

Unifying Points, Beams, and Paths in Volumetric Light Transport Simulation

Jaroslav Křivánek¹ Iliyan Georgiev² Toshiya Hachisuka³ Petr Vévoda¹
Martin Šik¹ Derek Nowrouzezahrai⁴ Wojciech Jarosz⁵

¹Charles University in Prague ²Light Transportation Ltd. ³Aarhus University ⁴Université de Montréal ⁵Disney Research Zürich




Figure 1: We present a new rendering algorithm, unified points, beams, and paths (UPBP), for light transport simulation in participating media. UPBP combines a number of volumetric estimators based on points and beams with Monte Carlo path sampling techniques, all of which have complementary strengths and weaknesses. The bottom row shows the contributions of the various volumetric transport techniques to the final image at the top. As shown in Fig. 2, UPBP excels at rendering scenes with different kinds of media, where previous specialized techniques each fail in complementary ways. (Scene inspired by Gkioulekas et al. [2013].)

Abstract

Efficiently computing light transport in participating media in a manner that is robust to variations in media density, scattering albedo, and anisotropy is a difficult and important problem in realistic image synthesis. While many specialized rendering techniques can efficiently resolve subsets of transport in specific media, no single approach can robustly handle all types of effects. To address this problem we unify volumetric density estimation, using point and beam estimators, and Monte Carlo solutions to the path integral formulation of the rendering and radiative transport equations. We extend multiple importance sampling to correctly handle combinations of these fundamentally different classes of estimators. This, in turn, allows us to develop a single rendering algorithm that correctly combines the benefits and mediates the limitations of these powerful volume rendering techniques.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing

Keywords: global illumination, light transport, participating media, bidirectional path tracing, photon mapping, photon beams

Links: [DL](#) [PDF](#) [WEB](#)

1 Introduction

Light scattering in participating media is responsible for many important visual phenomena. However, simulation of these effects can incur a large computational cost as variations in media density (e.g., haze vs. skin), scattering albedo (wine vs. milk), and scattering anisotropy result in significantly different light interaction behaviors. As such, designing a single light transport simulation algorithm that is *robust* to these variations remains an important and open problem.

Two classes of widely adopted light transport simulation approaches excel at rendering complex volumetric shading effects: those based

©ACM, 2014. This is the authors' version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in ACM Transactions on Graphics, 33, 4, July 2014. doi.acm.org/10.1145/2601097.2601219

on Monte Carlo integration [Georgiev et al. 2013] and those based on photon density estimation [Jarosz et al. 2011a]. Several different volumetric photon density estimators specialize in handling complex effects such as indirect caustics, where bidirectional Monte Carlo path-tracing (BPT) performs poorly. While previous work has empirically demonstrated the power of the density estimation techniques, no well-founded analysis of their relative performance exists.

We perform a canonical variance analysis of the various volumetric radiance estimators categorized by Jarosz et al. [2011a] in order to quantify the behavior of their relative performance under different configurations. We show that while photon beam-based estimators are beneficial for relatively sparse media, dense media are still often better handled by traditional point-based density estimators.

Based on our findings, we seek to combine the strengths of all potential volumetric estimators with the versatility of BPT in a principled way. To that end, we develop an extension of multiple importance sampling (MIS) that accommodates the combination of these fundamentally different approaches into a single rendering algorithm (see Fig. 1). The algorithm, which we call *unified points, beams, and paths (UPBP)*, excels at rendering scenes with different kinds of media, where previous specialized techniques each fail in complementary ways (see Fig. 2). In summary, we make the following major contributions:

- We present a variance analysis showing the complementary merits of point- and beam-based volumetric estimators.
- We extend multiple importance sampling to enable the combination of point-based, beam-based, and Monte Carlo volumetric estimators.
- We implement a practical combined algorithm for robust light transport simulation in scenes with different media.

2 Related Work

Light transport simulation. Photorealistic image synthesis requires finding solutions to the equations of light transport [Immell et al. 1986; Kajiyama 1986; Veach 1997; Pauly et al. 2000]. The two dominant classes of light transport simulation methods are photon density estimation [Jensen 1996; Hachisuka et al. 2008; Jarosz et al. 2011a] and Monte Carlo (MC) path integration [Kajiyama 1986; Veach 1997; Pauly et al. 2000]. Our work builds upon the path integral formulation developed by Veach [1997] in order to frame light transport simulation as an MC process. We review its extension to media [Pauly et al. 2000] in Sec. 4.1.

The formulation of light transport simulation as a Monte Carlo integration problem forms the basis of the seminal multiple importance sampling (MIS) framework [Veach and Guibas 1995]. MIS exploits the fact that the same light-carrying path can be sampled by several path sampling approaches (or “techniques”) with different probability densities. Veach and Guibas [1995] showed how these different techniques correspond to different lengths of the eye and light sub-paths of the path, which led to the formulation of *bidirectional path tracing*. Since some estimators capture certain light transport paths more efficiently than others, the idea of combining several estimators into one algorithm has proven extremely useful.

While the original MIS framework considered only Monte Carlo path integration, recent work has successfully extended it to include surface photon density estimation [Georgiev et al. 2012; Hachisuka et al. 2012]. Such a combination is particularly effective when handling paths that are (nearly) impossible to sample with Monte Carlo path integration, such as specular reflections of caustics from small light sources [Veach 1997; Hachisuka et al. 2008]. Our work is inspired by these recent advances but we propose a new framework for combining photon density estimation and Monte Carlo path inte-




Figure 2: Equal-time comparison of our UPBP algorithm against previous work on the Still Life scene from Fig. 1. Only volumetric transport is shown. Sec. 9 provides details on how these images have been rendered.

gration. We also show how it can be applied to the more challenging problem of rendering participating media.

Rendering algorithms for participating media have traditionally employed direct extensions of surface rendering algorithms [Lafortune and Willems 1996; Jensen and Christensen 1998]. Here, path vertices are placed not only on surfaces but also in the volume of the media. The recently proposed “beams” [Jarosz et al. 2011a] representation revealed that this traditional approach is in fact a legacy of surface-based approaches, and that significant improvements [Novák et al. 2012a; Novák et al. 2012b] are possible once path segments (i.e., beams) are considered as fundamental entities of the light transport simulation process. More recently, Georgiev et al. [2013] exploited the path segment interpretation to derive more efficient estimators for MC path integration in media to jointly sample multiple path vertices. Our work reformulates beam-based radiance estimates [Jarosz et al. 2011a] in a manner that permits their incorporation in a novel, combined photon density estimation and Monte Carlo path integration rendering approach.

Neutron transport. Several works in graphics have observed the ties between neutron transport simulation and participating media rendering [Arvo and Kirk 1990; Arvo 1993; D’Eon and Irving 2011; Georgiev et al. 2013]. The use of kernel density estimation, however, has started relatively recently in the neutron transport community. The work of Banerjee [2010] estimates collision densities from a set particle collision sites, roughly corresponding to volume photon mapping [Jensen and Christensen 1998]. Dunn and Wilson [2012] extended this idea to the estimation from particle tracks (or trajectories), corresponding to photon beams [Jarosz et al. 2011a]. Our variance analysis of different volumetric estimators in Sec. 5 is inspired by a similar analysis by MacMillan [1966] for the collision, track-length, and expected value estimators [Spanier 1966].

3 Overview

Our main goal is to develop a volumetric light transport simulation algorithm that is robust to a wide range of media properties. To achieve this goal, we draw inspiration from previous work which combines estimators with complementary strengths and weaknesses [Veach and Guibas 1995; Georgiev et al. 2012; Hachisuka et al. 2012].

Which estimators should we combine? Jarosz et al. [2011a] derived a number of point- and beam-based volumetric radiance estimators from photon density estimation (Sec. 4.2). To gain an intuition about their relative performance, we analyze their variance (Sec. 5). We additionally exploit the analysis setup to relate the point and beam estimators to the collision, track-length, and expected value estimators from the neutron transport literature [Spanier 1966; Spanier and Gelbard 1969]. The analysis shows that the estimators’ relative variance depends on the configuration (specifically, on the ratio of their kernel size to the mean free path of the media), and that beam estimators may in fact (somewhat counter-intuitively) yield higher variance than point estimators. These results provide a strong incentive for combining point and beam-based estimators in a robust light transport simulation algorithm for volumetric media. In addition, previous work [Vorba 2011; Georgiev et al. 2012; Hachisuka et al. 2012] has shown that using photon density estimators in a bidirectional context and combining them with BPT can provide further benefits. All these approaches rely on the path integral framework, reviewed in Sec. 4.1.

How can we combine the estimators? While multiple importance sampling [Veach and Guibas 1995] is a powerful tool for combining estimators, previous work [Georgiev et al. 2012; Hachisuka et al. 2012] has identified fundamental difficulties that arise when applying MIS to combine photon mapping with Monte Carlo path integral estimators, such as BPT. To resolve those problems, we extend MIS in Sec. 6 in a manner that permits combination of estimators of integrals over spaces of different dimension (such as the various point and beam estimators and MC path integral estimators). Our MIS extension has potentially broader implications outside of the context of light transport in participating media.

The combined algorithm. In Sec. 7, we map the volumetric point and beam estimators to a form that is compatible with our extended MIS. This results in a provably good combination of the estimators that we implement into a single algorithm for robust volumetric light transport simulation, detailed in Sec. 8.

4 Background

We begin by defining our notation and reviewing the path integral formulation of light transport and volumetric radiance estimators.


