling strategy. The zero-variance path sampling theory [Hoogenboom 2008] provides a set of particle state change rules, which can be used for guided sampling of incrementally constructed paths. Applying these rules provably results in all particles having the exact same contribution (cf. Equation 4), resulting in an estimator with zero variance. A remarkable property

of such schemes is that by following the prescribed set of *local* sampling decisions, a *globally optimal* sampling strategy emerges. This is enabled by the fact that each of these local decisions relies on (is ‘guided’ by) information about the global transport solution.

In contrast to surface-based transport – where the generation of a next particle state only relies on directional sampling – the presence of participating media introduces a set of additional sampling decisions (illustrated in Figure 2). At this point we define the PDFs for these four decisions derived from Equation 1 according to the zero-variance theory:

1. Given a particle state (\mathbf{x}_j, ω_j) , the first step of the state change is to determine the new particle position \mathbf{x}_{j+1} . This involves two sub-steps:
 - a. First, a discrete decision determines whether the path explores the contribution L_s from the nearest surface, or the volumetric contribution L_m from the medium (see Equation 1). The optimal probability for sampling L_m between \mathbf{x}_j and the next surface is

$$P_m(\mathbf{x}_j, \omega_j) = \frac{L_m(\mathbf{x}_j, \omega_j)}{L_s(\mathbf{x}_j, \omega_j) + L_m(\mathbf{x}_j, \omega_j)}. \quad (7)$$

- b. Next, assuming that the volume contribution was selected in Step 1a, the collision distance d_{j+1} to the next scattering position at \mathbf{x}_{j+1} has to be sampled. The optimal PDF for sampling d_{j+1} starting at (\mathbf{x}_j, ω_j) is

$$p_d(d_{j+1} | \mathbf{x}_j, \omega_j) = \frac{T(\mathbf{x}_j, \mathbf{x}_{j+1})\sigma_s(\mathbf{x}_{j+1})L_i(\mathbf{x}_{j+1}, \omega_j)}{L_m(\mathbf{x}_j, \omega_j)}. \quad (8)$$

2. At the new scattering position \mathbf{x}_{j+1} the random walk needs to explore the integral of in-scattered radiance L_i , i. e., scattered into the direction ω_j . Therefore, a new direction ω_{j+1} needs to be sampled, with the optimal PDF

$$p_\omega(\omega_{j+1} | \mathbf{x}_{j+1}, \omega_j) = \frac{f(\mathbf{x}_{j+1}, \omega_j, \omega_{j+1})L(\mathbf{x}_{j+1}, \omega_{j+1})}{L_i(\mathbf{x}_{j+1}, \omega_j)}. \quad (9)$$

3. Obtaining L_i itself involves an evaluation of the incoming radiance L (Equation 3). Steps 1–2 are thus repeated, until we reach a position \mathbf{x}_j which emits energy in the direction ω_{j-1} and the decision is made to terminate the random walk and return its weighted emission $a(r) \cdot L_e$. The optimal termination probability at \mathbf{x}_j is

$$P_{RR}(\mathbf{x}_j, \omega_{j-1}) = \frac{L_e(\mathbf{x}_j, \omega_{j-1})}{L_e(\mathbf{x}_j, \omega_{j-1}) + L_r(\mathbf{x}_j, \omega_{j-1})} = \frac{L_e(\mathbf{x}_j, \omega_{j-1})}{L_o(\mathbf{x}_j, \omega_{j-1})}, \quad (10)$$

with the resulting survival probability of

$$q_j = 1 - P_{RR}(\mathbf{x}_j, \omega_{j-1}). \quad (11)$$

In the case there is no emission at \mathbf{x}_j the zero-variance probability of continuing the walk is one.

Putting everything together, the resulting PDF for a state change (Equation 6) is composed of the conditional PDFs for each of these




Fig. 3. *Left:* Visualization of the (kd) tree for the spatial cache component. *Right:* A schematic and false-color depiction of the directional caches for the incident and in-scattered radiance represented via vMF mixtures (orange). The sample volume has a slightly forward-peaked phase function.

sampling decisions:

$$p(\mathbf{x}_{j+1}, \omega_{j+1} | \mathbf{x}_j, \omega_j) = P_m(\mathbf{x}_j, \omega_j) \cdot p_d(d_{j+1} | \mathbf{x}_j, \omega_j) \cdot p_\omega(\omega_{j+1} | \mathbf{x}_{j+1}, \omega_j). \quad (12)$$

In the following sections, we present our core contribution: techniques for efficiently performing the guided sampling decisions outlined in Steps 1 through 3, on the basis of the zero-variance framework. Sections 4–6 describe the algorithmic design, followed by practical considerations and evaluation in Section 7.

4 METHODOLOGY

We build on the core principle that a globally optimal sampling of the transport path space can be achieved by making the ideal sampling decision for every individual integrand of the governing transport equation (Section 3.1). Section 3.2 thus provides a consistent set of optimal Monte Carlo path sampling strategies unified under the zero-variance theoretical framework.

To that end, our method needs to rely on estimating the adjoint transport solution that *guides the stochastic path construction* process, since all the partial PDFs (Equations 7–10) require the knowledge of the global energy distribution. As described in Section 5, we store this estimate in a hybrid cache, using a kd-tree to represent the spatial component of the adjoint solution, and von-Mises-Fisher model mixtures to represent its directional component.

Building on the adjoint solution estimate we describe the individual sampling procedures in Sections 6.1–6.3, reflecting the theoretical subdivision into partial PDFs from Steps 1–3 in Section 3.2. Incrementally applying these sampling decisions leads to a guided construction of a near-optimal transport path, subject to the representation accuracy of the cached adjoint solution (Section 5). Note that we only address the sampling of the volumetric light transport, since sampling on surfaces has been addressed previously (see Section 2). On the other hand, the distinct advantage of our approach is its zero-variance-theoretical foundation, which makes it readily compatible with other surface guiding approaches built on this theoretical framework (see Section 8).

5 ADJOINT SOLUTION ESTIMATE

As stated, zero-variance sampling requires prior knowledge of the light transport solution. Given this cyclic dependency, our method resorts to approximations, specifically of the incident radiance L and in-scattered radiance L_i (Equations 1 and 3). These quantities are stored in a hybrid spatial and directional cache structure (see Figure 3), similar to recent methods in surface path guiding [Bus and Boubekeur 2017; Müller et al. 2017; Vorba et al. 2014]. We initialize our prior estimates in a short pre-processing phase, where photons are traced from the light sources and then clustered using a spatial subdivision (kD) tree. Photons in each leaf node are used to build the directional representation of the incoming radiance for the volume contained within that node.

We represent directional distributions using a parametric mixture model based on the von-Mises-Fisher distribution (vMF) widely used in directional statistics. This has important advantages:

- similarly to the main quantities of the volumetric transport the distributions are naturally defined on the sphere and therefore no additional mapping is required, cf. [Herholz et al. 2016; Müller et al. 2017; Vorba et al. 2014];
- the in-scattered radiance estimate is calculated through an efficient *closed-form convolution*;
- the product distribution between the incident radiance and the phase function for *any directional configuration* is also obtained in a closed-form.

To fit the vMF distributions within maximum a-posteriori expectation maximization (MAP-EM) algorithm, which is similar to the method presented by Vorba et al. [2014]. Since we use vMFs, as opposed to their Gaussians, our maximum-likelihood parameter estimation and the associated priors are different (as described in Appendix A.1). For more details on MAP-EM for vMF mixtures please refer to Bangert et al. [2010].

Incoming radiance estimate. To represent the incoming radiance distributions at different locations \mathbf{x} inside a medium, we use vMF mixtures \mathcal{V} described by the spatially-varying parameter set $\Theta(\mathbf{x})$. Such a mixture represents a spherical distribution via a set of K weighted vMF lobes v :

$$\mathcal{V}(\omega | \Theta) = \sum_{i=1}^K \pi_i \cdot v(\omega | \mu_i, \kappa_i). \quad (13)$$

The parameter set Θ contains the weights $\{\pi_1, \dots, \pi_K\}$, the mean directions $\{\mu_1, \dots, \mu_K\}$ and precisions $\{\kappa_1, \dots, \kappa_K\}$ for each component (‘lobe’). Each vMF lobe is a normalized spherical distribution, rotationally symmetric around its mean, with its spread being inversely proportional to κ . Appendix A provides additional details.

The incident radiance estimate \tilde{L} is related to the normalized vMF mixture through scaling by the scalar irradiance (i. e., fluence):

$$\tilde{L}(\mathbf{x}, \omega) = \Phi(\mathbf{x}) \cdot \mathcal{V}_L(\omega | \Theta_L(\mathbf{x})). \quad (14)$$

Fluence is given by $\Phi(\mathbf{x}) = \int_S L(\mathbf{x}, \omega') d\omega'$. It is determined using a standard photon map estimate, during the fitting process of the vMF mixture for the incoming radiance, and stored along the mixture parameter set in the spatial structure.




Fig. 4. Our adjoint solution estimates for the volumetric radiance quantities L and L_i for two positions below the surface of the Buddha statue. The center column shows the ground-truth spherical function of L and L_i evaluated using path tracing with 100k samples per pixel. The right column shows our estimates \tilde{L} and \tilde{L}_i fitted from 1k photons. Note that since L_i is an integrated quantity, its estimate is much more reliable than that of L .

In-scattered radiance estimate. Functionally, the in-scattered radiance L_i is given by spherically convolving the incident radiance (represented by a vMF mixture) with the phase function (Equation 3):

$$\tilde{L}_i(\mathbf{x}, \omega) = \Phi(\mathbf{x}) \cdot \int_S f(\mathbf{x}, \omega, \omega') \mathcal{V}_L(\omega' | \Theta_L(\mathbf{x})) d\omega'. \quad (15)$$

Since phase functions are typically rotationally invariant they can be represented with one or more vMF lobes. The convolution of two vMF lobes results in another vMF lobe, which can be calculated analytically (Appendix A.2), providing a highly efficient means to obtain the integrated in-scattered radiance estimate:

$$\begin{aligned} \tilde{L}_i(\mathbf{x}, \omega) &= \Phi(\mathbf{x}) \cdot \mathcal{V}_{L_i}(\omega | \Theta_{L_i}), \\ \mathcal{V}_{L_i}(\omega) &= (\mathcal{V}_f * \mathcal{V}_L)(\omega). \end{aligned} \quad (16)$$

This spatio-angular caching structure thus provides the estimates \tilde{L} and \tilde{L}_i in constant time, minimizing the overhead of our sampling procedures. Examples of our estimates compared to the ground-truth measurements are given in Figure 4.

6 SAMPLING PROCEDURES

We now proceed with details of the procedures for realizing the zero-variance-based sampling decisions described in Section 3.2. We begin with our core contribution: the volume-specific sampling decision that concerns finding a scattering location along a particle’s propagating direction, given its current position.

6.1 Guided Product Distance Sampling

To find the next scattering location, the tracing algorithm first needs to decide if the path should explore the medium contribution L_m or the surface contribution L_s (Figure 2, 1a). In the former case, the distance to the next medium interaction (‘collision distance’) has to be sampled (Figure 2, 1b).

