Beyond Points and Beams: Higher-Dimensional Photon Samples for Volumetric Light Transport

BENEDIKT BITTERLI and WOJCIECH JAROSZ, Dartmouth College

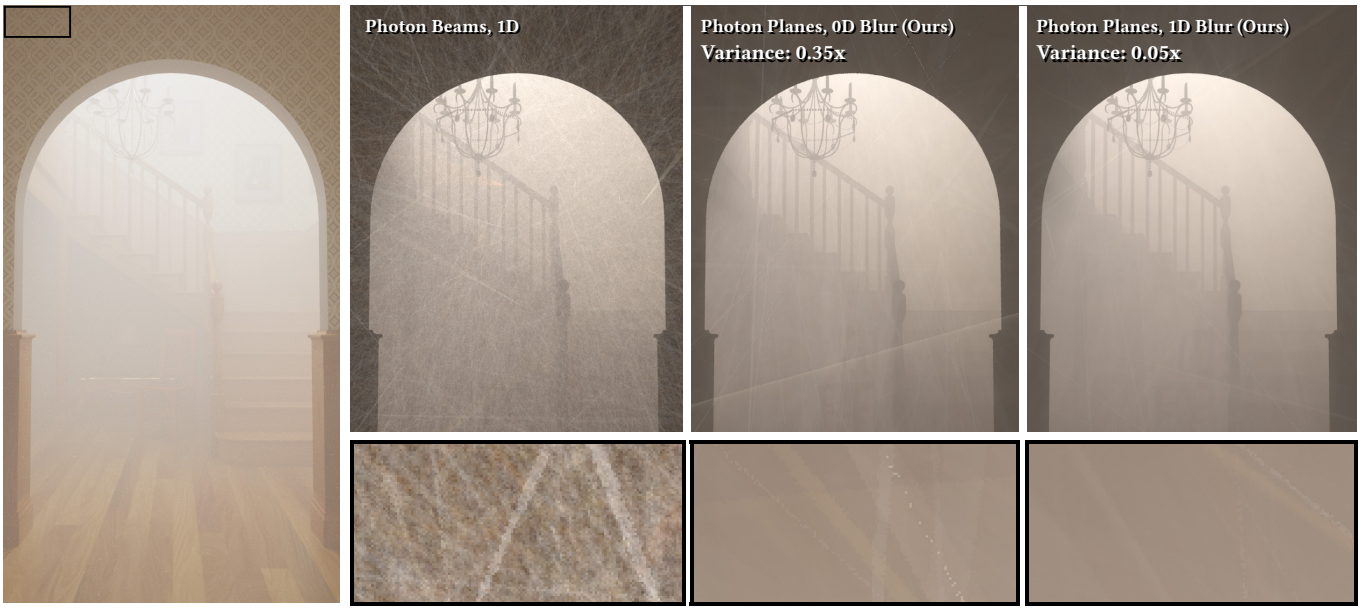

Fig. 1. We compare two of the estimators predicted by our theory (photon planes, 0D blur and 1D blur) to photon beams on the Doorway scene. We show the full light transport (left image) and estimates of the multiply scattered volumetric transport obtained from each of the three techniques (right images, insets taken from black rectangle). At equal render time, both of our estimators provide substantial image quality improvement and variance decrease.

We develop a theory of volumetric density estimation which generalizes prior photon point (0D) and beam (1D) approaches to a broader class of estimators using "nD" samples along photon and/or camera subpaths. Volumetric photon mapping performs density estimation by point sampling propagation distances within the medium and performing density estimation over the generated points (0D). Beam-based (1D) approaches consider the expected value of this distance sampling process along the last camera and/or light subpath segments. Our theory shows how to replace propagation distance sampling steps across *multiple bounces* to form higher-dimensional samples such as photon planes (2D), photon volumes (3D), their camera path equivalents, and beyond. We perform a theoretical error analysis which reveals that in scenarios where beams already outperform points, each additional dimension of nD samples compounds these benefits further. Moreover, each additional sample dimension reduces the required dimensionality of the blurring needed for density estimation, allowing us to formulate, for the first time, *fully unbiased* forms of volumetric photon mapping. We demonstrate practical implementations of several of the new estimators our theory predicts, including both biased and unbiased variants, and show that they outperform state-of-the-art beam-based volumetric photon mapping by a factor of 2.4–40×.