4.1 Path Integral Framework

The path integral framework [Veach 1997; Pauly et al. 2000] expresses image pixel intensity I as an integral over the space Ω of light transport paths: $I = \int_{\Omega} f(\bar{\mathbf{x}}) d\mu(\bar{\mathbf{x}})$. A length- k path $\bar{\mathbf{x}} = \mathbf{x}_0 \dots \mathbf{x}_k$ has $k \geq 1$ segments and $k + 1$ vertices, with its first vertex \mathbf{x}_0 on a light source, its last vertex \mathbf{x}_k on the eye sensor, and the $\mathbf{x}_1 \dots \mathbf{x}_{k-1}$ scattering vertices on surfaces and/or in media. The differential path measure $d\mu(\bar{\mathbf{x}})$ is a product measure corresponding to area and volume integration for surface and medium vertices, respectively. The measurement contribution function $f(\bar{\mathbf{x}})$ is the product of emitted radiance $L_e(\mathbf{x}_0) = L_e(\mathbf{x}_0 \rightarrow \mathbf{x}_1)$, path throughput $T(\bar{\mathbf{x}})$, and detector sensitivity $W_e(\mathbf{x}_k) = W_e(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k)$:

$$f(\bar{\mathbf{x}}) = L_e(\mathbf{x}_0) T(\bar{\mathbf{x}}) W_e(\mathbf{x}_k). \quad (1)$$

The path throughput $T(\bar{\mathbf{x}})$ is the product of the geometry and transmittance terms for path segments, and scattering function for the inner path vertices, as expressed and illustrated below:

$$T(\bar{\mathbf{x}}) = \left[\prod_{i=0}^{k-1} G(\mathbf{x}_i, \mathbf{x}_{i+1}) T_r(\mathbf{x}_i, \mathbf{x}_{i+1}) \right] \left[\prod_{i=1}^{k-1} \rho(\mathbf{x}_i) \right]. \quad (2)$$



The geometry term for a path segment $\mathbf{x}\mathbf{y}$ is given by $G(\mathbf{x}, \mathbf{y}) = V(\mathbf{x}, \mathbf{y}) \frac{D(\mathbf{x} \rightarrow \mathbf{y}) D(\mathbf{y} \rightarrow \mathbf{x})}{\|\mathbf{x} - \mathbf{y}\|^2}$, where $D(\mathbf{x} \rightarrow \mathbf{y}) = |n_{\mathbf{x}} \cdot \omega_{\mathbf{x}\mathbf{y}}|$ if \mathbf{x} is on a surface, and $D(\mathbf{x} \rightarrow \mathbf{y}) = 1$ if \mathbf{x} is in a medium, and likewise for $D(\mathbf{y} \rightarrow \mathbf{x})$. Here $n_{\mathbf{x}}$ is the surface normal at \mathbf{x} and $\omega_{\mathbf{x}\mathbf{y}}$ is a unit-length vector from \mathbf{x} to \mathbf{y} . $V(\mathbf{x}, \mathbf{y})$ is the visibility indicator function. The transmittance of segment $\mathbf{x}\mathbf{y}$ is given by

$$T_r(\mathbf{x}, \mathbf{y}) = \exp \left(- \int_0^{\|\mathbf{x} - \mathbf{y}\|} \sigma_t(\mathbf{x} + t\omega_{\mathbf{x}\mathbf{y}}) dt \right). \quad (3)$$

We define the scattering function as

$$\rho(\mathbf{x}_i) = \begin{cases} \rho_s(\mathbf{x}_{i-1} \rightarrow \mathbf{x}_i \rightarrow \mathbf{x}_{i+1}) & \text{if } \mathbf{x}_i \text{ on surface} \\ \rho_p(\mathbf{x}_{i-1} \rightarrow \mathbf{x}_i \rightarrow \mathbf{x}_{i+1}) \sigma_s(\mathbf{x}_i) & \text{if } \mathbf{x}_i \text{ in medium,} \end{cases} \quad (4)$$

where ρ_s and ρ_p are the bidirectional scattering distribution function (BSDF) and phase function, respectively. σ_s and σ_t denote the scattering and extinction coefficients.

The path integral can be evaluated with an unbiased MC estimator $\langle I \rangle = \frac{1}{m} \sum_{j=1}^m f(\bar{\mathbf{x}}_j) / p(\bar{\mathbf{x}}_j)$ that averages estimates from m random paths $\bar{\mathbf{x}}_j$ sampled using a *path sampling technique* with probability distribution $p(\bar{\mathbf{x}}) d\mu(\bar{\mathbf{x}})$. The path pdf $p(\bar{\mathbf{x}})$ is given by the joint density of the individual path vertices, i.e., $p(\bar{\mathbf{x}}) = p(\mathbf{x}_0, \dots, \mathbf{x}_k)$, and it is determined by the path sampling technique employed to generate the path. For example, *bidirectional path tracing* (BPT) [Lafortune and Willems 1993; Veach and Guibas 1994] generates paths by independently sampling one subpath from a light and another from the eye, optionally connecting them with an edge. The different path sampling techniques in BPT for generating a given path correspond to the different lengths of the light and eye subpaths. The full path pdf is then given by the product of the pdfs for the two subpaths, $p(\bar{\mathbf{x}}) = p(\mathbf{x}_0 \dots \mathbf{x}_s) p(\mathbf{x}_t \dots \mathbf{x}_k)$. The subpath pdf reflects the local sampling techniques used to generate the individual subpath vertices, and can be written as a product of vertex pdfs $p(\mathbf{x}_i | \text{vertices sampled before } \mathbf{x}_i)$.

In our notation, we express directional pdfs $p(\omega)$ w.r.t. the solid angle measure, distance pdfs $p(t)$ w.r.t. the Euclidean length on \mathbb{R}^1 , and volume vertex pdfs $p(\mathbf{x})$ w.r.t. the Euclidean volume on \mathbb{R}^3 . In participating media, converting from the solid angle \times length product measure to the volume measure involves multiplication by the geometry term G .

We define the *subpath contribution*, or *weight*, for light and eye subpaths as the partial evaluation of a path integral estimator:

$$C_l(\mathbf{x}_0 \dots \mathbf{x}_i) = L_e(\mathbf{x}_0) \frac{T(\mathbf{x}_0 \dots \mathbf{x}_i)}{p(\mathbf{x}_0 \dots \mathbf{x}_i)} \quad (5)$$

$$C_e(\mathbf{x}_j \dots \mathbf{x}_k) = \frac{T(\mathbf{x}_j \dots \mathbf{x}_k)}{p(\mathbf{x}_j \dots \mathbf{x}_k)} W_e(\mathbf{x}_k). \quad (6)$$

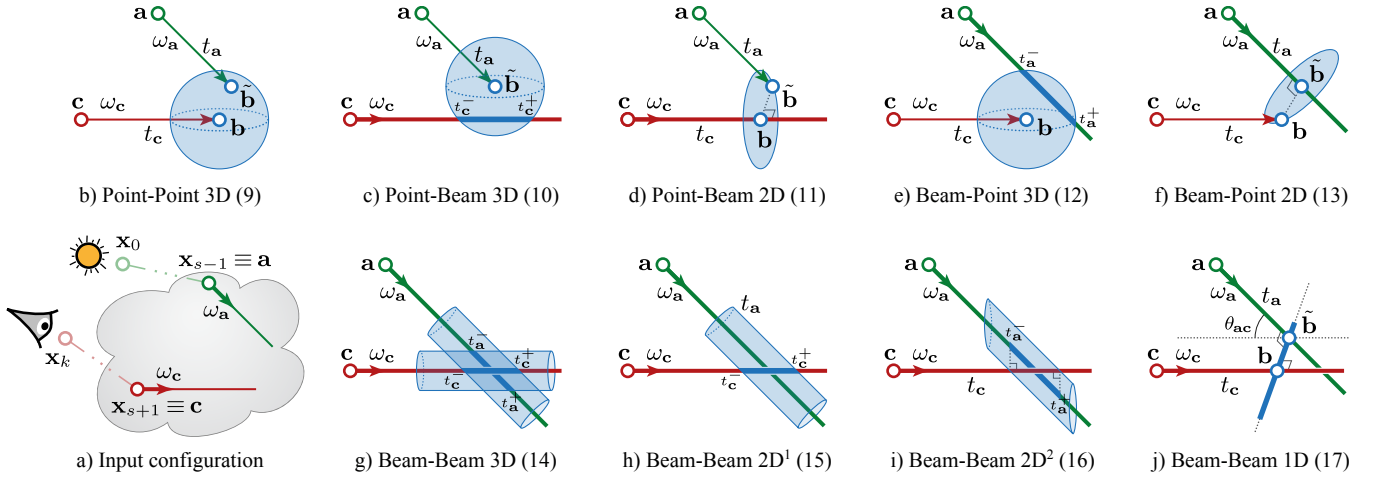


Figure 3: Illustration of the volumetric radiance estimators considered by Jarosz et al. [2011a], generalized for light and eye subpaths of arbitrary length. The number in each label denotes the equation number of the corresponding estimator.

4.2 Volumetric Estimators

Jarosz et al. [2011a] introduced nine distinct volumetric radiance estimators that differ in the representation of equilibrium radiance they employ, the radiometric quantities they estimate, and the dimension of blur employed when reconstructing the desired quantity from the samples. We express these estimators using the path integral notation, permitting their evaluation on any segment along a subpath. In contrast to the aforementioned unbiased path integral estimators, Jarosz et al.’s volumetric estimators are *not* given as general path sampling techniques, i.e. they have no notion of path contribution function and path pdf. Providing such a formulation is one contribution of this paper.

Fig. 3a shows the shared geometric setup for our formulation of Jarosz et al.’s estimators. A light subpath extends to vertex \mathbf{x}_{s-1} , hereafter denoted as \mathbf{a} for brevity, and a direction ω_a is sampled from \mathbf{a} . The ray (\mathbf{a}, ω_a) defines a *photon beam*, whose energy is given by the light subpath weight including scattering at \mathbf{a}

$$C_l(\mathbf{x}_0 \dots \mathbf{a}] = C_l(\mathbf{x}_0 \dots \mathbf{a}) \frac{\rho(\mathbf{a})}{p(\omega_a)}. \quad (7)$$

Similarly, an eye subpath extends up to vertex \mathbf{x}_{s+1} , denoted \mathbf{c} , and a direction ω_c is sampled from \mathbf{c} . The ray (\mathbf{c}, ω_c) defines a *query beam* with weight w.r.t. the pixel estimate given by

$$C_e[\mathbf{c} \dots \mathbf{x}_k] = \frac{\rho(\mathbf{c})}{p(\omega_c)} C_e(\mathbf{c} \dots \mathbf{x}_k). \quad (8)$$

By sampling a distance t_a along the ray (\mathbf{a}, ω_a) , we could create a *photon* at position $\tilde{\mathbf{b}} \equiv \tilde{\mathbf{x}}_s$ with weight¹ $C_l(\mathbf{x}_0 \dots \mathbf{a}] \frac{T_r(t_a)}{p(t_a)}$. We could similarly create a *query point* at $\mathbf{b} \equiv \mathbf{x}_s$ by sampling a distance t_c along the ray (\mathbf{c}, ω_c) . We could then follow Jarosz et al. [2011a] and write some of the estimators in terms of the photons and query points. We instead treat the photon beam (\mathbf{a}, ω_a) and the query beam (\mathbf{c}, ω_c) as the common input to *all* the estimators, and include the terms involved with the calculation of the photon or query point weights into the estimator expressions themselves. With this convention in place, every estimator operates on a shared input, and estimators with the same blur dimension calculate the same pixel value estimate, allowing their direct comparison.