These two decisions can be *merged into a single distance sampling decision*, where either the interval between the current particle position \mathbf{x} and the next surface interaction at \mathbf{x}_s , or the full distance to \mathbf{x}_s is sampled. The zero-variance PDF for this joint distance

sampling is the combination of Equation 7 and Equation 8:

$$p_d(d | \mathbf{x}, \omega) = \frac{T(\mathbf{x}, \mathbf{x}_d) \cdot \sigma_s(\mathbf{x}_d) \cdot L_i(\mathbf{x}_d, \omega)}{L(\mathbf{x}, \omega)}. \quad (17)$$

This merger changes the denominator of the zero-variance PDF for the collision distance from L_m to $L_s + L_m = L$. Since existing tracking methods (Section 2) do not have prior knowledge of the incident and in-scattered radiance, they sample proportionally to the local ('medium') terms only: $p_d(d | \mathbf{x}, \omega) \propto \sigma_s(\mathbf{x}_d)T(\mathbf{x}, \mathbf{x}_d)$. Our aim, in contrast, is to sample the full product PDF in Equation 17. We first explain a naïve *static* solution to this problem, and follow by presenting a significantly more efficient *incremental* proposal.

6.1.1 Naïve Solution. A straightforward way to use our volumetric radiance estimates to sample according to the zero-variance PDF is to explicitly build a tabulated discrete PDF and its corresponding CDF [Fong et al. 2017; Kulla and Fajardo 2012] by a uniform sampling of Equation 17. The downside of this method is that it requires stepping along the *entire length* of the ray to evaluate *all the bins* of the discrete PDF, and calculate its normalization factor – before making any sampling decision (Figure 5, top).

To illustrate the inefficiency of this approach, consider an optically dense medium such as the one in Figure 1. Most of the light transport in such objects takes place near the surface [Křivánek et al. 2014], at depths much smaller compared to the size of the object. The overhead of the naïve approach is then proportional to the ratio of the transport and object scales. Another example are large sparse media (for example in Figure 11), where the important illumination features (light shafts, caustics, etc.) might be too spaced-out to be sufficiently captured by a uniform tracking.

Our key technical contribution is a novel guided distance sampling method that overcomes this limitation by performing only the necessary number of steps to reach the sampled collision distance.

6.1.2 Proposed Incremental Distance Sampling. Instead of constructing the complete discrete PDF along the ray upfront, the tracking procedure *implicitly* estimates it on the fly while stepping through the medium. This is done through a repeated stochastic decision whether the collision event should be located somewhere within the currently examined bin, or in the remainder of the ray (Figure 5, bottom). In the latter case the stepping examines the successive bin in the same fashion, leading to a procedure that can be conceptualized as a *tail recursion*. In the following exposition we then unfold this recursive process into an incremental tracking procedure (Algorithm 1).

Note that such a procedure requires a normalization factor for each individual step. Conveniently, we can approximate this factor thanks to the cached quantities \tilde{L} for the incoming radiance and L_i for the in-scattered radiance (Section 5).

To derive our incremental sampling method we first have a look at the cumulative density function (CDF) of the combined zero-variance PDF from Equation 17. Since the sampling proceeds along the direction $-\omega$ starting at \mathbf{x} , we will for the remainder of this section operate within the 1D distance space with the origin at \mathbf{x} . The CDF defines the probability that a random variable D (representing

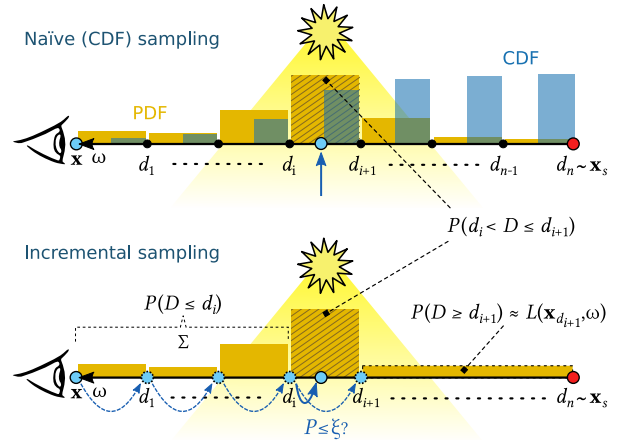


Fig. 5. Naïve (top) versus our incremental guided distance sampling (bottom). The naïve solution generates a discrete PDF (and its corresponding CDF) by first sampling along the entire ray to select a single collision location (blue arrow). Our incremental technique (Algorithm 1) steps along the ray, proposing candidate bins according to the current PDF estimate and the estimated probability mass along the remainder of the ray without explicitly sampling it. This approach is especially beneficial for optically dense media, since on average fewer steps are taken due to the rapidly decreasing transmittance function with regard to increasing d .

the scattering distance) is below or equal to the distance d :

$$P(D \leq d) = \frac{\int_0^d T(0, d') \sigma_s(d') L_i(d', \omega) dd'}{L(0, \omega)}. \quad (18)$$

To rephrase, since the in-scattered radiance contribution of the ray interval from \mathbf{x} to \mathbf{x}_d is a subset of incoming radiance arriving at \mathbf{x} , the probability is the ratio between these two quantities.

Per-bin sampling. Consider first, that the stepping procedure has already arrived (collision-free) to d_i (Figure 5, bottom). We now propose the transition to the next point d_{i+1} , with the goal to determine the probability of the collision being between d_i and d_{i+1} , or alternatively, after d_{i+1} . If the former case occurs, the tracking stops, whereas in the latter case, it restarts again from distance d_{i+1} .

Let us therefore take a look at an individual bin (ray interval) between d_i and d_{i+1} . We will denote the conditional probability of generating a collision within this bin as

$$P_i(D \leq d_{i+1}) \equiv P(D \leq d_{i+1} | d_i < D). \quad (19)$$

Since the bins are typically defined on the scale of the medium or cache resolution, we can reasonably assume a constant L_i and σ_s for the current bin, allowing a closed-form integration of the local transmittance. From Equation 18 we then get

$$P_i(D \leq d_{i+1}) \approx \frac{1 - T(d_i, d_{i+1})}{\sigma_i(d_i)} \cdot \frac{\sigma_s(d_i) \cdot L_i(d_i)}{L(d_i, \omega)}, \quad (20)$$

the first factor being the transmittance integral from d_i to d_{i+1} .

Now, the probability of reaching the point d_i collision-free (i. e., the condition in Equation 19) can be updated incrementally while

stepping over the preceding bins:

$$P(d_i < D) = \prod_{j=0}^{i-1} 1 - P_j(D \leq d_{j+1}). \quad (21)$$

The final PDF for sampling d between d_i and d_{i+1} is then:

$$p_d(d | \mathbf{x}, \omega) = P(d_i < D) \cdot P_i(D \leq d_{i+1}) \cdot p(d | d_i, d_{i+1}, \sigma_t), \quad (22)$$

where the last factor is the PDF for generating the collision distance inside the bin. For details on sampling a distance inside one bin in a homogeneous medium we refer the reader to the work of Kulla and Fajardo [2012].

Replacing the radiance quantities from Equation 20 by our estimates enables us to build an incremental tracking algorithm that samples each bin according to Equation 17. Starting at the first bin ($i = 0$), the algorithm sequentially decides with the probability of $P_i(D \leq d_{i+1})$ if a scattering event should be generated in the current bin. If not, the algorithm updates the probability of not scattering before d_{i+1} (Equation 21) and continues to the next bin. In the case an event is generated in the current bin a distance d between d_i and d_{i+1} is sampled, assuming a constant bin-specific σ_t . The pseudocode of this algorithm is presented in Algorithm 1 and an example of the resulting PDF in Figure 6.

Algorithm 1 Algorithm for incremental guided distance sampling.

```

1: procedure GUIDEDDISTANCESAMPLING( $\mathbf{x}, \omega, dMax$ )
2:   //  $\mathbf{x}$  ... starting position inside or at the beginning of the medium
3:   //  $\omega$  ... direction in which the walk continues
4:   //  $dMax$  ... distance to the next surface intersection
5:    $d_i := 0$ 
6:    $pdf := 1$ 
7:    $scatter := False$ 
8:   while not  $scatter$  and  $d_i < dMax$  do
9:      $\mathbf{x}_i := \mathbf{x} - d_i \omega$ 
10:     $[\tilde{L}, \tilde{L}_i] := \text{LOOKUPESTIMATES}(\mathbf{x}_i, \omega)$  ▷ Sec. 5
11:     $[\sigma_t, \sigma_s] := \text{GETMEDIUMCOEFFICIENTS}(\mathbf{x}_i)$ 
12:     $d_{i+1} := \text{GETDISTANCETO NEXTBIN}(\mathbf{x}_i, \omega)$ 
13:     $P := \text{CALCBINPROBABILITY}(\tilde{L}, \tilde{L}_i, \sigma_s, \sigma_t, d_i, d_{i+1})$  ▷ Eq. 20
14:     $\xi := \text{GETRANDOMVALUE}()$ 
15:    if  $\xi \leq P$  then ▷ current bin selected
16:       $pdf *= P$ 
17:       $[d, pdfBin] := \text{SAMPLEDISTANCEINBIN}(d_i, d_{i+1}, \sigma_t)$ 
18:       $pdf *= pdfBin$ 
19:       $scatter := True$ 
20:    else ▷ continue tracking
21:       $pdf *= (1 - P)$  ▷ Eq. 21
22:       $d_i := d_{i+1}$ 
23:    end if
24:  end while
25:  return [ $scatter, d, pdf$ ]
26: end procedure

```

Sampling stability: local versus global estimates. A key feature of our incremental distance sampling is that it only relies on local estimates of the volumetric radiance quantities $\tilde{L}(d_i, \omega)$ and $\tilde{L}_i(d_i, \omega)$. Especially, using the latter quantity as local normalization factor instead of a single global one (i. e., $\tilde{L}(0, \omega)$ evaluated at the start of the tracking) has multiple advantages:

- Inaccuracies of the local estimates almost solely affect the local scattering decision at the current bin.
- As both estimates derive from the same representation (Section 5), any repeated per-bin misestimation is only manifested in the fluence term which cancels out in Equation 20. In other words, the quality of the per-bin sampling only depends on the representation accuracy of two estimates' ratio.

A demonstration of the PDF resulting from our incremental distance sampling compared to the standard transmittance-based sampling is given in Figure 6. Since our sampling procedure is directly derived from the optimal PDF (Equation 17), it will converge to it when the adjoint estimates are accurate and the step size decreases so that the assumption of a per-bin constant σ_t and L_i holds. Given that the naïve method samples the same PDF, these two procedures will also lead to similar sampling patterns – again, granted that the radiance estimates are accurate enough.

On the other hand, in the sporadic cases when the local adjoint estimates are corrupted (due to insufficient cache data or misfit records), our strategy can suffer from inaccurate estimates for the incident and in-scattered radiance, influencing the sampling quality. Although such inaccuracies *do not* affect the consistency between the generated samples and the calculated PDF used for normalization, sampling bias might still be introduced in scenarios where the generated samples distribution would not cover the *entire integrand*. Such bias results from 'early exits' caused in particular by underestimation of the incident radiance $\tilde{L}(d_i, \omega)$. Similarly to other guiding methods [Dahm and Keller 2017; Müller et al. 2017; Vorba et al. 2014], we use MIS (balance heuristic with $\beta = 0.5$) with the traditional sampling method (i. e., transmittance-based sampling) to circumvent this problem. In addition, we also heuristically bound the probability of scattering inside a bin (Equation 20) to be ≤ 0.9 .