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

Additional Key Words and Phrases: global illumination, light transport, participating media, photon mapping, photon beams

ACM Reference format:
Benedikt Bitterli and Wojciech Jarosz. 2017. Beyond Points and Beams: Higher-Dimensional Photon Samples for Volumetric Light Transport. *ACM Trans. Graph.* 36, 4, Article 112 (July 2017), 12 pages.
DOI: http://dx.doi.org/10.1145/3072959.3073698

1 INTRODUCTION

Light scattering in participating media is responsible for the appearance of many materials (such as clouds, fog, juices, or our own skin). Unfortunately, simulating these effects accurately and efficiently remains a notoriously difficult problem since it requires solving not only the rendering equation [Immel et al. 1986; Kajiya 1986], but also its volumetric generalization, the radiative transfer equation [Chandrasekhar 1960].

Due to decades of research, practitioners now have an arsenal of competing and complementary rendering techniques to tackle this problem. Path tracing [Lafortune and Willems 1993; Veach and Guibas 1994, 1997] and its volumetric variants [Georgiev et al. 2013; Lafortune and Willems 1996] have become one of the dominant forms of Monte Carlo (MC) light transport simulation in graphics due to their simplicity, generality, and ability to produce unbiased results where the only error is noise. Photon density estimation

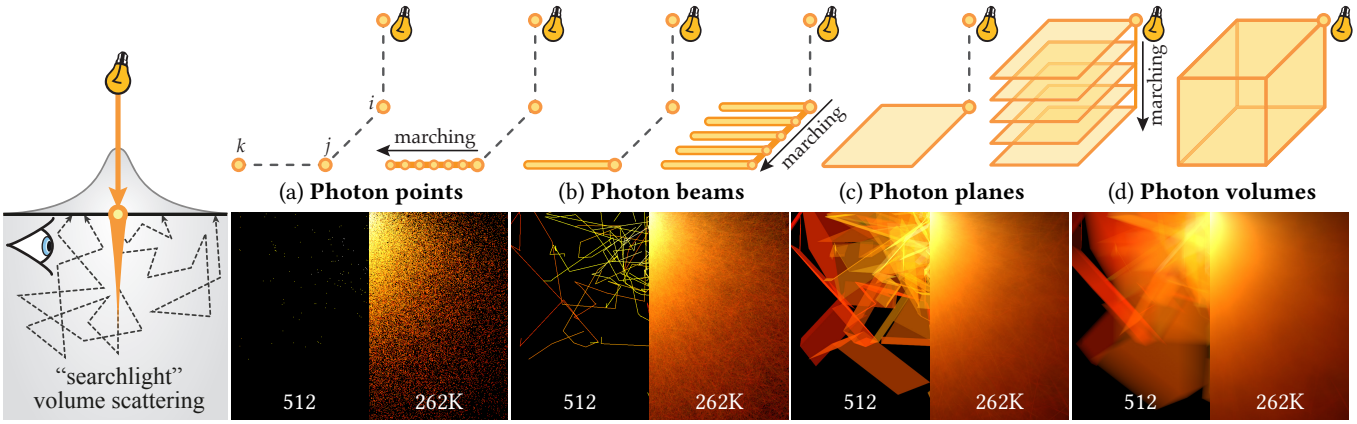

Fig. 2. We generalize 0D photon points (a) and 1D beams (b) to produce progressively higher-dimensional nD samples such as 2D "photon planes" (c) and 3D "photon volumes" (d). We form these estimators by computing the limit process of "marching" or sweeping photons along preceding light path segments, which allows us to progressively reduce variance and bias. The motivational experiment in the bottom row uses these successive estimators (each shown vertically split at two sample counts) on a searchlight problem setup (left), confirming that higher-order nD samples have the potential to dramatically improve quality in volumetric light transport.