¹We use “weight” instead of the more common term “power” since the quantity associated with volume photons in our notation has units of $[\text{W} \cdot \text{m}]$.

$$\langle I \rangle_{\text{P-P3D}} = \underbrace{\frac{T_r(t_a)}{p(t_a)}}_{\text{photon sampling}} \rho(\tilde{\mathbf{b}}, \mathbf{b}) K_3(\tilde{\mathbf{b}}, \mathbf{b}) \underbrace{\frac{T_r(t_c)}{p(t_c)}}_{\text{query point sampling}} \quad (9)$$

$$\langle I \rangle_{\text{P-B3D}} = \underbrace{\frac{T_r(t_a)}{p(t_a)}}_{\text{photon sampling}} \int_{t_c^-}^{t_c^+} \rho(\tilde{\mathbf{b}}, \mathbf{b}_{t_c}) K_3(\tilde{\mathbf{b}}, \mathbf{b}_{t_c}) T_r(t_c) dt_c \quad (10)$$

$$\langle I \rangle_{\text{P-B2D}} = \underbrace{\frac{T_r(t_a)}{p(t_a)}}_{\text{photon sampling}} \rho(\tilde{\mathbf{b}}, \mathbf{b}) K_2(\tilde{\mathbf{b}}, \mathbf{b}) T_r(t_c) \quad (11)$$

$$\langle I \rangle_{\text{B-P3D}} = \int_{t_a^-}^{t_a^+} T_r(t_a) \rho(\tilde{\mathbf{b}}_{t_a}, \mathbf{b}) K_3(\tilde{\mathbf{b}}_{t_a}, \mathbf{b}) dt_a \underbrace{\frac{T_r(t_c)}{p(t_c)}}_{\text{query point sampling}} \quad (12)$$

$$\langle I \rangle_{\text{B-P2D}} = T_r(t_a) \rho(\tilde{\mathbf{b}}, \mathbf{b}) K_2(\tilde{\mathbf{b}}, \mathbf{b}) \underbrace{\frac{T_r(t_c)}{p(t_c)}}_{\text{query point sampling}} \quad (13)$$

$$\langle I \rangle_{\text{B-B3D}} = \int_{t_c^-}^{t_c^+} \left[\int_{t_a^-}^{t_a^+} T_r(t_a) \rho(\tilde{\mathbf{b}}_{t_a}, \mathbf{b}_{t_c}) K_3(\tilde{\mathbf{b}}_{t_a}, \mathbf{b}_{t_c}) dt_a \right] T_r(t_c) dt_c \quad (14)$$

$$\langle I \rangle_{\text{B-B2D}^1} = T_r(t_a) \int_{t_c^-}^{t_c^+} \rho(\tilde{\mathbf{b}}, \mathbf{b}_{t_c}) K_2(\tilde{\mathbf{b}}, \mathbf{b}_{t_c}) T_r(t_c) dt_c \quad (15)$$

$$\langle I \rangle_{\text{B-B2D}^2} = \left[\int_{t_a^-}^{t_a^+} T_r(t_a) \rho(\tilde{\mathbf{b}}_{t_a}, \mathbf{b}) K_2(\tilde{\mathbf{b}}_{t_a}, \mathbf{b}) dt_a \right] T_r(t_c) \quad (16)$$


$$\langle I \rangle_{\text{B-B1D}} = T_r(t_a) \rho(\tilde{\mathbf{b}}, \mathbf{b}) \frac{K_1(\tilde{\mathbf{b}}, \mathbf{b})}{\sin \theta_{ac}} T_r(t_c) \quad (17)$$

Eqns. (9–17) give the expressions of the estimators, which are also illustrated in Figs. 3b–j. While we only use the minimum-blur estimators (i.e., P-P3D, P-B2D and B-B1D) in our implementation, we list all the estimators for the sake of completeness. All the estimators share the same prefix $C_l(\mathbf{x}_0 \dots \mathbf{a}]$ and postfix $C_e[\mathbf{c} \dots \mathbf{x}_k]$ which we purposefully omit for notational brevity. K_d denotes a normalized d -dimensional kernel. The scattering function ρ at a query location \mathbf{x}_j is evaluated with the direction of the photon beam or photon, which may not pass through this location. To describe this behavior, we amend our definition of ρ as

$$\rho(\mathbf{x}_i, \mathbf{x}_j) = \begin{cases} \rho_s(\mathbf{x}_{i-1} \rightarrow \mathbf{x}_i, \mathbf{x}_j \rightarrow \mathbf{x}_{j+1}) & \mathbf{x}_j \text{ on surface} \\ \rho_p(\mathbf{x}_{i-1} \rightarrow \mathbf{x}_i, \mathbf{x}_j \rightarrow \mathbf{x}_{j+1}) \sigma_s(\mathbf{x}_j) & \mathbf{x}_j \text{ in medium.} \end{cases}$$

We opt to fold the scattering coefficient σ_s , which appears in Jarosz et al.'s [2011a] estimator formulas, into the scattering term ρ . The estimator abbreviations provide information about the radiance data (photon **P**oints or photon **B**eams), query type (**P**oint or **B**eam) and the kernel dimension (1D, 2D or 3D), in exactly this order (e.g., B-P3D refers to beam data \times point query, 3D kernel.) This follows our convention that mathematical terms appear in left-to-right order according to the natural flow of light; note that this opposes the convention used by Jarosz et al. [2011a].

“Long” and “short” beams. Jarosz et al. [2011a] derived photon beams assuming a beam extends until the closest surface, with the transmittance along the beam appearing as a part of the estimator. We refer to this as “long” beams (Fig. 4a). Jarosz et al. [2011b] proposed an unbiased approximation of transmittance by several step functions, and coined this approach “progressive deep shadow maps.” Approximating transmittance by a single step function yields beams of finite extent where the transmittance vanishes (it is replaced by a constant 1). We refer to this as “short” beams (Fig. 4b). The same idea can be applied to query beams. To distinguish between these options, we will use B_s and B_l to denote short and long beams respectively, either of which can be used in place of any B in the estimators in Eqns. (9–17). These extra choices result in a total of 25 distinct estimators. In the next section we derive explicit formulas for the variance of all of these estimators in a canonical configuration. In Sec. 7.3 we derive the effect of the step approximation on the pdfs of the respective path sampling techniques.



a) Long photon beams b) Short photon beams c) Photon points

Figure 4: “Long” and “short” photon beams, and photon points.

5 Variance Analysis of Volumetric Estimators

We present a variance analysis of the estimators in Eqns. (9–17), along with their “short-beam” variants, in a canonical configuration representative of their use in light transport simulation. Our goal is to gain insights into the relative performance of the different point- and beam-based estimators, and to motivate their combination in a robust rendering algorithm. We also establish a firm relation of these estimators to the neutron transport literature [Spanier and Gelbard 1969]. Since the analysis serves mainly as a motivation for the rest of the paper, readers interested in practical results can skip to Sec. 6.

Analysis setup. Since our goal is to compare the variance of the estimators under the same settings, we consider a canonical configuration corresponding to the common estimators’ input depicted in Fig. 3a, where the rays (\mathbf{a}, ω_a) and (\mathbf{c}, ω_c) are given and held fixed. The distances t_a and t_c from the origins of the two rays to the point closest to the other ray (where the kernel is centered) are, therefore, also fixed. The shape of the kernel is an arbitrary choice, so we use a d -dimensional cube kernel for all the estimators as its separability simplifies the analysis. We use a line for 1D estimators, a square for 2D estimators, and a cube for 3D estimators (see Fig. 5). Finally, to facilitate the derivation of closed-form variance expressions, we assume that the two input rays are orthogonal and the medium is homogeneous. These simplifying assumptions are acceptable since our goal is to form an insight into the general behavior of variance, not a generic result that would hold for any configuration.




Figure 5: Canonical configuration with orthogonal light and eye rays and a 1D kernel (a), 2D kernel perpendicular to the light ray (used in the B-P2D and B-B2D¹ estimators) (b) and to the eye ray (in the P-B2D and B-B2D² estimators) (c), and a 3D kernel (d).

5.1 Variance Derivation

In the above setup, any of the analyzed estimators can be written as

$$\langle I \rangle_* = \langle A \rangle \rho w^{-d} \langle C \rangle, \quad (18)$$

where the phase function value ρ is fixed in our setup, so it can be disregarded in the variance analysis, w is the kernel width, d is the kernel dimension, and the definition of the $\langle A \rangle$ and $\langle C \rangle$ terms depends on the kernel type used by the estimator (see Fig. 5):

$$\langle A \rangle = \begin{cases} \langle T_r(t_a) \rangle & \text{for the 1D and 2D}^1 \text{ kernels} \\ \langle \hat{T}_r(t_a^-, t_a^+) \rangle & \text{for the 2D}^2 \text{ and 3D kernels} \end{cases} \quad (19)$$

$$\langle C \rangle = \begin{cases} \langle T_r(t_c) \rangle & \text{for the 1D and 2D}^2 \text{ kernels} \\ \langle \hat{T}_r(t_c^-, t_c^+) \rangle & \text{for the 2D}^1 \text{ and 3D kernels.} \end{cases} \quad (20)$$

Here $\langle T_r(t) \rangle$ denotes a transmittance estimator, and $\langle \hat{T}_r(t^-, t^+) \rangle$ denotes an estimator of transmittance integrated over the kernel:

$$\hat{T}_r(t^-, t^+) \equiv \int_{t^-}^{t^+} T_r(t') dt' = \frac{e^{-\sigma_t t^-} - e^{-\sigma_t t^+}}{\sigma_t}. \quad (21)$$

We consider two possible transmittance estimators, corresponding to either 1) short beams or 2) long beams, and three integrated transmittance estimators, corresponding to either 3) points, 4) short beams, or 5) long beams. The different combinations of these five estimators yield all the 25 pixel estimators that we wish to analyze.