Sampling stability: incoming radiance filter. Figure 4 shows that we usually can expect the $\tilde{L}_i(d_i, \omega)$ estimate to be more accurate than $\tilde{L}(d_i, \omega)$. To improve the incident radiance estimate during the tracking, we can reuse the already evaluated estimates from previous bins to refine our approximation within the current bin at d_i . We use a prediction of the incident radiance in combination with the *exponential moving average* filter for the refinement:

$$\begin{aligned} \tilde{L}^I(d_i, \omega) &= (1 - A) \cdot \tilde{L}^{II}(d_i, \omega) + A \cdot \tilde{L}(d_i, \omega) \\ \tilde{L}^{II}(d_i, \omega) &= \frac{\tilde{L}^I(d_{i-1}, \omega) - (1 - T(d_{i-1}, d_i)) \cdot \alpha(d_{i-1}) \cdot \tilde{L}_i(d_{i-1}, \omega)}{T(d_{i-1}, d_i)} \end{aligned} \quad (23)$$

The extrapolated prediction \tilde{L}^{II} at d_i is obtained from the previous estimate \tilde{L}^I at d_{i-1} , by subtracting the contribution of in-scattered radiance along the bin from d_{i-1} to d_i and inverting the attenuation from d_{i-1} to d_i . The factor $\alpha(d_{i-1}) = \sigma_s(d_{i-1})/\sigma_t(d_{i-1})$ is the scattering albedo of the medium at d_i . To avoid an occasional negative radiance estimate, \tilde{L}^{II} is bounded to be ≥ 0 . The parameter $A = 0.75$ defines how much the refined estimate \tilde{L}^I should rely on the cached estimate \tilde{L} .




Fig. 6. Comparison of different distance sampling strategies for a didactic scene with a heterogeneous medium (3D grid with alternating density). The ray starts from the camera and goes straight through the medium (*red point*). Since the standard transmittance-based sampling (*blue*) ignores the illumination component, it overestimates the importance of the first dense cube of the grid structure. Our incremental guided sampling (*green*) on the other hand is close to the simulated ground-truth PDF (*orange*).

6.2 Guided Product Directional Sampling

To explore the in-scattered radiance at a location in the volume, the third sampling decision in the random walk generation chooses in which *direction* the walk should continue. Traditionally, this decision is only based on the phase function and therefore only covers a part of the integrand in Equation 3.

Our approach is based on the zero-variance PDF in Equation 9: it follows the idea of guiding based on the product of the phase function and the incoming radiance, and as such can be viewed as a direct extension of directional product guiding from surfaces [Herholz et al. 2016] to volumes. The main difference: in contrast to the Gaussian mixtures used by Herholz et al. we exploit the vMF model which better represents the involved spherical (rather than hemi-spherical) distributions.

The product mixture \mathcal{V}_{\otimes} is calculated based on the mixture \mathcal{V}_L of the incoming radiance and the mixture \mathcal{V}_f of the phase function:

$$\mathcal{V}_f(\omega)\mathcal{V}_L(\omega) = (\mathcal{V}_f \otimes \mathcal{V}_L)(\omega) = \mathcal{V}_{\otimes}(\omega). \quad (24)$$

The new scattering direction ω is obtained by importance-sampling this product mixture following the stable procedure presented by Jakob [2012]. We use the same vMF approximation of the phase function as when calculating in-scattered radiance (Section 5). To account for inaccuracies in the estimates of incoming radiance we use multiple importance sampling (MIS) between our product mixture and the actual phase function. Formulas for calculating the product mixture, including our derivation of numerically stable approximation of the integral of the product of the vMF models, are given in Appendix A.2. Figure 7 shows a comparison of different directional sampling methods for a dense anisotropic medium.

6.3 Guided Russian Roulette and Splitting

Stochastic path termination via Russian roulette (RR) is a fundamental part of algorithms based on random walks: it ensures correct termination without having to impose hard limits on path lengths, and prevents wasteful evaluation of long paths with disproportionately low or even null contributions. Using the correct *guided* RR probability (cf. Equation 10) is especially important in the context of zero-variance-theoretical sampling, since an overly aggressive

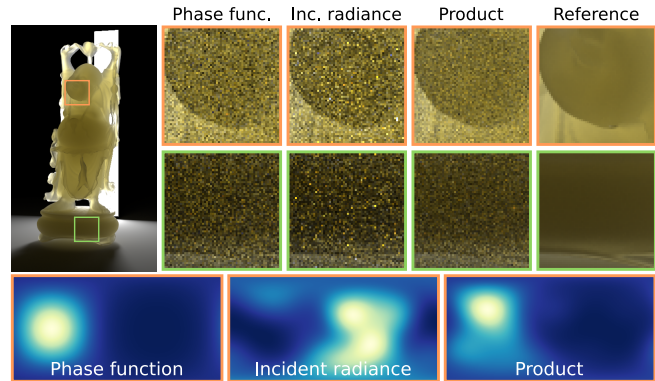


Fig. 7. Comparison of phase function-based sampling, MIS-based sampling between the incoming radiance and the phase function (akin to Vorba et al. [2014]) and our guided product directional sampling using MIS. For the MIS-based methods the same conservative probability of 0.5 is used. The phase function of the medium is anisotropic with a mean cosine of $g = 0.6$. The bottom row shows the PDFs of each methods (without MIS) at the (*orange*) sample location below the surface.

path termination can dramatically increase variance [Vorba and Krivánek 2016].

An intuitive way to reason about guided Russian roulette is through an analogy with *rejection sampling*: the distance and directional sampling procedures (Sections 6.1 and 6.2) can be thought of as processes that jointly yield path proposals. These proposals are accepted or rejected using RR such that the accepted paths' distribution more closely matches the desired, zero-variance distribution.

We therefore seek an acceptance (survival) probability to arrive at this path distribution. For an estimator to have zero variance, it follows that the contribution of every path to a given pixel should be equal. This suggests an ideal path survival probability proportional to the path's expected contribution: low-contribution paths are more likely to be terminated (rejected), so that surviving paths all contribute equally (that is, after normalizing by the survival probability). This notion then naturally extends to *path splitting*: if a path's contribution is higher than the corresponding pixel expectancy, we split the path to explore the path space more evenly.

The intuition above has been formalized for surface transport by Vorba and Krivánek [2016] in their *adjoint-driven Russian roulette and splitting* approach. As they point out, RR and splitting are related to the *go with the winners method* [Grassberger 2002]. Instead of heuristically determining the RR and splitting factor [Szirmay-Kalos 2005], they use an adjoint estimate of the actual light transport. We build atop their formalism and extend it to volumetric media, referring to it as *guided Russian roulette and splitting* (GRR&S) to be consistent with the remaining strategies presented here. For simplicity, we discuss the case of paths traced from the camera, for which the adjoint solution is given by the equilibrium radiance (incident L and in-scattered L_i). A generalization to light paths is straightforward [Vorba and Krivánek 2016].

Volumetric guided RR and splitting. GRR&S computes the Russian roulette and splitting factor $q \in [0, \infty)$ (Equations 6 and 11) for a partially-generated path as the ratio between the path's *expected*

contribution $E[r]$, and pixel's true value I , i. e., $q = E[r]/I$. Whenever $q \leq 1$, we apply RR to terminate the path with the probability of $1 - q$. On the other hand, if $q > 1$, the path's expected contribution is higher than the actual pixel value, and we split it into $n = \lfloor q \rfloor + 1$ sub-paths (the extra sub-path with the probability $q - \lfloor q \rfloor$).

What sets volumetric transport apart from surface transport is that generating a new path vertex now involves *two* steps: distance and direction sampling. Unlike on surfaces, Russian roulette and splitting can be applied before either of these two steps, yielding *distance RR&S* and *directional RR&S* processes.

Distance GRR&S. Consider a partially constructed path reaching a position \mathbf{x}_j . Suppose we have already sampled the direction ω_j and that the current particle weight is a_r . For distance GRR&S, we compute the expected path contribution as the product of a_r and the appropriate adjoint (here, the *incident radiance*) [Hoogenboom 2008] as:

$$E[r]_{\text{dist}} \approx a_r \cdot \tilde{L}(\mathbf{x}_j, \omega_j). \quad (25)$$

If we RR-terminate the path, we stop right at \mathbf{x}_j ; splitting, on the other hand, leads to several independent distance samples distributed along the ray $(\mathbf{x}_j, -\omega_j)$, as illustrated in Figure 8a,b.

Directional GRR&S. Suppose the path has reached the next position \mathbf{x}_{j+1} from direction ω_j with the weight:

$$a'_r = a_r \frac{T(\mathbf{x}_j, \mathbf{x}_{j+1})\sigma_s(\mathbf{x}_{j+1})}{p_d(d | \mathbf{x}_j, \omega_j)}, \quad (26)$$




Fig. 8. Volumetric guided RR and splitting. A volume split (marked green) can generate both new sampled distances (*left*) and directions (*right*).

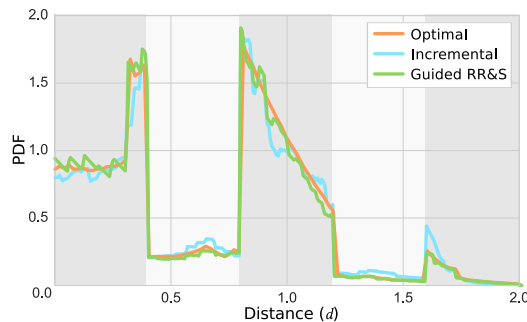


Fig. 9. Guided RR&S (green) compensating for inaccurate adjoint estimates used in our incremental distance sampling (blue), in the configuration from Figure 6. By using guided RR&S the resulting samples distribution matches the optimal one (orange) more closely.

which encapsulates the weight a_r for reaching the state (\mathbf{x}_j, ω_j) with the addition of all factors encompassing the sampling of the distance $d = \|\mathbf{x}_{j+1} - \mathbf{x}_j\|$. The appropriate adjoint for this case is the *in-scattered radiance* L_i and we have

$$E[r]_{\text{dir}} \approx a'_r \cdot \tilde{L}_i(\mathbf{x}_{j+1}, \omega_j). \quad (27)$$

Splitting generates several independent scattering directions ω_{j+1} (and, consequently, multiple new sub-paths – see Figure 8c,d).

Figure 9 shows an example of the directional GRR&S compensating for the inaccuracy in the adjoint estimate for \tilde{L} used in our incremental collision sampling. Section 7.3 presents additional results for both the distance and directional GRR&S.




Fig. 10. Comparison of the pixel estimates \tilde{I} (generated by accumulating the in-scattered radiance from our adjoint) to the ground-truth solution. Through the use of a *weight window* [Vorba and Krivánek 2016] even this rough estimate is sufficient to guide our guided RR&S.