approaches [Jensen 2001; Jensen and Christensen 1998], can robustly simulate light paths that are difficult for path tracing, though at the cost of introducing bias due to blurring. Recent work has combined path integration and photon density estimation [Georgiev et al. 2012; Hachisuka et al. 2012] to retain each one's complementary benefits.

Despite their increasing sophistication, all of the aforementioned methods can be viewed in Veach's path integral formulation [Pauly et al. 2000; Veach 1997] as MC integration using *point sampling*. In this view, the concatenated coordinates of a light path's vertices correspond to a single high-dimensional point, where the integrand expresses the amount of light transported along the path.

Recently, researchers have argued that such point sampling is an unnecessary legacy from surface rendering, leading to the development of line or "beam" samples in the volumetric rendering context. Jarosz et al. [2008] proposed (1D) line samples along camera rays, and later [2011a] formed a generalized theory of volumetric density estimators using lines along camera rays, photon path segments, or both. This offered substantial variance reduction, but also bias reduction since such beams lower the required blur dimensionality. Progressive formulations [Jarosz et al. 2011b] and combinations [Křivánek et al. 2014] with volumetric path tracing have helped to further reduce bias and improve robustness. Unfortunately, despite these improvements, variance continues to be a problem, and since some amount of blur is still required to employ such beam estimators, bias remains.

In this paper we show that existing estimators (Sec. 3) using points (0D) and beams (1D) are just special cases of a more general theory (Sec. 4) of volumetric light transport simulation using higher-dimensional (nD) samples. Our theory operates by successively replacing propagation distance point samples with lines across *multiple bounces* to produce 2D (plane/quad) samples, 3D (volume/parallelepiped) samples, their camera path equivalents, and beyond. Fig. 2 (top) illustrates this process, and (bottom) demonstrates the potential of using the corresponding estimators for a simple volumetric lighting simulation. We perform a theoretical error analysis (Sec. 5) which reveals that in scenarios where 1D beams provide error reduction compared to 0D points, we should expect higher-dimensional samples to provide asymptotically greater improvements. Moreover, higher-dimensional samples provide greater flexibility to reduce the dimensionality of density estimation blur, allowing us to formulate completely *unbiased volumetric density estimators*. Fig. 1 shows that photon planes allow for a 3× reduction in variance compared to photon beams while simultaneously producing an *unbiased* result, or a 20× improvement when allowing for the same amount of bias as photon beams. We discuss how to develop these estimators into practical rendering algorithms (Sec. 6) and show how they outperform photon beams by a factor of 2.4–40× for a wide range of scenes (Sec. 7). Since our theory predicts an infinite set of new estimators, we discuss the ample opportunities for future work in Sec. 8.

2 PREVIOUS WORK

We focus primarily on prior work for participating media rendering and related problems which leverage higher-dimensional samples. Refer to Pharr et al. [2016] for a more extensive and general survey.

Lines/beams in photon mapping. Havran et al. [2005] were the first to propose using photon path segments, and not vertices, to perform density estimation, in an effort to reduce boundary bias for surface illumination. Jarosz and colleagues [2008] initially proposed using beams to find photons along entire camera rays in volumes and later [2011a] generalized this idea to a collection of estimators that allow using beams along either the camera ray or light ray, or both. Sun et al. [2010] concurrently proposed a method for caustics and single scattering that corresponds to one of these estimators. Jarosz et al. [2011b] later derived beam estimators akin to progressive photon mapping [Hachisuka and Jensen 2009; Hachisuka et al. 2008; Knaus and Zwicker 2011] and introduced an unbiased "short" form of beam to avoid expensive and biased heterogeneous transmittance evaluations. All of these volumetric approaches can be viewed as

using line samples to collapse one or two dimensions of the volume rendering equation. Our theory generalizes this concept by collapsing an arbitrary number of propagation distance dimensions.