Under our orthogonality assumption, the $\langle A \rangle$ and $\langle C \rangle$ terms are statistically independent, so the expected value and variance of any of the analyzed pixel estimators can be calculated from the first and second moments of $\langle A \rangle$ and $\langle C \rangle$ using the standard relations:

$$E[\langle I \rangle_*] = \rho w^{-d} E[\langle A \rangle] E[\langle C \rangle] \quad (22)$$

$$V[\langle I \rangle_*] = \rho^2 w^{-2d} (E[\langle A \rangle^2] E[\langle C \rangle^2] - E[\langle A \rangle]^2 E[\langle C \rangle]^2). \quad (23)$$

In other words, to derive the variance, we only need to derive the first and second moments of the different estimators of the transmittance value, $T_r(t)$, and its integral, $\hat{T}_r(t^-, t^+)$, which we do next.

5.1.1 Moments of the transmittance value estimators

Short beams. Short beams estimate the transmittance $T_r(t_0)$ at a distance t_0 from the ray origin by sampling a distance t from the pdf proportional to transmittance, $p(t) = \sigma_t T_r(t)$, and evaluating:

$$\langle T_r \rangle_{B_s} = \begin{cases} 1 & \text{if } t \geq t_0 \\ 0 & \text{otherwise.} \end{cases} \quad (24)$$

This estimator’s expected value,

$$E[\langle T_r \rangle_{B_s}] = \int_{t_0}^{\infty} 1 p(t) dt = T_r(t_0), \quad (25)$$

is also equal to its second moment, $E[\langle T_r \rangle_{B_s}^2]$, because $1^2 = 1$.

Long beams. Long beams evaluate the transmittance directly:

$$\langle T_r \rangle_{B_1} = T_r(t_0), \quad (26)$$

so we have $E[\langle T_r \rangle_{B_1}] = T_r(t_0)$ and $E[\langle T_r \rangle_{B_1}^2] = T_r(t_0)^2$.

5.1.2 Moments of the transmittance integral estimators

We now shift our attention to the estimation of the transmittance integral (21). To establish the relation of the point and beam estimators to the collision, track-length, and expected value estimators from the neutron transport field [Spanier and Gelbard 1969], we use the latter three to estimate this integral. We consider a variant of the neutron transport estimators used to calculate collision rates of particles inside an (imaginary) detector, which in our case corresponds to the integration region defined by the kernel. Traditionally, in a random walk, the *collision estimator* contributes (or “scores”) $1/\sigma_t$ upon every particle-medium collision that occurs in the detector volume, the *track-length estimator* scores the length of every particle trajectory (or path) inside the detector, and the *expected value estimator* scores the expectation of every track-length (or path segment) in the detector. In our setup each particle will only contribute upon its first interaction in the integration region, since the shared configuration of the density estimators established in Sec. 4.2 integrates only single scattering events between the given light and eye subpaths. Our analysis bears some similarity to that of MacMillan [1966] for a slab, but we will arrive at different conclusions due to the different analysis setup.

Points (the collision estimator). Points (photons) correspond to sampling a distance t from a pdf proportional to transmittance, $p(t) = \sigma_t T_r(t)$, and evaluating the collision estimator

$$\langle \hat{T}_r \rangle_P = \begin{cases} 1/\sigma_t & \text{if } t \in [t^-, t^+] \\ 0 & \text{otherwise.} \end{cases} \quad (27)$$

It is easy to see that the expected value gives the correct result:

$$E[\langle \hat{T}_r \rangle_P] = \int_{t^-}^{t^+} \frac{1}{\sigma_t} p(t) dt = \int_{t^-}^{t^+} T_r(t) dt = \hat{T}_r. \quad (28)$$

(We omit the arguments of \hat{T}_r for brevity.) The second moment is:

$$E[\langle \hat{T}_r \rangle_P^2] = \int_{t^-}^{t^+} \left[\frac{1}{\sigma_t} \right]^2 p(t) dt = \frac{\hat{T}_r}{\sigma_t}. \quad (29)$$

Short beams (the track-length estimator). Short beams correspond to sampling a distance t using a pdf proportional to transmittance and evaluating the track-length estimator:

$$\langle \hat{T}_r \rangle_{B_s} = \Delta t \quad (30)$$

where Δt , the track length of t through the integration region, is:

$$\Delta t = \begin{cases} 0 & \text{if } t < t^-, \\ t^+ - t^- & \text{if } t > t^+, \\ t - t^- & \text{otherwise.} \end{cases} \quad (31)$$

Again, the expected value of this estimator gives the correct result:

$$E[\langle \hat{T}_r \rangle_{B_s}] = E[\Delta t] = \int_{-\infty}^{\infty} \Delta t p(t) dt \quad (32)$$

$$= \int_{t^-}^{t^+} (t - t^-) p(t) dt + \int_{t^+}^{\infty} (t^+ - t^-) p(t) dt = \hat{T}_r, \quad (33)$$

and the second moment is:

$$E[\langle \hat{T}_r \rangle_{B_s}^2] = E[\Delta t^2] = \int_{-\infty}^{\infty} \Delta t^2 p(t) dt \quad (34)$$

$$= \frac{2}{\sigma_t^2} \left[e^{-\sigma_t t^-} + e^{-\sigma_t t^+} ((t^- - t^+) \sigma_t - 1) \right]. \quad (35)$$

Long beams (the expected value estimator). Long beams correspond to the expected value estimator, which scores the expectation of the track-length through the integration region:

$$\langle \hat{T}_r \rangle_{B_1} = E[\Delta t] = \hat{T}_r. \quad (36)$$

Since this estimator scores the expected value directly, clearly its expectation is $E[\langle \hat{T}_r \rangle_{B_1}] = \hat{T}_r$ and $E[\langle \hat{T}_r \rangle_{B_1}^2] = \hat{T}_r^2$.

5.1.3 Relation to volumetric density estimators.

By using one of the five estimators in Eqns. (24, 26, 27, 30, 36) in place of the $\langle A \rangle$ or $\langle C \rangle$ terms in Eqn. (18), we obtain, in our canonical configuration, any of the 25 different pixel estimators. For instance, we obtain $\langle I \rangle_{P-P3D}$ (9) by using the collision estimator (27) for both $\langle A \rangle$ and $\langle C \rangle$; or, we can obtain $\langle I \rangle_{P-B2D}$ (11) by using the collision estimator for $\langle A \rangle$, and the direct evaluation of transmittance via long beams, Eqn. (26), for $\langle C \rangle$.

5.2 Variance Comparison and Discussion

Given the derivations above, we can now analyze the estimators' variance. One useful comparison is to see how this variance behaves as we change the width w of the blur kernel (expressed in units of mean free paths, $\text{mfp} \equiv 1/\sigma_t$). Since the expected value of most of the estimators depends on w , it is more useful to examine their normalized standard deviation (NSD) $\sigma_\mu[\langle I \rangle_*] = \sqrt{V[\langle I \rangle_*]}/E[\langle I \rangle_*]$. Fig. 6 shows the NSD as a function of kernel width w for the estimators with the lowest possible kernel dimension (P-P3D and the long/short-beam variants of P-B2D, B-P2D, and B-B1D), since those are the most useful for practical implementations. A similar plot for all the 25 estimators is included in the supplemental document. To examine the effect of distance from \mathbf{a} and \mathbf{c} to the center of the kernel, we set $t_a \equiv (t_a^+ + t_a^-)/2 = 5 \text{ mfp}$ and $t_c \equiv (t_c^+ + t_c^-)/2 = 10 \text{ mfp}$. From this experiment we can gain some interesting, and sometimes unintuitive, insights.




Figure 6: Normalized standard deviation (NSD) as a function of the kernel width w for selected volumetric estimators. The $B_l-B_l 1D$ estimator does not appear in the plot because its NSD is zero. A similar plot for all the 25 estimators we consider is included in the supplemental document.




Figure 7: Comparison of the results of different volume radiance estimators for different kernel width w . For small kernels (left column), the beam-based estimators perform better, while for larger kernels (right column) the point-based estimators provide smoother results. When $w \approx 1$ mfp (middle), all the estimators perform roughly the same. The same number of photon points and beams has been used to render these images. The brightness differences among the images are due to the different effects of boundary bias in the various estimators. The noise level differences among the three B_s - B_s 1D images, which are not predicted by our analysis, are due to the additional variance from random sampling of the input configuration. Sec. 10 provides a more detailed discussion.

Points vs. short beams: From the four estimators with highest NSD (P-P, P- B_s , B_s -P, B_s - B_s) we see that short beams perform better for smaller kernels (less than about 1 mfp), whereas points perform better than short beams for larger blur kernels. As the blur goes to zero, however, the variance of point-based estimators diverges to infinity, but the variance of short beams is bounded. Hence, points can be better than short beams in certain cases but they can also be arbitrarily worse. The opposite is not true, however, since short beams have bounded variance for any blur width. Also note that P- B_s and B_s -P have identical variance, in spite of $t_a \neq t_c$.

Long beams: With long beams, either the $\langle A \rangle$ or the $\langle C \rangle$ term contributes zero variance, so the estimators with one long beam appear lower on the plots. Due to the asymptotically poor behavior of points at small kernel widths, the benefit from using long beams in one dimension can be diminished by the high variance of points in the other dimension.

Long beams vs. short beams: Let us now compare B_s -* and *- B_s to their respective long-beam variants. We see that replacing a short beam by a long one along the ray with the origin further away from the kernel ($t_c > t_a$ in our case) reduces variance more than when this change is made along the other ray. This is natural, because the farther away, the lower the probability that point or short beams will reach the kernel, whereas long beams always reach it.

Summary. The practical result of the variance analysis is that points can be better than beams and vice versa, depending on the relative size of the blur kernel to the mean free path. Dense media are better handled by points and thin media by beams. This is illustrated in Fig. 7, where we render single scattering in a simple scene with a candle illuminated from the side with the P-P3D, P- B_s 2D, and

B_s - B_s 1D estimators using the same number of photon points and photon beams. First, when the kernel width is 0.2 mfp of the medium, the performance of the B_s - B_s 1D estimator is far superior to the other two. However, when we set the kernel radius to 5 mfp, the B_s - B_s 1D image becomes significantly more noisy than the other two. As we will see later, our combined algorithm is able to automatically assign high weight to the appropriate estimator. We note that while our analysis assumes homogeneous media, and heterogeneity may have some impact on the estimators' relative variance, it will not change the fact that no single estimator is superior to all others in all circumstances, which is the core of our observation.