Practical considerations. To calculate the termination/splitting factor q , we require an estimate \tilde{I} of the true pixel value I . As with Vorba and Krivánek [2016], we pre-compute these estimates from our cached adjoint (radiance) solution: we trace several camera rays per pixel and, for each, gather the in-scattered radiance by marching through the volume. Figure 10 illustrates these solution estimates. An alternative way to estimate I would be to use a denoised version of the rendered image after a small initial batch of samples, which then could be refined over time after each sampling iteration.

In practice, we do not enforce the q -based termination and splitting rules strictly, in the form discussed above: instead, we use q to center a *weight window* [Vorba and Krivánek 2016] to accommodate for imprecisions in our cached adjoint estimates.

Finally, note that since GRR&S works as an *ex-post* compensation for any deviations from the desired zero-variance scheme, it is desirable to apply it as often as possible along a path. This happens in an alternating fashion: the directional GRR&S follows the distance sampling (Section 6.1), and vice versa, the distance GRR&S is applied after the directional sampling step (Section 6.2).

GRR&S and zero-variance sampling. An important distinction between the zero-variance path termination (as presented in Section 3.2) and our volumetric GRR&S lies in the fact that we can ever only operate with an approximation of the adjoint solution. Because the distance and directional sampling decisions have no notion of path history, any sub-optimal sampling decisions (due to adjoint representation inaccuracy or MIS with traditional unguided sampling methods) can accumulate and result in deviations from the true zero-variance solution. *Conducting guided RR and splitting*

Scene	Pre-process times (sec)				Cache stats		
	PT	Fit	\bar{I}	Total	Dep	#Ent	Mem
INFINITESCAN HEAD	22	66	<1	45	22	149k	104MB
NATURAL HISTORY	32	25	<1	57	18	67k	48MB
VOLUME CAUSTIC	84	29	1.2	104	22	75k	52MB

Table 1. *Left*: pre-processing times for the main result scenes from Figure 11, broken down to individual stages – photon tracing (‘PT’), EM-fitting of the directional mixtures (‘Fit’), pixel expectancy estimate (‘ \bar{I} ’). *Right*: adjoint cache statistics – kD-tree depth (‘Dep’), number of cache entries (‘#Ent’) and memory occupancy (‘Mem’). Notably, while the cost of the photon-tracing stage is strongly scene-dependent, the remaining steps have approximately linear dependence w. r. t. the configured accuracy of the cache.

based on the expected path contribution compensates for any such deviations, since the expected contributions incorporate both a notion of path history (in the path weight) as well as its expected future contribution (in the cached adjoint solution).

As such, RR and splitting are a key balancing mechanism to any approximate zero-variance solution – for that reason we apply GRR&S after each sampling decision, rather than just on light sources as implied in Equation 11. Note that, in the hypothetical case of ideal distance and directional sampling this would not be necessary, and in fact splitting itself could be omitted entirely.

7 RESULTS AND EVALUATION

7.1 Implementation

We implemented our guided volume path tracer in the Mitsuba renderer [Jakob 2010]. Although we explain our method in the context of the last-event estimator (Equation 4), we in fact utilize next event estimation for direct illumination and single scattering. Its inclusion is analogous with surface guiding, we therefore refer the reader to Vorba et al. [2014; 2016]. To account for remaining inaccuracies in our radiance estimates we use multiple importance sampling (MIS) between our guided sampling methods and traditional phase function and transmittance-based sampling. For both methods we use the balance heuristic with a conservative weight of $\beta = 0.5$. Unless specified otherwise, we limit the maximum path length to 40. Whenever the guided Russian roulette (RR) is not used, we disable RR altogether; this way we avoid additional variance caused by standard throughput-based RR.

Stepping through heterogeneous media uses the DDA algorithm on top of the OpenVDB library. All tests were performed on a PC with two Intel Xeon Gold 5115 CPUs (20 physical / 40 logical cores) and 512 GB of RAM. We make use of SSE instructions for accelerated vMF fitting and evaluation [Vorba et al. 2014]. Since we only rely on one training pass and are not updating the mixtures on-the-fly, we use a batch MAP-EM algorithm. For all tests our vMF mixtures use 8 components and 50M–100M photons to fit the cache. The minimum and maximum number of photons per leaf node are set to 50 and 1000. Table 1 lists the statistics on the pre-processing step.

7.2 Scenes

We evaluate our method in several scenes representing different configurations (especially with regard to the effective optical thickness of the media, to gauge different proportions of low- and high-order scattering) and challenging transport features. Equal-time

renderings are the primary form of benchmark, comparing our approach with a standard unidirectional path tracer which importance-samples directions according to the phase function and collision distances based on transmittance. We show that our guided path tracer produces superior results to this industry-standard solution.

The primary application of our guided sampling is indeed unidirectional path tracing due to its numerous advantages in production [Fascione et al. 2017]. Our sampling techniques can, however, be equally well employed in light tracing, photon tracing, bidirectional path tracing, and combinations thereof [Křivánek et al. 2014].

Since our scenes contain both high-intensity areas (e. g., light shafts or caustics) and low-intensity areas, we use the mean relative squared error metric (relMSE) [Rousselle et al. 2011] with a 0.5-percentile outlier removal for regularization. This metric treats both types of regions equally compared to the standard mean squared error (MSE), which is far more sensitive to variance in high-intensity regions. A definition of the metric is given in Appendix B.

INFINITESCAN HEAD. The head (Figure 1 and Figure 11) consists of an optically dense medium, so subsurface scattering will occur near to the surface. For the volume parameters we used the measured data from Jensen et al. [2001]: $g = 0$, $\sigma_t = (1.103, 1.66, 1.935)$ and $\alpha = (0.988, 0.957, 0.925)$. The size of the head corresponds on average to 64 mean free path (MFP) lengths. As suggested by Donner and Jensen [2006], we use a rough dielectric boundary interface. The refractive nature of this interface prevents the use of next event estimation to estimate the volumetric transport in this case. Since the volumetric properties have strong color-dependence, each RGB channel is rendered separately for 10 minutes and the results are merged afterward. Both our guided incremental collision distance sampling and directional sampling improve the efficiency of the path generation. In addition, GRR&S avoids expensive generation of long but weakly contributing paths deep inside the dense medium. In particular, splitting reduces the variance caused by either sub-optimal guiding decisions or MIS with standard sampling techniques.

NATURAL HISTORY. In the museum (Figure 11) light shafts in the homogeneous medium lead to dramatically varying in-scattering. The museum is filled with an anisotropically scattering medium ($g = 0.4$, $\sigma_t = 0.05$ and $\alpha = 0.75$), which has the size of 2 MFP lengths. Our method explores the cached in-scattered radiance for the distance sampling and therefore will more often predict a scattering event inside a shaft. Directional guiding will then predominantly choose a sampling direction pointing towards the light source. Since the light shafts contribute the most to the volumetric transport in the scene, the directional guiding will steer the sampling decisions towards these light shafts.

VOLUME CAUSTIC. Caustic light paths in a heterogeneous medium (Figure 11) are a serious challenge for unguided path tracing. Our approach considers the heterogeneity during the distance sampling, and explores the cached radiance during the guided product distance and product directional sampling to resolve the volumetric caustic properly. The isotropically scattering heterogeneous medium has an albedo of $\alpha = 0.75$ and varying optical density of $\sigma_t \in [0.001, 0.04]$ resulting in an average size of 1/3 MFP lengths (i. e., dominated by single scattering). Even after four days the reflections of the caustics

Scene	Average contributing path length (CPL)			
	Reference	No guiding	Dist.+Dir.	+RR
INFINITESCAN HEAD	6.2	12.3	10.5	5.3
NATURAL HISTORY	7.8	29.0	27.7	9.8
VOLUME CAUSTIC	3.7	3.2	3.8	3.7

Table 2. Statistics about the contributing path lengths for the main result scenes (Figure 11). With the help of our guided sampling strategies the path lengths become significantly closer to the target reference values.

are not fully converged in the reference image, which was generated using bidirectional path tracing.

7.3 Contribution of Individual Decisions

Guided product distance and directional sampling. As visible in Figure 11, independently applying the guided collision distance sampling or the scattering direction sampling can sometimes provide only a minor decrease of solution variance (e. g., in the NATURAL HISTORY scene). Due to the additional overhead it might even be counter-productive at times (see distance sampling in the INFINITESCAN HEAD scene). That is because, the effort to sample a suitable scattering location might be wasted if the following directional sampling is just following the phase function, and vice versa. However, combining both sampling decisions always improves the results. This is in line with the observations made by Dwivedi [1982] in the context of deep penetration transport problems, that ‘synergistic’ directional and distance sampling is necessary to leverage the full potential of the zero variance-based sampling schemes (cf. also [Křivánek and d’Eon 2014; Meng et al. 2016]).

Incremental distance sampling. Despite being asymptotically comparable, the advantage of our distance sampling approach over the naïve method is that it allows us to avoid the complete traversal along the ray before making a sampling decision. Especially in dense media, most of the light transport happens close to the surface, and generally we expect the next collision to be close to the current one due to the exponentially decreasing transmittance function. Building the complete discrete PDF therefore creates an unnecessary overhead, which depends on the transport characteristics in each individual scene. While the naïve sampling always has to access all the caches intersected by the sampled ray, our incremental method stops on average after a third of the full ray distance in our tested scenes, or specifically, after 38 % in NATURAL HISTORY, 16 % in INFINITESCAN HEAD, 45 % in VOLUME CAUSTIC and 30 % in the BUMPY SPHERE scene (see below in Section 7.5). In the equal-time comparison in Figure 12 our method is therefore able to make sampling decisions shortly after entering the medium, resulting in more evaluated samples and consequently smoother images.

Guided RR. An informed Russian roulette strategy guided by the adjoint further improves on top of the other sampling decisions (Figure 11). Especially in dense or bounded media such as in the INFINITESCAN HEAD or the NATURAL HISTORY scenes, the generated paths can experience several scattering events before leaving the scene or the medium. In these cases guided RR identifies and terminates long paths with low potential contribution, leading to more samples evaluated at an equal time.

In Table 2 we compare the average per-pixel *contributing* path length (tracked while generating the reference solution), in each scene with the actual generated average path length produced by different configurations of our estimator. *In a truly zero-variance estimator, the relative number of paths of a certain length should be exactly proportional to the total image contribution due to paths with the corresponding length.* Since the average contributing path lengths are strongly scene-dependent, manually setting an optimal RR starting depth would be prohibitively complicated, and in any case requiring some prior knowledge of the transport solution. By using the adjoint solution to estimate the future contribution of a path, the guided RR automatically leads to a termination strategy close to the zero-variance optimum.

Table 3, left lists the relative ratios of paths terminated by the two guided RR methods. Notably, the ratio between paths terminated by directional RR versus distance RR depends strongly on the particular scene-dependent transport characteristics. For instance, the high optical density in INFINITESCAN HEAD causes a strong in-scattered radiance gradient underneath the surface boundary. It is therefore more likely for the directional RR to terminate paths which would progress deeper inside the volume. In VOLUME CAUSTIC then, the distance RR tends to terminate those samples inside the volumetric caustics which are pointing away from the source.