Lines/beams in MC path sampling. Novák and collaborators [2012a; 2012b] applied the insights from photon beams to virtual ray/beam lights to provably reduce the singularities that often plague many-light rendering techniques [Dachsbacher et al. 2014]. Georgiev et al. [2013] extended these ideas by reformulating volumetric path tracing and bidirectional path tracing to operate on path segments instead of vertices, reducing variance and singularities. Křivánek et al. [2014] combined point and beam density estimators with bidirectional path tracing using MIS [Veach and Guibas 1995]. All of these approaches resorted to MC point sampling to leverage beams in an unbiased context. We show that the extra degrees of freedom in volumetric light transport allow devising higher-dimensional density estimators that can avoid point sampling while obtaining unbiased results. Moreover, our theory could have potential future applications for, or could be combined with MC path sampling techniques, much as these prior methods have done for beams.

Lines and nD samples for other rendering problems. Several researchers have proposed MC-like approaches using 1D samples to collapse other dimensions of the rendering integral. Applications have included anti-aliasing [Jones and Perry 2000], hemispherical visibility and motion blur [Gribel et al. 2011, 2010], depth of field [Tzeng et al. 2012], hair rendering [Barringer et al. 2012], soft shadows [Billen and Dutré 2016], and masked environment lighting [Nowrouzezahrai et al. 2014]. Since evaluating nD samples effectively involves computing an nD integral, harnessing analytic solutions to rendering sub-problems can be viewed as a form of line or nD sampling, including semi-analytic solutions to the so-called airlight integral in homogeneous participating media [Pegoraro and Parker 2009; Sun et al. 2005], seminal work on closed-form area lighting on diffuse and glossy surfaces [Arvo 1995a,b; Chen and Arvo 2001, 2000], or analytic solutions to form factors [Schröder and Hanrahan 1993]. We focus on the volumetric problem, devising the theory that allows collapsing multiple dimensions of the light transport integral.

Neutron transport. Graphics has a long history of importing techniques from related fields such as neutron transport [Arvo 1993; Arvo and Kirk 1990; D'Eon and Irving 2011]. Křivánek et al. [2014] recently established a firm mathematical link between photon points [Jensen and Christensen 1998], short beams [Jarosz et al. 2011b] and long beams [Jarosz et al. 2011a], and "collision", "track-length", and "expected valued" transmittance estimators developed in neutron transport some fifty years ago [Lux 1978; Spanier 1966; Spanier and Gelbard 1969] (we briefly review these in Sec. 3). Our theory generalizes track-length and expected value estimators to evaluate transmittance across *multiple bounces*, and could have potential applications in related fields.

3 BACKGROUND

We begin by defining our notation and briefly reviewing volumetric light transport. We then summarize existing volumetric estimators most relevant to our work using our notation.

3.1 Light Transport in Participating Media

Light transport can be most generally described using the path integral framework [Veach 1997], where the intensity of some pixel measurement I is defined as an integral, $I = \int_\Omega f(\bar{\mathbf{z}})\mathrm{d}\mu(\bar{\mathbf{z}})$, over the space Ω of light transport paths $\bar{\mathbf{z}}$ with path throughput $f(\bar{\mathbf{z}})$.

The path integral can be approximated by an MC estimator $I \approx \frac{1}{m}\sum_{i=1}^{m}\frac{f(\bar{\mathbf{z}}_i)}{p(\bar{\mathbf{z}}_i)}$, which averages the contribution of m transport paths sampled with probability density $p(\bar{\mathbf{z}})$.