6 Extended Multiple Importance Sampling

The analysis in the previous section suggests that it would be beneficial to combine the point- and beam-based estimators. Multiple importance sampling (MIS) [Veach and Guibas 1995] is one approach for combining estimators. However, MIS only considers estimators of integrals on the same space, which is not the case for the different volumetric estimators reviewed in Sec. 4.2. Previous work [Hachisuka et al. 2012; Georgiev et al. 2012] addresses the issue by reformulating the estimators so that they all operate on the same space and standard MIS can be applied. We instead extend the MIS framework to accept estimators of integrals on spaces of different dimension and we devise a corresponding balance heuristic that yields provably good weights for these estimators. To accommodate the increased dimension due to the blurring in the photon density estimators, we build on the idea of *extended space* discussed in the previous work and incorporate it directly into the formulation of a combined estimator. Since this extension may have potential applications beyond the combination of the volumetric estimators considered here, we purposefully keep the initial exposition generic. In Sec. 7 we apply the theory to our specific rendering problem.

Review of multiple importance sampling. Consider an integral $I = \int_{\mathcal{D}_x} f(x) dx$. An estimator $\langle I \rangle = \frac{1}{m} \sum_{j=1}^m f(X_j) / p(X_j)$ is constructed by taking m random variables X_j , $j = 1, \dots, m$ from a *sampling technique* with the pdf $p(x)$. Multiple importance sampling (MIS) [Veach 1997] combines (weighted instances of) estimators with different sampling techniques $p_i(x)$:

$$\langle I \rangle_{\text{MIS}} = \sum_{i=1}^n \frac{1}{n_i} \sum_{j=1}^{n_i} w_i(X_{i,j}) \frac{f(X_{i,j})}{p_i(X_{i,j})}, \quad (37)$$

where n_i is the number of samples taken from the sampling technique with $p_i(x)$, and $w_i(x)$ are heuristic weighting functions that must sum up to one for any x . A provably good choice for $w_i(x)$, in terms of minimizing the variance of $\langle I \rangle_{\text{MIS}}$, is the *balance heuristic*:

$$\hat{w}_i(x) = \frac{n_i p_i(x)}{\sum_{k=1}^n n_k p_k(x)}. \quad (38)$$

6.1 Combining Extended Space Estimators

Let us now consider *extended space* integrals of the form $I_i^E = \int_{\mathcal{D}_x} \int_{\mathcal{D}_{y_i}} f(x, y_i) dy_i dx$ with estimators

$$\langle I_i^E \rangle = \frac{1}{m} \sum_{j=1}^m \frac{f(X_j, Y_{i,j})}{p(X_j, Y_{i,j})}, \quad (39)$$

which can be interpreted as biased estimators of I . In our application, integrals over the different \mathcal{D}_{y_i} will be used to model blurring by the different kernels, and the random variables $Y_{i,j}$ will correspond to photon points and beams. Our goal is to combine estimators of

I with estimators of the potentially different I_i^E 's while minimizing the variance. The main challenge is that these estimators generally have different expected values and, furthermore, the estimators of I have no notion of the integrals over the domains \mathcal{D}_{y_i} .

We begin by extending the combined estimator (37) to accommodate this setup. Suppose we have $u > 0$ sampling techniques $p_1(x), \dots, p_u(x)$, and $n - u > 0$ sampling techniques $p_{u+1}(x, y_{u+1}), \dots, p_n(x, y_n)$ each of potentially different dimension. We define the *extended combined estimator* F^C as

$$F^C = \sum_{i=1}^n \frac{1}{n_i} \sum_{j=1}^{n_i} F_{i,j}^C \quad (40)$$

where the random variables $F_{i,j}^C$ are defined as

$$F_{i,j}^C = \begin{cases} F_{i,j}^E & \text{if } 1 \leq i \leq u \\ F_{i,j}^C & \text{if } u < i \leq n, \end{cases} \quad (41)$$

$$F_{i,j}^E = w_i(X_{i,j}) \frac{f(X_{i,j})}{p_i(X_{i,j})} \quad F_{i,j}^C = w_i(X_{i,j}) \frac{f_i(X_{i,j}, Y_{i,j})}{p_i(X_{i,j}, Y_{i,j})}.$$

This estimator combines u regular and $n - u$ extended space sampling techniques, each taking n_i samples. Note that the functions $f_i(x, y)$ are allowed to be different for each individual extended space technique. This allows us to model estimators with different bias. We have chosen weighting functions that only depend on x and not on y_i , because \mathcal{D}_x is the common domain for both regular and extended space estimators.

6.2 Extended Balance Heuristic

The key to the efficient combination of different estimators with MIS is the use of a good weighting heuristic. In the spirit of Veach [1997], we solve for the weights by minimizing the variance of the estimator in Eqn. (40), arriving at the *extended balance heuristic*:

$$\tilde{w}_i(x) = \frac{n_i / \kappa_i(x)}{\sum_{k=1}^n n_k / \kappa_k(x)}, \quad (42)$$

where

$$\kappa_i(x) = \begin{cases} \frac{f^2(x)}{p_i(x)} & \text{if } 1 \leq i \leq u \\ \int_{\mathcal{D}_{y_i}} \frac{f_i^2(x, y_i)}{p_i(x, y_i)} dy_i & \text{if } u < i \leq n. \end{cases} \quad (43)$$

The different form of $\kappa_i(x)$ for the extended space techniques is due to the extra integration over \mathcal{D}_{y_i} they perform. They estimate this integral by MC sampling (the random variables $Y_{i,j}$), which contributes to the estimator variance and is thus reflected in the weight. In our application, the $Y_{i,j}$ variables correspond to photon points and beams, which are, in fact, samples used to estimate the integral that corresponds to blurring by a kernel.

This heuristic yields a combination with variance within the bounds given by the following theorem. The proof of the theorem, given in the supplemental document, provides a derivation of the heuristic.

Theorem 1. Let F^C be a combined estimator of the form (40), and \tilde{F}^C be the estimator that uses the weighting functions \tilde{w}_i (42). Then

$$V[\tilde{F}^C] - V[F^C] \leq \frac{1}{\min_i n_i} (\mu^+)^2 - \frac{1}{\sum_i n_i} (\mu^-)^2, \quad (44)$$

with $\mu^+ = I + \beta^+$ and $\mu^- = I - \beta^-$, where β^- and β^+ are the lower and upper bounds on the bias of the estimator \tilde{F}^C (please refer to

the supplemental document for the exact definitions of these terms). This theorem is similar to Veach's [1997] Theorem 9.2, and gives a bound on the extra variance that the extended balance heuristic may introduce compared to an unknown, hypothetical optimal weighting function that truly minimizes the variance of F^C .

7 Combining the Estimators

We now show how to apply our extended MIS framework to combine the volumetric estimators in Eqns. (9–17) and unbiased path integral estimators (such as BPT) into a single estimator of the form in Eqn. (40). We need to express each of the estimators in a form accepted by that formula, i.e., to define the variables x and y_i , the function $f(x)$ or $f_i(x, y_i)$, and the pdf $p_i(x)$ or $p_i(x, y_i)$. We take x to be a light transport path \bar{x} as defined in the path integral framework. This way the unbiased path sampling techniques are already given in the desired form $f(x)/p_i(x)$, so we only need to recast the volumetric estimators into the appropriate form. We provide a detailed discussion for the estimators with the minimum blur dimension (P-P3D, P-B2D, B-P2D, B-B1D), as those are the most useful for practical implementations. We also briefly discuss the extension to the remaining estimators.

7.1 Expressing Estimators as Sampling Techniques

We consider the configuration shown in Fig. 3a, where we are given two vertices \mathbf{a} and \mathbf{c} at the end of a light and an eye subpath, respectively, so the full path x would be $\mathbf{x}_0 \dots \mathbf{a} \mathbf{b} \mathbf{c} \dots \mathbf{x}_k$. For notational brevity, we omit all terms on the two subpaths up to and including \mathbf{a} and \mathbf{c} . What remains is the vertex \mathbf{b} where the estimator kernels are centered (Fig. 8), so we have $x \equiv \mathbf{b}$. The y_i variable is, roughly speaking, the point of evaluation of the kernel, as detailed below.

Estimator pdfs. For the **P-P3D** estimator, we set $y_{\text{P-P3D}} \equiv \tilde{\mathbf{b}}$, that is, the 3D photon location (Fig. 8a). The technique's joint pdf is

$$p_{\text{P-P3D}}(x, y) = p(\mathbf{b}, \tilde{\mathbf{b}}) = p(\mathbf{b})p(\tilde{\mathbf{b}}), \quad (45)$$

because \mathbf{b} and $\tilde{\mathbf{b}}$ are sampled independently.

For the **P-B2D** estimator, we first express the photon location $\tilde{\mathbf{b}}$ in a coordinate frame $(\tilde{t}, \tilde{u}, \tilde{v})$ aligned with the query ray, i.e., $\tilde{t} = \omega_c$, as illustrated in Fig. 8b. The 2D variable $y_{\text{P-B2D}}$ is then defined as $y_{\text{P-B2D}} \equiv (\tilde{u}, \tilde{v})$, the photon's u and v coordinates in this frame. With this definition, the joint pdf of the P-B2D technique is:

$$p_{\text{P-B2D}}(x, y) = p(\mathbf{b}, \tilde{u}, \tilde{v}) \stackrel{(1)}{=} p(t_c, \tilde{u}, \tilde{v} | \omega_c) G(\mathbf{b}, \mathbf{c}) p(\omega_c) \stackrel{(2)}{=} p(\tilde{\mathbf{b}}) G(\mathbf{b}, \mathbf{c}) p(\omega_c), \quad (46)$$

where in step (1), we factor out the direction sampling pdf for the query beam using the relationship $p(\mathbf{b}) = p(\omega_c) p(t_c | \omega_c) G(\mathbf{b}, \mathbf{c})$. In step (2) we use $p(\tilde{\mathbf{b}}) = p(t_c, \tilde{u}, \tilde{v} | \omega_c)$ which holds because the




Figure 8: Definition of the variable y for three estimators.

transform from the global frame to the local frame around the query ray is rigid and has no effect on the pdf (the change of coordinates has a unit Jacobian). The derivation for the **B-P2D** estimator is analogous to P-B2D, and its joint pdf reads

$$p_{\text{B-P2D}}(x, y) = p(\omega_{\mathbf{a}})G(\mathbf{a}, \tilde{\mathbf{b}})p(\mathbf{b}). \quad (47)$$