Guided splitting. Complementing the RR, splitting helps in high-variance regions: here our sampling strategies might still leave some residual variance, due to sub-optimal sampling decisions caused by inaccurate local radiance estimates or by the MIS with traditional sampling methods (e. g., in the VOLUME CAUSTIC scene). Table 3, middle, lists the ratios of paths split by the different methods – in contrast to guided RR, the splitting ratios tend to be similar. Note that by splitting a path the *effective* number of samples grows for the corresponding pixel (Table 3, right).

Scene	Guided RR		Guided splitting		SPP	
	Dir.	Dist.	Dir.	Dist.	Orig.	Effect.
INFINITESCAN HEAD	0.27	0.06	0.22	0.17	1312	1829
NATURAL HISTORY	0.48	0.29	0.29	0.28	1340	2104
VOLUME CAUSTIC	0.05	0.16	0.23	0.25	3796	5685

Table 3. Termination and splitting ratios (i. e., the relative proportions of paths terminated or split by each respective strategy) of the different guided RR&S methods employed in our test scenes. Splitting generates additional paths, on average leading to higher amounts of *effective* samples per pixel.

7.4 Adjoint Solution Accuracy

As discussed in Section 5, the accuracy of the approximate adjoint representation directly impacts the guided estimator variance. Here we investigate how different configurations of our cache system influence the resulting solution quality, using the NATURAL HISTORY scene as a testing ground.

In Figure 13, top we analyze the relationship between the number of mixture components and the variance of our guided estimator. A small number of components (e. g., 4) is not sufficient to represent the complexity of the transport, while with 32 the EM algorithm tends to overfit. In our experiments 8 or 16 components have therefore proven to provide the best tradeoff; we use 8 components to

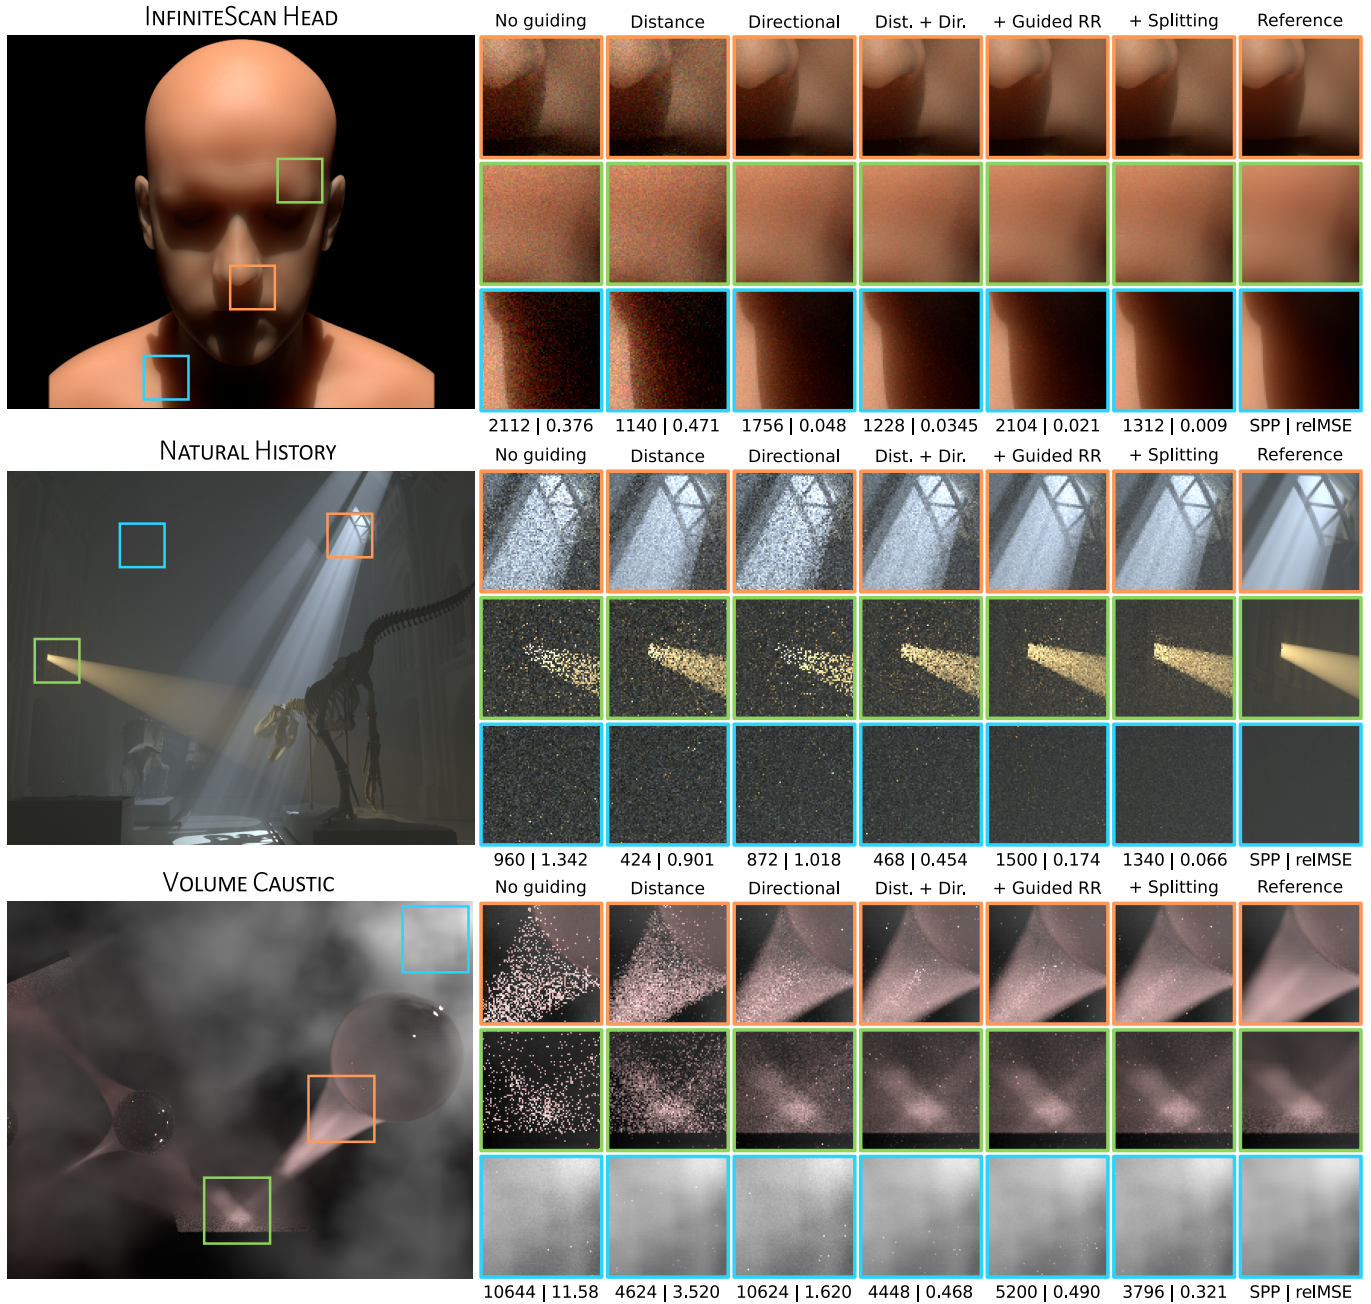


Fig. 11. Contribution of the individual guided sampling decisions to the final image. Note how the image quality gradually improves by considering the adjoint-guided distance sampling, directional sampling, RR and splitting. The images show equal-time comparisons (INFINITESCAN HEAD: 3×10 min, NATURAL HISTORY: 45 min, VOLUME CAUSTIC: 60 min) and compare to the reference results of a bidirectional path tracer running for 3×12 hrs for the INFINITESCAN HEAD scene, 12 hrs for the NATURAL HISTORY and 96 hrs for the VOLUME CAUSTIC.

generate our results, since the faster evaluation of the model makes up for the slightly higher variance of the estimator.

In Figure 13, bottom we then analyze the dependency of our adjoint estimate on the spatial cache resolution. To test this we train our estimates with a large number (400M) of photons, but limit the kD-tree depth. While denser adjoints produce lower variance, we

can see that even with a rather inexpensive pre-processing phase and rough adjoint, our solution clearly surpasses the quality of unguided path tracing. In addition, since in our current implementation the cache density is mainly a consequence of the photon distribution, a higher spatial resolution is generated closer to the light sources. This influences the efficiency of our incremental distance sampling,




Fig. 12. Our incremental distance sampling explores the volume only up to the collision location. Significantly more samples can thus be generated in dense media (compared to the naïve sampling which first generates the full PDF along the ray), leading to better results in (after 10 minutes of rendering). Here, the naïve sampling needs to perform 18 cache look-ups to produce a collision event, while our approach only needs 4, on average.

causing unnecessarily small stepping near sources. To avoid this bottleneck we limit our spatial tree depth between 18 and 22 as a simple but effective heuristic.

A general problem our estimates share with other related guiding approaches [Dahm and Keller 2017; Müller et al. 2017; Vorba et al. 2014] is that, due to the averaging of the directional information over the volume of each cache record (node), high angular frequencies in L may not be reproduced well. An example is provided in Figure 14, bottom. While this has only minor effects when it comes to directional guiding, we sometimes observed negative impact on our incremental distance sampling. In particular, overestimation of L around the small high-frequency areas leads to the expectancy of a higher contribution from behind the volume, and therefore potentially prevents scattering inside the volume. In such cases our method has to rely on MIS with the standard transmittance-based

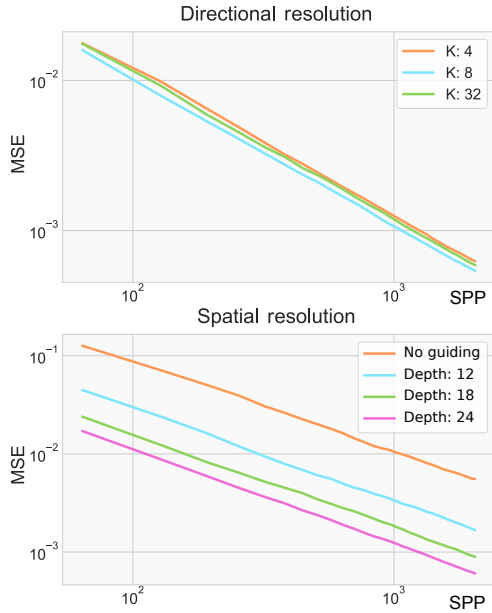


Fig. 13. Evaluating the influence of the directional (top) and spatial (bottom) resolution of the adjoint estimates on the variance of our guided unidirectional path tracer used in the NATURAL HISTORY scene. The log-log plots show the convergence behavior as a function of the number of samples per pixel (up to 2k SPP).




Fig. 14. The effects of averaging the directional data over the volume of a cache. *Bottom:* The averaging blurs high-frequency signals (left) over a larger angular footprint (center), which is also represented in our directional vMF caches (right). *Top:* This spread causes an overestimation of the incident radiance around sharp regions, which might lead to sub-optimal distance sampling decisions.

sampling to generate collisions at the right position, leading to sub-optimal convergence as demonstrated in Figure 14.