For the **B-B1D** estimator, we set $y \equiv \tilde{d}$, that is, the 1D distance between the two beams (see Fig. 8c). The joint pdf is:

$$\begin{aligned} p_{\text{B-B1D}}(x, y) &= p(\mathbf{b}, \tilde{d}) = p(t_{\mathbf{c}}, \tilde{d} | \omega_{\mathbf{c}})G(\mathbf{b}, \mathbf{c})p(\omega_{\mathbf{c}}) \\ &\stackrel{(1)}{=} p(\omega_{\mathbf{a}}) \left| \frac{dt_{\mathbf{c}} d\tilde{d}}{d\omega_{\mathbf{a}}} \right| G(\mathbf{b}, \mathbf{c})p(\omega_{\mathbf{c}}) \\ &\stackrel{(2)}{=} p(\omega_{\mathbf{a}}) G(\mathbf{a}, \tilde{\mathbf{b}}) \sin \theta_{\text{ac}} G(\mathbf{b}, \mathbf{c}) p(\omega_{\mathbf{c}}), \end{aligned} \quad (48)$$

where the equality $p(t_{\mathbf{c}}, \tilde{d} | \omega_{\mathbf{c}}) = p(\omega_{\mathbf{a}}) \left| \frac{dt_{\mathbf{c}} d\tilde{d}}{d\omega_{\mathbf{a}}} \right|$ used in step (1) is due to a reparametrization of the photon beam direction. In step (2), we have used $\left| \frac{dt_{\mathbf{c}} d\tilde{d}}{d\omega_{\mathbf{a}}} \right| = G(\mathbf{a}, \tilde{\mathbf{b}}) \sin \theta_{\text{ac}}$, illustrated in Fig. 8c. To see why this holds, note that the differential patch $dt_{\mathbf{c}} \times d\tilde{d}$ is observed from the point \mathbf{a} under the solid angle $d\omega_{\mathbf{a}} \cos \alpha / t_{\mathbf{a}}^2 = d\omega_{\mathbf{a}} \sin \theta_{\text{ac}} G(\mathbf{a}, \tilde{\mathbf{b}})$, where $\alpha = \pi/2 - \theta_{\text{ac}}$.

Contribution function. Given these definitions of y_i and the estimators' pdfs $p_i(x, y_i)$, we can now devise the contribution functions f_i from the relation $\langle I \rangle_i = f_i(x, y_i) / p_i(x, y_i)$. Doing this for all the considered estimators yields the contribution functions in the form $f_i(x, y_i) = h_i(x, y_i) K_i(x, y_i)$, where K_i is a 3D, 2D or 1D kernel, respectively, and

$$h_i(x, y_i) = \rho(\mathbf{a})T_r(\mathbf{a}, \tilde{\mathbf{a}})G(\mathbf{a}, \tilde{\mathbf{b}})\rho(\tilde{\mathbf{b}}, \mathbf{b})G(\mathbf{b}, \mathbf{c})T_r(\mathbf{b}, \mathbf{c})\rho(\mathbf{c}). \quad (49)$$

By taking the ratio of the above contribution functions f_i and the pdfs p_i , we indeed obtain the original estimators in Eqns. (9–17), with the additional terms $\rho(\mathbf{a})/p(\mathbf{a})$ and $\rho(\mathbf{c})/p(\mathbf{c})$ that we considered part of the input setup in Fig. 3a and excluded from the estimator expressions (9–17).

Estimators with higher-dimensional kernels. The estimators B-B2D, B-B3D, P-B3D, and B-P3D can be analogously mapped to our framework, depending on how the integrals along the additional dimensions are estimated. If they are computed analytically, then only f_i and K_i need to be adjusted accordingly. If they are estimated using MC integration, we also need to extend y_i to include the additional random variables, and the extended MIS weights will automatically take into account the additional variance due to the stochastic evaluation of these integrals.

7.2 Practical Approximation of the Balance Heuristic

Having expressed each volumetric estimator as a ratio of a contribution function f_i and a pdf p_i , we could use the extended balance heuristic (42) to evaluate the estimator weights in a combined algorithm. However, the evaluation of the $\kappa_i(x)$ term is not practical as it involves evaluating an integral over \mathcal{D}_{y_i} which in our case is the kernel support. We now find an approximation of $\kappa_i(x)$ that leads to an efficient implementation. Assuming that, for any given x , the functions $h_i(x, y_i)$ and $p_i(x, y_i)$ are approximately constant within the support of the kernel K_i (the same assumption as in prior work [Georgiev et al. 2012; Hachisuka et al. 2012]), we have:

$$\begin{aligned} \kappa_i(x) &= \int_{\mathcal{D}_{y_i}} \frac{h_i^2(x, y_i) K_i^2(x, y_i)}{p_i(x, y_i)} dy_i \\ &\approx \frac{h_i^2(x, y_i^*)}{p_i(x, y_i^*)} \int_{\mathcal{D}_{y_i}} K_i^2(x, y_i) dy_i = \kappa'(x, y_i^*), \end{aligned} \quad (50)$$




Figure 9: Weights for selected estimators calculated by the extended balance heuristic, Eqn. (51), in the canonical setup from Sec. 5.

where y_i^* is an arbitrary y_i within the kernel support. The dependency on y_i^* appears in $\kappa'(x, y_i^*)$ because we have pulled the terms in front of the integral. If we furthermore assume that $h_i(x, y_i^*) \approx f(x)$, then all the h_i 's and f 's from Eqn. (42) cancel out and we can write the *extended balance heuristic* in a form suitable for practical implementation:

$$\tilde{w}'_i(x, y_i^*) = \frac{n_i p'_i(x, y_i^*)}{\sum_{k=1}^n n_k p'_k(x, y_k^*)}, \quad (51)$$

where $p'_i(x, y_i^*) = p_i(x)$ for the unbiased path sampling techniques and $p'_i(x, y_i^*) = p_i(x, y_i^*) / \int K_i^2(x, y_i) dy_i$ for the techniques corresponding to the volumetric radiance estimators.

Under the aforementioned assumptions, y_i^* can be set to an arbitrary y_i within the kernel support. In practice, when an extended space technique i evaluates $p'_i(x, y_i^*)$, it uses the actual realization of the photon variable (e.g., the photon location). Unbiased techniques would use y_i^* corresponding to the kernel center (e.g., the query point in photon mapping). These are the same practical approximations as in previous work [Georgiev et al. 2012; Hachisuka et al. 2012].

If we use a *constant kernel* in the form $K_i(x, y_i) = |S_i(x)|^{-1}$, where $S_i(x)$ is the support of $K_i(x, y_i)$ for a given x , then we have

$$\int_{\mathcal{D}_{y_i}} K_i^2(x, y_i) dy_i = |S_i(x)|^{-2} \int_{S_i(x)} dy_i = |S_i(x)|^{-1}. \quad (52)$$

This in turn yields $p'_i(x, y_i^*) = p_i(x, y_i^*) |S_i(x)|$, which exactly corresponds to the vertex merging pdf of Georgiev et al. [2012].

Summary. Let us now summarize the practical outcomes of our derivations so far. Each estimator considered in the mixture yields a full family of sampling techniques corresponding to evaluating the estimator at different vertices along the path. Any technique in the combined estimator (40) weights its estimates by Eqn. (51). This involves calculating the extended pdfs $p'_i(x, y_i^*) = p_i(x, y_i^*) / \int K_i^2(x, y_i) dy_i$ for all the combined techniques, where the pdfs p_i for the volumetric estimators are given by Eqns. (45–48). The constant $\int K_i^2(x, y_i) dy_i$ can be precomputed. In practice, Eqn. (51) can be evaluated by taking pdf ratios [Veach 1997, Sec. 10.2], which are simple expressions because the pdfs differ by only a few terms. Note that the unweighted estimates are evaluated using the estimators in their original form, without any of the approximations used to derive the MIS weights.

In Fig. 9 we plot the weights calculated using the above formulas for selected estimators in the canonical setup described in Sec. 5. Comparison to the NSD plot in Fig. 6 reveals that high NSD corresponds

to low weight and vice-versa. This demonstrates the ability of our extended balance heuristic to promote low-variance estimators while lessening the contribution of the high-variance ones.

7.3 PDFs of Long and Short Beams

In the above derivations of the estimator pdfs, we have assumed the use of long beams (see Sec. 4.2). To derive the impact of short beams (i.e., the replacement of transmittance by a step function, dubbed “progressive deep shadow maps” by Jarosz et al. [2011b]), we use a new interpretation of short beams as a Russian roulette (RR) decision on the outcome of the long-beam estimator.

Consider an arbitrary pdf $p(t)$ used to sample the length of a short beam. The probability that the beam contributes at some distance t_0 from its origin is the probability that the beam length l is at least t_0 :

$$\Pr\{l > t_0\} = \int_{t_0}^{\infty} p(t') dt'. \quad (53)$$

Any long beam estimator can be converted to a corresponding short beam estimator by making a zero contribution whenever the beam is too short to make the considered contribution, else by the original estimator divided by the RR probability (53). For example, take the B-B1D estimator (17) and consider the use of short photon beams and long query beams. The resulting estimator reads:

$$\langle I \rangle_{B_s-B_11D} = \langle I \rangle_{B_1-B_11D} \frac{H(l_a - t_a)}{\Pr\{l_a > t_a\}}, \quad (54)$$

where l_a is the length of the photon beam, and the Heaviside step function H indicates whether or not the beam is long enough to make a contribution at t_a . If we further assume that the beam length is sampled from a pdf proportional to transmittance, $p(t) = \sigma_t(t)T_r(t)$ (as is usually the case), we get $\Pr\{l_a > t_a\} = T_r(t_a)$, and therefore

$$\langle I \rangle_{B_s-B_11D} = H(l_a - t_a) \rho(\tilde{\mathbf{b}}, \mathbf{b}) \frac{K_1(\tilde{\mathbf{b}}, \mathbf{b})}{\sin \theta_{ac}} T_r(t_c). \quad (55)$$

This is the original B_1 - B_1 estimator (17), with the transmittance $T_r(t_a)$ along the photon beam replaced by a step function [Jarosz et al. 2011b]. We could also apply the same procedure to the query beam.

By interpreting short beams as Russian roulette, it becomes clear that the pdf of a short beam estimator is given by the product of the pdf of the corresponding long beam estimator and the RR probability (53), usually given by the transmittance. This approach provides an intuitive explanation, as well as a generalization of, Jarosz et al.’s [2011b] unbiased progressive deep shadow maps. We could use any suitable pdf to sample the beam length, giving rise to the possibility of sampling “medium-length” beams. This could be an interesting avenue for future work.