One way to overcome this problem in the future could be to extend the lobes of the vMF mixtures to include a weighted-average distance to the source of the incident radiance. The mixture fitting would then incorporate the parallax with respect to the source, which would then function as an on-the-fly adjustment of the lobes’s directions when evaluating the mixture.

7.5 Varying Scattering Anisotropy

In order to evaluate our guiding for different angular frequencies of the transport, we simulate the BUMPY SPHERE scene (Figure 15) with varying degrees scattering anisotropy, ranging from near-isotropic ($g = 0.3$) to highly anisotropic ($g = 0.9$) values. The BUMPY SPHERE is filled with a homogeneous medium with $\alpha = [0.95, 0.67, 0.45]$ and the sphere has a diameter of approximately 6 MFP lengths. To be able to render the scene using a unidirectional path tracer we replaced the point light source by a small spherical area light. Its main difficulty is the specular dielectric interface, which disallows the possibility of direct light connections.

This is most apparent in the single-scattering scenario (Figure 15, top), as in this case all illumination inside the medium consists of refractive caustics, leading to very poor performance of the unguided path tracer. Our sampling (guided by the single-scattering adjoint) is able to localize the caustics, and then lead the tracer towards the source. Even though for the higher scattering anisotropy values the quality of our cached radiance estimates is somewhat decreased (due to spiky energy distributions), our solution remains robust in contrast to the unguided tracer.

For the multi-scattered case (Figure 15, bottom), the diffusion of the caustics makes it easier for the unguided tracer to sample the transport. Still though, the guided solution yields significantly smoother results with at least an order-of-magnitude lower error.

7.6 Alternative Distance Sampling Strategies

Since the incremental guided distance sampling (Section 6.1) is the key novel component of our method, an evaluation against other illumination-aware sampling strategies is warranted. Out of the volume-specific sampling approaches reviewed in Section 2, we compare to the equiangular sampling by Kulla and Fajardo [2012]. Other related methods [Georgiev et al. 2013; Kutz et al. 2017] have the potential to incorporate the knowledge of the adjoint, but this has not been proposed to date.

The equiangular strategy [Kulla and Fajardo 2012] places samples according to the *direct incident illumination* footprint of a *point light source* along the ray, neglecting both the phase function and the attenuation by the medium. As plotted in Figure 16, left for the CHECKERBOARD CUBE scene (cf. Figure 6), this method has the desired ability to place samples in the high-contribution region of the spotlight. The PDF (green) is however far from the optimal one (orange), due to the neglected terms and the fact that it only accounts for direct (single-scattered) illumination. MIS with the standard transmittance-based sampling is therefore necessary, and this procedure still leads to a PDF (purple) which is far from optimal.

Figure 16, right then demonstrates the method in the NATURAL HISTORY scene. Here the equiangular sampling achieves good results in directly illuminated regions (i. e., light shafts), but leads to excessive variance in areas dominated by multiple scattering. This is further exacerbated by the light-selection problem – the strategy has no prior information about the relative contribution of the individual sources. Our sampling, in contrast, does not use a separate strategy for sampling direct illumination, but instead relies on the adjoint to *guide paths towards light sources*.

8 DISCUSSION AND LIMITATIONS

Utility and applicability. Our work shows the importance of regarding the individual sampling decisions in volume rendering as

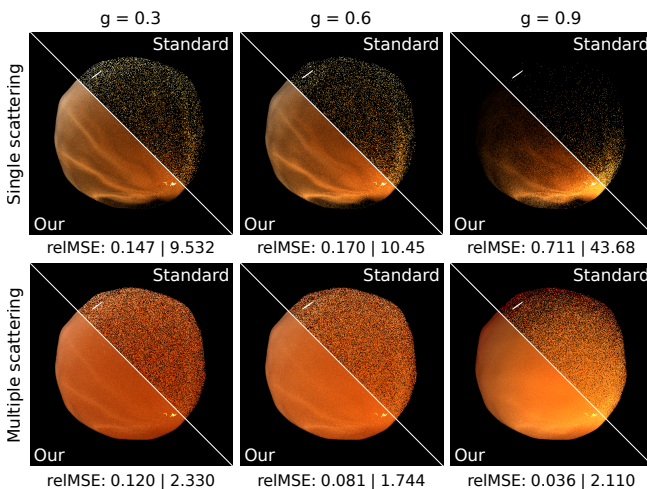


Fig. 15. BUMPY SPHERE scene simulated with 2k SPP at different scattering anisotropies, comparing our guided sampling (below diagonal) against the unguided one (above diagonal). The unguided tracer struggles to sample high-contribution paths, which is especially visible when limiting the transport to a single volumetric bounce (top row).




Fig. 16. Left: Comparing equiangular distance sampling [Kulla and Fajardo 2012] without (green) and with MIS (purple) against the optimal (orange) and our (blue) incremental distance sampling, for the configuration in Figure 6. Right: Comparing our guided sampling against equiangular and standard distance sampling in the NATURAL HISTORY scene.

parts of a holistic sampling strategy. We demonstrated that doing so substantially improves the baseline industry-standard approach in the context of volume rendering, i. e., unguided path tracing based on local sampling decisions. In agreement with the zero-variance sampling theory, we provide apt evidence (Section 7) that phenomena inherently difficult for a unidirectional solver (such as volumetric caustics or light shafts) can be handled efficiently when aided by globally informed, guided sampling.

Furthermore, any method based on incremental path generation will benefit from our results. For instance, guided light tracing would have the ability to distribute photons to visually important locations. Because all sampling decisions in our framework rely on the zero-variance theory, the estimator variance is directly linked with the residual error of our adjoint solution’s representation and the used MIS weights between our guided methods and traditional unguided sampling. Therefore, any optimization directed at the adjoint’s accuracy (fitting quality, spatial distributions, measurement estimates, etc.) in turn improves the accuracy of the estimator itself.

Adjoint estimate accuracy. Despite proving adequate in all our experiments, the presented single-pass photon-based training is arguably not the optimal way to obtain a robust adjoint estimate. Multi-pass methods, such as the bidirectional online learning of Vorba et al. [2014] or Q-learning by Dahm and Keller [2017], could prove to be a better approach. Alternatively, an adaptation of forward learning (akin to Müller et al. [2017]) should be considered, especially since it matches the directionality of our solver.

Spatial cache distribution. In contrast to Müller et al. [2017], the spatial subdivision of our kD-tree caching structure depends on the photon energy distribution from the light sources, rather than their importance to the estimated quantity (that is, the rendered image). The caches’ distribution is therefore skewed towards the sources, causing a performance penalty as unnecessarily small stepping is used during the guided distance sampling in the sources’ vicinity. A more sophisticated caching scheme refined via illumination gradients [Jarosz et al. 2008], or other well chosen heuristics, could therefore positively impact our method’s performance.

MIS weights. Given that the knowledge of the perfect adjoint solution is merely a hypothetical limit case, we must, to an extent,

rely on a combined sampling with the standard ‘local’ methods. Similarly as other related approaches [Dahm and Keller 2017; Müller et al. 2017; Vorba et al. 2014] we employ MIS for this purpose, using a rather conservative weight of 0.5. Lowering this weight to favor the guided sampling strategies can lead to better estimates (as in Herholz et al. [2016]), but can also increase the occurrence of high-intensity outliers (‘fireflies’) in the cases when an occasional badly-fitted adjoint is encountered. An adaptive way to obtain the MIS weights, most likely based on an independently estimated *confidence* of the adjoint fit, would therefore be of great benefit to our and other existing guiding methods.

Additional sampling strategies. The efficiency of each sampling decision could also be individually improved in combination with existing approaches. One example is the work of Kutz et al. [2017] whose tracking approach formulates probabilities for scattering, absorption, and ‘null’ events. Biasing these probabilities according to the adjoint solution could generate samples following the desired product distribution, further increasing the efficiency of our current differential stepping. Another interesting direction would be using the L_1 approximation as a control variate. This is similar to Pegoraro et al. [2008], but using our adjoint estimate for L_1 would overcome their restrictions on the scattering anisotropy.

Even though our strategies are robust enough for sampling low-order effects (e. g., in the scenes from Figures 11 and 15), they still predominantly focus on indirect illumination. Specialized learning-based methods tackling the light-selection problem [Donikian et al. 2006; Vévoda et al. 2018] could therefore be another useful addition in guided transport solvers, especially to handle large production scenes with many sources and optically thin media.

Surface guiding. Finally, as mentioned before, our sampling techniques only cover the volumetric portion of light transport. To obtain a full estimator, the complementary surface-oriented sampling decisions could be performed on the basis of the existing works of Vorba et al. [2014; 2016] and Herholz et al. [2016]. These works likewise rely on the zero-variance-theoretical framework, and store the adjoint solution in similar structures, which hints at a straightforward combination with our work.

Local versus global guiding. We can distinguish between *local* and *global* path generation strategies in path guiding. The recent *global* approaches [Guo et al. 2018; Müller et al. 2018; Zheng and Zwicker 2018], which are based on learning the guiding distributions in the *primary sample space* (PSS), consider sampling of the complete path as a whole. The benefit of the global guiding in PSS is the lack of need for an explicit spatial or directional representation of the radiance distribution, and that additional features such as product sampling with the scattering functions are implicitly encoded in the random number domain (i. e., the PSS). However, the dimension of the PSS grows with the length of the path and – due to the curse of dimensionality – eventually it becomes virtually impossible to gather sufficient training data for longer path lengths. The global guiding approach by Simon et al. [2018] identifies and guides only the – usually small – subset of paths that are hard to sample with traditional techniques, making the global approach feasible.

The *local* path sampling strategies, as defined by Veach [1997], rely solely on local low-dimensional information (e. g., spatial position and incident direction at a path vertex) to perform individual sampling decisions along a path. As such, they do not suffer from the curse of dimensionality and gracefully handle paths of arbitrary lengths. On the other hand, they often make suboptimal sampling decision because of the limited information available to them.

While our method seemingly belongs among the local guiding strategies, it is important to realize that the adjoint information it accesses at the path vertices is the partial solution of the *global* transport problem. In fact, the main result of the zero variance random walk theory [Hoogenboom 2008] is that globally optimal path importance sampling can be achieved through a set of carefully selected local sampling decision. Our method relies on this result.

Relation to Metropolis light transport. While the described adjoint-guided RR and splitting strategies represent a path length control mechanism, there is a complementary view of them: that of an indirect *path density* regulation mechanism. As such they share similarities with Metropolis light transport (MLT) methods, which explore the path space by mutating important paths.

While both methods strive to generate paths from the same target distribution – that proportional to their image contribution – the key difference is in the underlying sampling mechanism, each coming with their own potential issues. Markov chains generated by MLT may suffer from correlation and insufficient global exploration [Šik and Křivánek 2018]. Splitting, on the other hand, may underperform due to its *ex post* nature: a splitting decision is taken only once the path weight has become large, and this can occasionally be too late to prevent an outlier sample from appearing in the image [Vorba and Křivánek 2016]. Still though, exploring the relationship of the two approaches could benefit both classes of methods.

9 CONCLUSION

We present a volume rendering framework in which every sampling decision is derived from the theory of zero-variance path sampling, i. e., is based on optimal importance sampling given the correct adjoint transport solution. Approximating the true adjoint by assembling a guiding structure using vMF mixtures, efficient sampling decision algorithms are proposed for guided distance and directional sampling as well as for path termination and splitting. Compared to previous approaches which only consider partial information during sampling, our decisions based on local estimates of the adjoint solution consider the full volume rendering equation and therefore yield significantly better convergence, while only adding a moderate overhead for maintaining and evaluating the guiding structure.

ACKNOWLEDGMENTS

We want to thank the following persons and institutions for providing some of the models and scenes used in this work: Stanford 3D scanning repository (BUDDHA), Alvaro Luna Bautista and Joel Andersdon (NATURAL HISTORY), Bruce Walter (BUMPY SPHERE), and Infinite Realities (INFINITESCAN HEAD). An additional thanks goes to Robert Hildebrandt for lighting some of the scenes. This work received funding from the European Union’s Horizon

2020 research and innovation program, under the Marie Skłodowska-Curie grant agreement No 642841 (DISTRO). It was further supported by the Czech Science Foundation grants 16-18964S and 19-07626S, and the Charles University grant SVV-2017-260452.

REFERENCES

- John Amanatides and Andrew Woo. 1987. A fast voxel traversal algorithm for ray tracing. In *Eurographics '87*, 3–10.
- Parthasarathy Bagchi and Irwin Guttman. 1988. Theoretical considerations of the multivariate von Mises-Fisher distribution. *Journal of Applied Statistics* 15, 2 (Jan 1988), 149–169.
- Arimand Banerjee, Inderjit S. Dhillon, Joydeep Ghosh, and Suvrit Sra. 2005. Clustering on the unit hypersphere using von Mises-Fisher distributions. *J. Mach. Learn. Res.* 6 (2005), 1345–1382.
- Mark Bangert, Philipp Hennig, and Uwe Oelfke. 2010. Using an infinite von Mises-Fisher mixture model to cluster treatment beam directions in external radiation therapy. In *Conference on Machine Learning and Applications*.
- Thomas Bashford-Rogers, Kurt Debattista, and Alan Chalmers. 2012. A significance cache for accelerating global illumination. *Comput. Graph. Forum* 31, 6 (2012), 1837–1851.
- Laurent Belcour, Kavita Bala, and Cyril Soler. 2014. A local frequency analysis of light scattering and absorption. *ACM Trans. Graph.* 33 (2014), 5.
- Benedikt Bitterli and Wojciech Jarosz. 2017. Beyond points and beams: higher-dimensional photon samples for volumetric light transport. *ACM Trans. Graph.* 36, 4 (2017), 112:1–112:12.
- Norbert Bus and Tamy Boubekur. 2017. Double hierarchies for directional importance sampling in Monte Carlo rendering. *Journal of Computer Graphics Techniques* 6, 3 (2017), 25–37.
- Florent Chatelain and Nicolas Le Bihan. 2013. Von Mises-Fisher approximation of multiple scattering process on the hypersphere. In *International Conference on Acoustics, Speech, and Signal Processing*.
- Per H. Christensen. 2003. Adjoints and importance in rendering: an overview. *IEEE Transactions on Visualization and Computer Graphics* 9, 3 (2003), 329–340.
- Per H. Christensen and Wojciech Jarosz. 2016. The path to path-traced movies. *Foundations and Trends in Computer Graphics and Vision* 10, 2 (2016), 103–175.
- Ken Dahm and Alexander Keller. 2017. Learning light transport the reinforced way. In *ACM SIGGRAPH 2017 Talks*, 73:1–73:2.
- Michael Donikian, Bruce Walter, Kavita Bala, Sebastian Fernandez, and Donald P. Greenberg. 2006. Accurate direct illumination using iterative adaptive sampling. *IEEE Transactions on Visualization and Computer Graphics* 12, 3 (2006), 353–364.
- Craig Donner and Henrik Wann Jensen. 2006. A spectral BSSRDF for shading human skin. In *Proceedings of the 17th Eurographics Conference on Rendering Techniques*, 409–417.
- SR Dwivedi. 1982. A new importance biasing scheme for deep-penetration Monte Carlo. *Annals of Nuclear Energy* 9, 7 (1982).
- Thomas Engelhardt and Carsten Dachsbacher. 2010. Epipolar sampling for shadows and crepuscular rays in participating media with single scattering. In *Proc. 13D*.
- Thomas Engelhardt, Jan Novák, Thorsten-W. Schmidt, and Carsten Dachsbacher. 2012. Approximate bias compensation for rendering scenes with heterogeneous participating media. *Comp. Graph. Forum* 31 (2012), 2145–2154.
- Luca Fascione, Johannes Hanika, Marcos Fajardo, Per Christensen, Brent Burley, and Brian Green. 2017. Path tracing in production. In *ACM SIGGRAPH 2017 Courses*.
- Julian Fong, Magnus Wrenninge, Christopher Kulla, and Ralf Habel. 2017. Production volume rendering. In *ACM SIGGRAPH 2017 Courses*, 2:1–2:79.
- Jeppe Revall Frisvad. 2011. Importance sampling the Rayleigh phase function. *J. Opt. Soc. Am. A* 28 (2011), 2436–2441.
- 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, C (2013), 57–68.
- Iliyan Georgiev, Jaroslav Krivánek, Toshiya Hachisuka, Derek Nowrouzezahrai, and Wojciech Jarosz. 2013. Joint importance sampling of low-order volumetric scattering. *ACM Trans. Graph.* 32, 6 (2013), 164:1–164:14.
- Peter Grassberger. 2002. Go with the winners: a general Monte Carlo strategy. *Computer Physics Communications* 147, 1 (2002), 64–70.
- Jerry Guo, Pablo Bauszat, Jacco Bikker, and Elmar Eisemann. 2018. Primary Sample Space Path Guiding. In *Eurographics Symposium on Rendering - EI & I*, 73–82.
- Sebastian Herholz, Oskar Elek, Jiří Vorba, Hendrik Lensch, and Jaroslav Krivánek. 2016. Product importance sampling for light transport path guiding. *Computer Graphics Forum* 35, 4 (2016), 67–77.
- Heinrich Hey and Werner Purgathofer. 2002. Importance sampling with hemispherical particle footprints. In *Proc. of SCCG*, 107–114.
- J. Eduard Hoogenboom. 2008. Zero-variance Monte Carlo schemes revisited. *Nuclear Science and Engineering* 160, 1 (2008), 1–22.
- Wenzel Jakob. 2010. Mitsuba renderer. (2010). <http://www.mitsuba-renderer.org>.
- Wenzel Jakob. 2012. *Numerically stable sampling of the von Mises-Fisher distribution on S^2 (and other tricks)*. Technical Report, Cornell University.
- Wojciech Jarosz. 2013. The perils of evolutionary rendering research: beyond the point sample. (2013). Eurographics Symposium on Rendering, Invited Talk.
- Wojciech Jarosz, Craig Donner, Matthias Zwicker, and Henrik Wann Jensen. 2008. Radiance caching for participating media. *ACM Trans. Graph.* 27, 1 (2008).
- Wojciech Jarosz, Derek Nowrouzezahrai, Iman Sadeghi, and Henrik Wann Jensen. 2011. A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM Transactions on Graphics* 30, 1 (2011), 5:1–5:19.
- Henrik Wann Jensen. 1995. Importance driven path tracing using the photon map. In *Rendering Techniques (Proc. of EGWR)*.
- Henrik Wann Jensen and Per H. Christensen. 1998. Efficient simulation of light transport in scenes with participating media using photon maps. In *Proc. of SIGGRAPH*, 311–320.
- Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. 2001. A practical model for subsurface light transport. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, 511–518.
- James T. Kajiya. 1986. The rendering equation. *SIGGRAPH Comput. Graph.* 20 (1986).
- Malvin H. Kalos and Paula A. Whitlock. 2008. *Monte Carlo methods*. Wiley-VCH.
- David Koerner, Jan Novák, Peter Kutz, Ralf Habel, and Wojciech Jarosz. 2016. Subdivision next-event estimation for path-traced subsurface scattering. In *Eurographics Symposium on Rendering - Experimental Ideas & Implementations*, 91–96.
- Christopher Kulla and Marcos Fajardo. 2012. Importance sampling techniques for path tracing in participating media. *Comput. Graph. Forum* 31, 4 (2012), 1519–1528.
- 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 (2017), 111:1–111:16.
- Jaroslav Krivánek and Eugene d'Eon. 2014. A zero-variance-based sampling scheme for Monte Carlo subsurface scattering. In *ACM SIGGRAPH 2014 Talks*, 66:1–66:1.
- Jaroslav Krivánek, Iliyan Georgiev, Toshiya Hachisuka, Petr Vévoda, Martin Šik, Derek Nowrouzezahrai, and Wojciech Jarosz. 2014. Unifying points, beams, and paths in volumetric light transport simulation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 33, 4 (2014).
- Eric P. Lafortune and Yves D. Willems. 1995. A 5D tree to reduce the variance of Monte Carlo ray tracing. In *Rendering Techniques*, 11–20.
- NJ McCormick and I. Kuscer. 1973. Singular eigenfunction expansions in neutron transport theory. *Advan. Nucl. Sci. Technol.* 7 (1973), 181–282.
- Johannes Meng, Johannes Hanika, and Carsten Dachsbacher. 2016. Improving the Dwivedi sampling scheme. *Computer Graphics Forum* 35, 4 (2016), 037–044.
- L.W.G. Morgan and D. Kotlyar. 2015. Weighted-delta-tracking for Monte Carlo particle transport. *Annals of Nuclear Energy* 85, C (2015), 1184–1188.
- Thomas Müller, Markus Gross, and Jan Novák. 2017. Practical path guiding for efficient light-transport simulation. *Computer Graphics Forum* 36, 4 (2017), 91–100.
- Thomas Müller, Brian McWilliams, Fabrice Rousselle, Markus Gross, and Jan Novák. 2018. Neural Importance Sampling. *CoRR* abs/1808.03856 (2018). [arXiv:1808.03856](https://arxiv.org/abs/1808.03856)
- R. F. Murray and Y. Morgenstern. 2010. Cue combination on the circle and the sphere. *Journal of Vision* 10, 11 (2010).
- 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 (2018).
- Jan Novák, Derek Nowrouzezahrai, Carsten Dachsbacher, and Wojciech Jarosz. 2012a. Progressive virtual beam lights. *Computer Graphics Forum (Proceedings of EGSR)* 31, 4 (2012).
- Jan Novák, Derek Nowrouzezahrai, Carsten Dachsbacher, and Wojciech Jarosz. 2012b. Virtual ray lights for rendering scenes with participating media. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 31, 4 (2012).
- 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 (2014).
- Vincent Pegoraro, Ingo Wald, and Steven G. Parker. 2008. Sequential Monte Carlo integration in low-anisotropy participating media. In *Proceedings of the Eurographics Conference on Rendering*, 1097–1104.
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically based rendering: from theory to implementation* (3rd ed.). Morgan Kaufmann.
- Mathias Raab, Daniel Seibert, and Alexander Keller. 2006. Unbiased global illumination with participating media. In *Proc. Monte Carlo and Quasi-Monte Carlo Methods*, 591–606.
- Zhong Ren, Kun Zhou, Stephen Lin, and Baining Guo. 2008. *Gradient-based interpolation and sampling for real-time rendering of inhomogeneous, single-scattering media*. Technical Report MSR-TR-2008-51, Microsoft Research.
- Fabrice Rousselle, Claude Knaus, and Matthias Zwicker. 2011. Adaptive sampling and reconstruction using greedy error minimization. In *Proceedings of the 2011 SIGGRAPH Asia Conference*, 159:1–159:12.