8 Combined Rendering Algorithm

We now have at our disposal a formulation that allows us to fuse all of the volumetric estimators in Eqns. (9-17) and their short-beam variants, along with the unbiased path sampling techniques from BPT into a combined algorithm, which we call *unified points, beams, and paths (UPBP)*. Each volumetric estimator yields a whole family of sampling techniques corresponding to the eye subpath vertex where the estimator is evaluated. Formally, our algorithm evaluates the extended combined estimator (40), where the unbiased techniques are the various path sampling techniques in BPT, and the extended space techniques correspond to evaluating P-P, P-B, and B-B estimators at different path vertices or segments. We also include the family of techniques corresponding to the P-P surface

photon mapping estimator. Our combined algorithm therefore fully subsumes the methods of Georgiev et al. [2012] and Hachisuka et al. [2012].

The BPT sampling techniques are an important constituent of the combined algorithm because they excel at connecting subpaths over long distances (e.g., direct illumination in media far away from the light). Our algorithm currently does not use the joint path importance sampling techniques of Georgiev et al. [2013]. Including them would further reduce variance, especially in scenes with highly anisotropically scattering media for which none of the currently used estimators is efficient.

Practical considerations. Combining all 25 variations of the volumetric estimators (9-17) may not be productive since not all of them have complementary advantages that the combined algorithm could benefit from. For instance, P-B2D and P-B3D both have very similar pdfs and only differ by the amount of bias. For this reason, we choose to only use the minimum-blur volumetric estimators, i.e., P-P3D, P-B2D and B-B1D, as they introduce less bias. This choice is in line with previous work [Jarosz et al. 2011a].

Similarly, combining the long- and short-beam variants of the same estimators would not be useful because the long-beam variant always has less variance. On the other hand, evaluating the long-beam estimators is more costly, so a judicious choice needs to be made. In our tests, the use of short photon beams and long query beams provided the best performance for a given time budget, so our implementation uses the P-P3D, P-B₁2D and B_s-B₁1D estimator variants. Opting for short photon beams is important for the overall performance, because long photon beams have much higher cost of queries and construction of an acceleration data structure, and this overhead is rarely compensated by a corresponding variance reduction.

While using both the P-B2D and B-P2D estimators could further improve robustness, our implementation only uses P-B2D (photon points and query beams). Again, this choice is motivated by efficiency considerations: the P-B2D estimator can be efficiently implemented as ray casting on a set of photon point spheres.

Implementation. Our implementation builds on top of the publicly available SmallVCM framework [Davidovič and Georgiev 2012]. It extends the VCM algorithm [Georgiev et al. 2012] to handle participating media and produces an image by averaging the results of independent two-stage iterations.

In the first stage, we trace a number of light subpaths, connect their vertices to the eye (light tracing), and then store the vertices (photon points) and segments (photon beams). We build separate hashed grids over the surface and medium vertices, which are later used for the P-P estimators. We also build an additional bounding volume hierarchy (BVH) over the medium vertices for the P-B estimator. Photon beams are organized in a uniform grid, which we found to be faster than the BVH used by Jarosz et al. [2011b], though its query performance still leaves much to be desired. We store fewer photon beams than points, as the computational and memory cost per beam is much larger.

In the second stage of every iteration, we trace one eye subpath per pixel and construct a number of estimates as follows. Each vertex, surface or medium, is connected to a light source and to the vertices of a light subpath in order to evaluate the different unbiased estimators from BPT. Furthermore, at each eye subpath vertex we evaluate the P-P2D (surface) or the P-P3D (medium) estimator by looking up the photons from the corresponding grid. For each eye subpath segment passing through a medium, we evaluate the P-B₁2D and B_s-B₁1D estimators.




Figure 10: Equal-time comparison of our combined algorithm against previous work on a scene with a thin global medium and dense media enclosed in the spheres.

9 Results

We compare our UPBP algorithm to the individual volumetric estimators, Point-Point 3D (P-P3D), Point-Beam 2D (P-B₁2D), Beam-Beam 1D (B_s-B₁1D), and to bidirectional path tracing (BPT). Because the volumetric estimators are designed to capture only medium transport, we perform the comparison on a subset of light transport paths, described by the expression² $L(S|D|M)*MS^*E$, that all of those estimators can sample. To calculate this subset of transport, our algorithm is modified from the description in Sec. 8 in that while tracing an eye subpath, all non-specular surface contributions are disregarded until the path goes through the first medium interaction, from which point on the full combined algorithm is run. The implementation for P-P3D, P-B₁2D, and B_s-B₁1D estimators, designed to emulate the previous work [Jensen and Christensen 1998; Jarosz et al. 2008; Jarosz et al. 2011a], traces an eye subpath until the first non-specular interaction, and evaluates the respective estimators for the vertices or segments of this path. Our P-P3D estimator does *not* use ray marching as suggested by Jensen and Christensen [1998], so that the implementation is consistent with the derived theory. Ray marching would correspond to evaluating more P-P3D than P-B₁2D and B_s-B₁1D estimators (one evaluation per ray-marching step), which would complicate the interpretation of the results. We also include the comparison of our algorithm to BPT, which is modified to calculate only the subset of transport given by the above regular expression. All the tests were run on a Windows 7 PC with a 4-core Intel i7-2600K CPU and 16GB of RAM using 8 threads.

For our main results, we sample fewer light subpaths for generating photon beams than for generating photon points (see Table 1). We set the number of subpaths such that the total rendering time spent on evaluating the B-B estimator is about the same as the time spent on the other estimators. While this simple heuristic works well in our scenes, a more systematic analysis of estimator efficiency is an important avenue for future work.

To demonstrate the robustness of our method, we render three scenes containing media with a wide range of parameters (see Table 2), fea-

² L – light source, E – eye, S – purely specular surface interaction, D – diffuse/glossy surface interaction, M – medium interaction.




Figure 11: In this scene, our algorithm performs only slightly better than P-B₁2D due to the overhead induced by photon beams.

turing complex specular-medium-specular transport. Full-resolution images for all the scenes are provided in the auxiliary material.

The **Still life** scene in Fig. 1 features different kinds of media (from left to right: wax candle, glycerin soap bar on top of a block of a back-scattering medium, diluted wine, apple juice, and olive oil). Fig. 2 compares the result of our UPBP algorithm to the previous methods, implemented as described above. These results are in line with the observation we made in the canonical variance analysis, that beams (B_s-B₁1D) are not necessarily more efficient than the point-based estimators P-P3D and P-B₁2D. For instance, P-B₁2D renders a less noisy image than B_s-B₁1D for the dense wax, while the opposite is true for the thinner glycerin soap. The unbiased BPT techniques efficiently capture the light transport in the thin diluted wine and the apple juice. No single technique is able to efficiently render all the media in this scene, while our combined algorithm performs well.

The **Mirror balls** scene (Fig. 10) shows the good performance of photon beams at rendering (reflected) caustics in thin media. The number of photon beams per iteration is only 0.63% of the number of photons. Despite this significant difference in the numbers of samples, the variance of B_s-B₁1D in the thin medium that fills the space is very low. Although BPT is efficient at rendering volumetric caustics, their reflections are more efficiently captured by B_s-B₁1D. P-B₁2D and P-P3D produce good results in the two spheres with a dense medium. However, the variance of P-P3D in the thin medium is enormous, while P-B₁2D performs nearly as well as B_s-B₁1D. Our combined algorithm is able to take the best of all these estimators to produce a superior result.

The **Bathroom** scene in Fig. 11 has similar settings to the Mirror balls scene and the various estimators show similar performance. The B_s-B₁1D estimator excels at capturing the focused light around the complex light fixture, while having high variance in the thick media of the flasks contents and the washbasin. Our combined algorithm still has an edge over P-B₁2D, the best-performing previous method, though the advantage of the combined estimator is nearly offset by the overhead of evaluating many techniques.

While our implementation of the P-P3D and P-B₁2D estimators is fairly well optimized, the B-B1D implementation leaves much room for improvement. Speeding up B-B1D would strengthen the incentive for the combined algorithm, as the beams would more efficiently handle most effects in thin media. We however emphasize that, even if we could afford storing as many photon beams as photon points, the points would still remain an important component of the combined algorithm responsible for transport in dense media.




Figure 12: Contributions of the different estimators to the UPBP image of the Still life scene from Figs. 1 and 2. The individual images show a sum of a number of sampling techniques, corresponding to evaluating the respective estimators at all the medium vertices or segments on the eye path. Note that the images in the bottom rows have skewed relative brightness as they accumulate unweighted contributions, double-counting energy differently for different types of paths, and their exposure has been manually adjusted.

Table 1: Rendering settings and statistics.

	Still life	Mirror balls	Bathroom
Image resolution	1600 × 700	800 × 800	672 × 896
# light subpaths	1,120,000	640,000	602,112
# photon beam subpaths	23,000 (2.05%)	4,000 (0.63%)	3,000 (0.50%)
Maximum path length	80	12	20
Rendering time	25 min	60 min	60 min
# iterations OUR	637	575	531
# iterations BPT	1296	1697	1247
# iterations P-P3D	1809	2106	2846
# iterations B-P2D	1739	1690	2286
# iterations B-B1D	1711	1135	1416

Contribution of individual techniques. Fig. 12 compares the contributions of the individual volumetric estimators and BPT to our method, before and after the MIS weights are applied, for the Still Life scene from Figs. 1 and 2. We see that our MIS weights are effective at weighting down high-variance (noisy) contributions, while promoting lower-variance (smooth) ones. This empirically validates our MIS weight derivation in Sec. 6 and 7.

10 Discussion and Limitations

Variance analysis. Our canonical variance analysis in Sec. 5 only considers the source of variance that differs among the individual estimators, but disregards the variance from sampling the input configuration. This additional variance depends on the kernel width (the smaller the width, the lower the probability that the beams pass through the kernel) and diverges as $w \rightarrow 0$. Therefore, the result of zero variance for B_1 - B_1 in the restricted setup we consider is not in contradiction with Jarosz et al. [2011b], where the total variance of the B_1 - B_1 1D estimator is shown to be $O(1/w)$. Analysis of this extra variance for all the estimators is left for future work.

Bias and efficiency. Our variance analysis as well as the extended balance heuristic are only based on the estimators’ variance per sample taken. In practice, however, one is interested in the estimators’ total error (due to both variance and bias) per unit time. For example, while long photon beams always have lower variance than short beams, rendering with long beams may be less efficient due to the computational overhead. Likewise, the point and beam estimators with a large kernel produce smoother results (less variance) than BPT but the error may be higher because of bias. Despite some initial results [Hachisuka et al. 2012, supplemental document], analyzing these behaviors and designing practical bias- and efficiency-aware combination heuristics remains an important open problem.