- M. Šik and J. Jaroslav Krivánek. 2018. Survey of Markov Chain Monte Carlo Methods in Light Transport Simulation. *IEEE Transactions on Visualization and Computer Graphics* (2018).
- Florian Simon, Johannes Hanika, Tobias Zirr, and Carsten Dachsbacher. 2017. Line integration for rendering heterogeneous emissive volumes. *Computer Graphics Forum (Proc. of EGSR)* (2017).
- Florian Simon, Alisa Jung, Johannes Hanika, and Carsten Dachsbacher. 2018. Selective guided sampling with complete light transport paths. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 37, 6 (2018).
- Jerome Spanier and Ely Meyer Gelbard. 1969. *Monte Carlo principles and neutron transport problems*. Addison-Wesley.
- László Szirmay-Kalos. 2005. Go with the winners strategy in path tracing. *Journal of WSCG* 13, 1-3 (2005), 49–56.
- László Szirmay-Kalos, Iliyan Georgiev, Milán Magdics, Balázs Molnár, and Dávid Légrády. 2017. Unbiased light transport estimators for inhomogeneous participating media. *Computer Graphics Forum* 36, 2 (2017).
- László Szirmay-Kalos, Balázs Tóth, and Milan Magdics. 2011. Free path sampling in high resolution inhomogeneous participating media. *Computer Graphics Forum* 30 (2011), 85–97.
- Eric Veach. 1997. *Robust Monte Carlo methods for light transport simulation*. Ph.D. Dissertation. Stanford University.
- Petr Vévoda, Ivo Kondapaneni, and Jaroslav Krivánek. 2018. Bayesian online regression for adaptive direct illumination sampling. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2018)* 37, 4 (2018).
- P. von Radziewsky, T. Kroes, M. Eisemann, and E. Eisemann. 2017. Efficient stochastic rendering of static and animated volumes using visibility sweeps. *IEEE Transactions on Visualization and Computer Graphics* 23, 9 (2017), 2069–2081.
- Jiří Vorba, Ondřej Karlík, Martin Šik, Tobias Ritschel, and Jaroslav Krivánek. 2014. On-line learning of parametric mixture models for light transport simulation. *ACM Trans. Graph.* 33, 4 (2014), 101:1–101:11.
- Jiří Vorba and Jaroslav Krivánek. 2016. Adjoint-driven Russian roulette and splitting in light transport simulation. *ACM Trans. Graph.* 35, 4 (2016).
- Adolf N. Witt. 1977. Multiple scattering in reflection nebulae. *Astrophysical Journal Supplement Series* 35 (1977).
- E Woodcock, T Murphy, P Hemmings, and S Longworth. 1965. Techniques used in the GEM code for Monte Carlo neutronics calculations in reactors and other systems of complex geometry. In *Proc. Conf. Applications of Computing Methods to Reactor Problems*, Vol. 557.
- Yonghao Yue, Kei Iwasaki, Bing-Yu Chen, Yoshinori Dobashi, and Tomoyuki Nishita. 2010. Unbiased, adaptive stochastic sampling for rendering inhomogeneous participating media. *ACM Trans. Graph. (Proc. SIGGRAPH)* 29 (2010), 177:1–177:8.
- Quan Zheng and Matthias Zwicker. 2018. Learning to Importance Sample in Primary Sample Space. *CoRR* abs/1808.07840 (2018). arXiv:1808.07840 <http://arxiv.org/abs/1808.07840>

A VON MISES-FISHER DISTRIBUTION

We use parametric mixtures of von Mises-Fisher (vMF) distributions to represent directional volumetric radiance estimates. The vMF distribution v is a probability distribution on a $d - 1$ dimensional sphere in \mathbb{R}^d and describes the distribution of random, unit-length vectors ω . We use vMF distributions in three dimensions in the numerically robust form due to Jakob [2012]:

$$v(\omega|\mu, \kappa) = \frac{\kappa}{2\pi(1 - e^{-2\kappa})} e^{\kappa(\mu^T \omega - 1)}. \quad (28)$$

The distribution is parametrized by its mean direction μ and precision, a.k.a. concentration, κ . It is rotationally invariant around μ .

A.1 Parameter Estimation

To fit the mixture distributions for the radiance estimates described in Section 5 a *weighted batch MAP-EM* algorithm (similar to the one described by Vorba et al. [2014]) is used. Given an observed dataset (e.g., photons in a cache cell), MAP-EM estimates the component parameters based on a combination of their maximum likelihood estimates and priors associated with those proposed parameters. Since this work uses vMF mixtures instead of Gaussian mixtures [Vorba et al. 2014], we describe in the following the maximum likelihood

estimates for the vMF parameters, including their weighted version, and the used κ parameter prior.

Maximum likelihood. We can estimate the vMF parameters by maximizing the likelihood of an observed data vector $\{\omega_1, \dots, \omega_N\}$ as follows. The mean direction is given by normalizing the average of the observed directions, $\mu = \frac{\sum \omega_i}{\|\sum \omega_i\|}$. The precision κ cannot be estimated in a closed form but Banerjee et al. [2005] showed that it can be approximated from the average cosine \bar{r} of the data around the mean direction:

$$\kappa \approx \frac{\bar{r}(d - \bar{r}^2)}{1 - \bar{r}}, \quad \text{where } \bar{r} = \frac{\|\sum \omega_i\|}{N}. \quad (29)$$

Weighted ML estimation. We can estimate the ML parameters of a vMF model also for a *weighted* dataset, where each data point ω_i is associated with a weight w_i describing its contribution. To do so, the data points ω_i are replaced by their weighted version ω'_i ,

$$\omega'_i = N \frac{w_i}{\sum_j w_j} \omega_i; \quad (30)$$

we then use the same ML estimation formulas as for the unweighted case. The multiplication by N is a matter of convenience: It has no bearing on the weighted ML estimation, but allows us to treat the maximum a-posteriori estimation for the unweighted and weighted cases in the same way (see Equation 32 below).

Maximum a-posteriori. To avoid overfitting, one can approach parameter estimation as a maximum a-posteriori (MAP) problem, imposing priors on μ and κ . Since there is no prior indication that the volumetric light transport is biased towards a specific direction, no directional prior is used for the mean directions of the vMF components. We use a (conjugate) prior only for κ ; it is given by

$$\lambda(\kappa|\alpha, \beta) \propto \left(\frac{\kappa}{2\pi(e^\kappa - e^{-\kappa})} \right)^\alpha e^{\kappa(\alpha\beta)}. \quad (31)$$

The parameter α is the strength of the prior and can be interpreted as the number of observations drawn from the prior distribution with an average cosine of β . To integrate this prior in the parameter estimation of κ , we replace the calculation of the average cosine from Equation 29 with

$$\bar{r}_\lambda = \frac{\alpha\beta + \|\sum \omega_i\|}{\alpha + N}. \quad (32)$$

An overview of additional priors and (unweighted) MAP-EM for vMF mixtures is given by Bagchi and Guttman [1988] and Bangert et al. [2010].

A.2 Convolution and Product

Closed-form solutions for convolving and calculating the product distribution of two vMF models are respectively used to efficiently obtain the in-scattered radiance estimate (Section 5) and to sample the scattering direction according to the product of the incident radiance and the phase function (Section 6.2).

Single lobes. The *convolution* $v_i * v_j$ of two vMF lobes v_i and v_j can be approximated with another vMF lobe v_k . Chatelain and Le Bihan [2013] show that the average cosine of v_k is well approximated by the product of the average cosines of v_i and v_j , i.e., $\bar{r}_k = \bar{r}_i \bar{r}_j$. The precision κ_k can then be calculated using Equation 29. If the

average cosine of a vMF lobe is unknown, it can be approximated by $\bar{r} \approx \frac{1}{\tanh(\kappa)} - \frac{1}{\kappa}$. When estimating the in-scattered radiance via convolution of the incoming radiance vMF mixture, we use a vMF representation of the phase function as convolution kernel (cf. Equation 3). The mean directions of the mixture components stay unchanged in the case of a forward-scattering phase function and get inverted otherwise.

The *product* of two vMF lobes v_i and v_j forms another vMF lobe v_k multiplied by a scaling factor s_k , i. e., $v_i v_j = s_k v_k$, with

$$\kappa_k = \|\kappa_i \mu_i + \kappa_j \mu_j\| \quad \mu_k = \frac{\kappa_i \mu_i + \kappa_j \mu_j}{\kappa_k}. \quad (33)$$

The scaling factor s_k in fact represents the integral of the product. Following Murray and Morgenstern [2010] and combining it with the numerically stable evaluation of vMF by Jakob [2012] (Equation 28), we present a numerically stable formula for s_k :

$$s_k = \frac{\kappa_i \kappa_j (1 - e^{-2\kappa_k})}{2\pi \kappa_k (1 - e^{-2\kappa_i})(1 - e^{-2\kappa_j})} e^{\kappa_i(\mu_i^T \mu_k - 1) + \kappa_j(\mu_j^T \mu_k - 1)}. \quad (34)$$

Mixtures of vMF lobes. The product of two vMF mixtures \mathcal{V}_0 and \mathcal{V}_1 with K_0 and K_1 components results in a mixture \mathcal{V}_2 with $K_2 = K_0 \cdot K_1$ components. The vMF lobes of this new mixture are the

products of the combinations between the individual components of each mixture ($v_{2,k} = v_{0,i} \cdot v_{1,j}$). For these lobes the weights are derived by the individual component weights and the product scaling factor:

$$\pi_{2,k} = \frac{s_{2,k} \pi_{0,i} \pi_{1,j}}{\sum_{n=1} \sum_{m=1} s_{2,l} \pi_{0,n} \pi_{1,m}}. \quad (35)$$

The scaling factor $s_{2,k}$ stands for the integral of the product of the i -th and j -th components of the mixtures \mathcal{V}_0 and \mathcal{V}_1 and $s_{2,l}$ for the product of the n -th and m -th components.

B RELATIVE MEAN SQUARED ERROR METRIC

In Section 7 the relative mean squared error metric is used to compare the result images (I_a), produced by of our zero variance-based framework or alternatives, against ground-truth images (I_b) generated by a bidirectional path tracer. The metric calculates the average *relative* squared error over all image pixels (indexed by n):

$$\text{relMSE}(I_a, I_b) = \frac{1}{N} \sum \left(\frac{I_a(n) - I_b(n)}{I_b(n) + \epsilon} \right)^2. \quad (36)$$

The avoid numerical instabilities wherever $I_b(n)$ is close to zero, a small bias of $\epsilon = 1e^{-3}$ is added in the denominator.