Relative number of samples. We have shown that combining estimators yields an algorithm that is more robust to variations in media properties. However, for simple scenes that are well handled

by a single estimator, running the other ones only serves to incur an overhead. To resolve this issue, one would need to determine the appropriate number of samples to take from each estimator in a given scene. A solution based on Metropolis sampling has recently been proposed by Hachisuka et al. [2014]. It would be interesting to see how it performs with the estimators considered here.

Asymptotic analysis and consistency. Our algorithm uses a fixed kernel radius. A comprehensive asymptotic analysis of the trade-off between variance and bias, along the lines of the works of Knaus and Zwicker [2011], Jarosz et al. [2011b], and Georgiev et al. [2012], would yield the appropriate radius reduction scheme and make the combined algorithm consistent.

Heterogeneous media. While our implementation currently supports only homogeneous media, our theoretical framework supports heterogeneous media without any modifications.

11 Conclusion

We have proposed a way to combine within and across two classes of widely adopted volumetric light transport simulation approaches: those based on Monte Carlo integration and those based on photon density estimation using points and beams. Our variance analysis revealed that many of the point- and beam-based density estimators have complementary benefits dependent on the properties of the media being simulated. We also solidified the mathematical connection between these estimators and related ones used in neutron transport. In order to harness these benefits along with unbiased path sampling techniques, we developed an extended multiple importance sampling framework for combining estimators of integrals of different dimension and an extended balance heuristic which automatically weights such disparate estimators in order to minimize variance. We demonstrated the utility of this framework by combining several point- and beam-based density estimators together with unbiased bidirectional path tracing to form a practical algorithm for rendering participating media. Our resulting unified algorithm automatically exploits the complementary benefits of each technique to provide a combination more robust to scene variations and lighting scenarios than any single previous technique alone.

Acknowledgments. This work was supported by the Czech Science Foundation (grant P202-13-26189S) and Charles University in Prague (projects GA UK 1362413 and SVV-2014-260103). Derek Nowrouzezahrai acknowledges funding from the Natural Sciences and Engineering Research Council of Canada. Chaos Group kindly provided the Bathroom scene. Ondřej Karlík helped modeling the Still Life scene. We thank the anonymous reviewers for their suggestions on improving the paper.

Table 2: Medium parameters for the Still Life (top) and the Mirror Balls scenes (bottom). Media are listed in the order in which they appear in the respective scenes from left to right. The mean free path value of the media is given here as a multiple of the largest side of the enclosing object’s bounding box. The rightmost column gives the mean cosine of the Henyey-Greenstein phase function.

Medium	Scattering coefficient σ_s			Absorption coefficient σ_a			Single-scattering albedo			Mean free path	g
	R	G	B	R	G	B	R	G	B		
Wax candle	1.5000	1.5000	1.5000	0.0300	0.1000	0.2000	0.9803	0.9615	0.7500	0.032	0.8
Glycerin soap (top)	0.0201	0.0202	0.0221	0.0020	0.0040	0.0002	0.9090	0.8347	0.9910	2.300	0.6
Block (bottom)	0.0100	0.0100	0.0100	0.0100	0.0100	0.0100	0.5000	0.5000	0.5000	2.500	-0.9
Diluted wine	0.0150	0.0130	0.0111	0.1220	0.3510	0.4020	0.1094	0.0357	0.0268	0.365	0.9
Apple juice	0.0201	0.0243	0.0323	0.1014	0.1862	0.4084	0.1654	0.1154	0.0732	0.412	0.9
Olive oil	0.0410	0.0390	0.0120	0.0620	0.0470	0.3530	0.0042	0.4535	0.0995	0.581	0.9
Global medium (fog)	0.0002	0.0002	0.0001	0.0001	0.0001	0.0001	0.6667	0.600	0.5000	5.000	0.5
Green glass sphere	0.0004	0.0001	0.0004	0.0200	0.0010	0.0200	0.0196	0.0909	0.0196	12.62	0.0
Orange sphere	0.1000	0.1000	0.1000	0.0050	0.0600	0.2600	0.6250	0.9523	0.2777	0.043	-0.9
Dark amethyst	0.0600	0.1000	0.1000	0.0001	0.0100	0.0100	0.5964	0.0909	0.0909	0.100	-0.3

References

- ARVO, J., AND KIRK, D. 1990. Particle transport and image synthesis. In *Proc. of SIGGRAPH '90*, ACM, New York, NY.
- ARVO, J. 1993. Transfer functions in global illumination. In *ACM SIGGRAPH 93 Course Notes - Global Illumination*.
- BANERJEE, K. 2010. *Kernel Density Estimator Methods for Monte Carlo Radiation Transport*. PhD thesis, University of Michigan.
- DAVIDOVIČ, T., AND GEORGIEV, I., 2012. SmallVCM. <http://www.smallvcm.com/>.
- D’EON, E., AND IRVING, G. 2011. A quantized-diffusion model for rendering translucent materials. *ACM Trans. Graph. (Proc. of SIGGRAPH)* 30, 4.
- DUNN, K., AND WILSON, P. 2012. Kernel density estimators for Monte Carlo tallies on unstructured meshes. *Transactions of the American Nuclear Society* 107, 490–493.
- GEORGIEV, I., KŘIVÁNEK, J., DAVIDOVIČ, T., AND SLUSALLEK, P. 2012. Light transport simulation with vertex connection and merging. *ACM Trans. Graph. (Proc. of SIGGRAPH Asia)* 31, 6.
- GEORGIEV, I., KŘIVÁNEK, J., HACHISUKA, T., NOWROUZEZAHRAI, D., AND JAROSZ, W. 2013. Joint importance sampling of low-order volumetric scattering. *ACM Trans. Graph. (Proc. of SIGGRAPH Asia)* 32, 6 (Nov.).
- GKIOULEKAS, I., ZHAO, S., BALA, K., ZICKLER, T., AND LEVIN, A. 2013. Inverse volume rendering with material dictionaries. *ACM Trans. Graph.* 32, 6 (Nov.), 162:1–162:13.
- HACHISUKA, T., OGAKI, S., AND JENSEN, H. W. 2008. Progressive photon mapping. *ACM Trans. Graph. (Proc. of SIGGRAPH Asia)* 27, 5.
- HACHISUKA, T., PANTALEONI, J., AND JENSEN, H. W. 2012. A path space extension for robust light transport simulation. *ACM Trans. Graph. (Proc. of SIGGRAPH Asia)* 31, 6 (Nov.).
- HACHISUKA, T., KAPLANYAN, A. S., AND DACHSBACHER, C. 2014. Multiplexed Metropolis light transport. *ACM Trans. Graph. (Proc. of SIGGRAPH 2014)* 33, 4.
- IMMEL, D. S., COHEN, M. F., AND GREENBERG, D. P. 1986. A radiosity method for non-diffuse environments. *ACM SIGGRAPH Computer Graphics* 20, 4, 133–142.
- JAROSZ, W., ZWICKER, M., AND JENSEN, H. W. 2008. The beam radiance estimate for volumetric photon mapping. *Computer Graphics Forum (Proc. of Eurographics)* 27, 2.
- JAROSZ, W., NOWROUZEZAHRAI, D., SADEGHI, I., AND JENSEN, H. W. 2011. A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM Trans. Graph.* 30, 1, 5:1–5:19.
- JAROSZ, W., NOWROUZEZAHRAI, D., THOMAS, R., SLOAN, P.-P., AND ZWICKER, M. 2011. Progressive photon beams. *ACM Trans. Graph. (Proc. of SIGGRAPH Asia)* 30, 6.
- JENSEN, H. W., AND CHRISTENSEN, P. H. 1998. Efficient simulation of light transport in scenes with participating media using photon maps. In *Proc. of SIGGRAPH '98*.
- JENSEN, H. W. 1996. Global illumination using photon maps. In *Proc. of Eurographics Rendering Workshop*.
- KAJIYA, J. T. 1986. The rendering equation. In *Computer Graphics (Proc. of SIGGRAPH)*.
- KNAUS, C., AND ZWICKER, M. 2011. Progressive photon mapping: A probabilistic approach. *ACM Transactions on Graphics* 30, 3.
- LAFORTUNE, E. P., AND WILLEMS, Y. D. 1993. Bi-directional path tracing. In *Compugraphics '93*.
- LAFORTUNE, E. P., AND WILLEMS, Y. D. 1996. Rendering participating media with bidirectional path tracing. In *Proc. of the Eurographics Workshop on Rendering*.
- MACMILLAN, D. B. 1966. Comparison of statistical estimators for neutron Monte Carlo calculations. *Nuclear Science and Engineering* 26, 3 (Nov.), 366–372.
- NOVÁK, J., NOWROUZEZAHRAI, D., DACHSBACHER, C., AND JAROSZ, W. 2012. Progressive virtual beam lights. *Proc. of Eurographics Symposium on Rendering* 31, 4 (June).
- NOVÁK, J., NOWROUZEZAHRAI, D., DACHSBACHER, C., AND JAROSZ, W. 2012. Virtual ray lights for rendering scenes with participating media. *ACM Trans. Graph. (Proc. of SIGGRAPH)* 31, 4 (July).
- PAULY, M., KOLLIG, T., AND KELLER, A. 2000. Metropolis light transport for participating media. In *Rendering Techniques (Proc. of Eurographics Workshop on Rendering)*.
- SPANIER, J., AND GELBARD, E. M. 1969. *Monte Carlo principles and neutron transport problems*. Addison-Wesley.
- SPANIER, J. 1966. Two pairs of families of estimators for transport problems. *SIAM Journal on Applied Mathematics* 14, 4, 702–713.
- VEACH, E., AND GUIBAS, L. 1994. Bidirectional estimators for light transport. In *Proc. of Eurographics Rendering Workshop*.
- VEACH, E., AND GUIBAS, L. J. 1995. Optimally combining sampling techniques for Monte Carlo rendering. In *Proc. of SIGGRAPH '95*.
- VEACH, E. 1997. *Robust Monte Carlo methods for light transport simulation*. PhD thesis, Stanford, CA, USA.
- VORBA, J. 2011. Bidirectional photon mapping. In *Proc. of the Central European Seminar on Computer Graphics (CESCG '11)*.