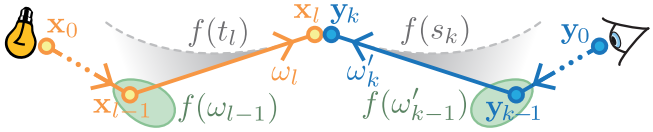

In photon mapping, complete transport paths $\bar{\mathbf{z}} = \bar{\mathbf{x}}_l\bar{\mathbf{y}}_k$ are decomposed into a photon subpath $\bar{\mathbf{x}}_l$ and a camera subpath $\bar{\mathbf{y}}_k$, where density estimation is performed at the vertex pair $\mathbf{x}_l\mathbf{y}_k$.

Photon subpath. The photon subpath $\bar{\mathbf{x}}_l = \mathbf{x}_0 \ldots \mathbf{x}_l$ is a length-l path with $l \geq 1$ segments and $l+1$ vertices. The first vertex $\mathbf{x}_0$ resides on the light, $\mathbf{x}_1 \ldots \mathbf{x}_{l-1}$ are scattering vertices on surfaces or in a medium, and the last vertex $\mathbf{x}_l$ will be connected to a camera subpath using density estimation.

We assume photon paths are constructed starting at the light by first sampling $\mathbf{x}_0$ with pdf $p(\mathbf{x}_0)$ and then sampling a sequence of directions and distances, which we denote $\bar{\omega} = \omega_1 \ldots \omega_l$ and $\bar{t} = t_1 \ldots t_l$ respectively. The i-th vertex is therefore $\mathbf{x}_i = \mathbf{x}_{i-1} + t_i\omega_i$, where t_i and ω_i denote the distance and direction *arriving* at vertex $\mathbf{x}_i$, not leaving it. The subpath pdf can therefore be decomposed as:

$$p(\bar{\mathbf{x}}_l) = p(\mathbf{x}_0)\,p(\bar{\omega}_l)\,p(\bar{t}_l), \text{ with:} \tag{1}$$

$$p(\bar{\omega}_l) = \prod_{i=1}^{l} p(\omega_i), \qquad \text{and} \qquad p(\bar{t}_l) = \prod_{i=1}^{l} p(t_i). \tag{2}$$

The throughput along a photon subpath $f(\bar{\mathbf{x}}_l)$ is the product of the light's emission L_e and the directional, $f(\bar{\omega}_l)$, and distance, $f(\bar{t}_l)$, throughput terms along each path vertex and segment, respectively:

$$f(\bar{\mathbf{x}}_l) = L_\mathrm{e}(\mathbf{x}_0, \omega_1)\,f(\bar{\omega}_l)\,f(\bar{t}_l), \text{ with:} \tag{3}$$

$$f(\bar{\omega}_l) = \prod_{i=1}^{l} f(\omega_i), \qquad \text{and} \qquad f(\bar{t}_l) = \prod_{i=1}^{l} f(t_i). \tag{4}$$

At each vertex, the directional throughput is a generalized scattering term that accounts for the phase function or BRDF:

$$f(\omega_i) = \begin{cases} \rho_\mathrm{s}(\omega_i \to \omega_{i+1}) & \text{if } \mathbf{x}_i \text{ is on a surface, and} \\ \rho_\mathrm{p}(\omega_i \to \omega_{i+1})\sigma_s & \text{if } \mathbf{x}_i \text{ is in a medium,} \end{cases} \tag{5}$$

where σ_s is the scattering coefficient.

Additionally, the propagation distance throughput accounts for binary visibility V and transmittance T_r along the i-th segment:

$$f(t_i) = V(t_i)\,T_\mathrm{r}(t_i), \text{ with } T_\mathrm{r}(t) = \mathrm{e}^{-t\sigma_t}, \tag{6}$$

with extinction coefficient σ_t, assuming homogeneous media.

We also define a shorthand for the photon subpath contribution, or *weight* up to vertex l:

$$C(\bar{\omega}_l) = \frac{f(\bar{\omega}_l)}{p(\bar{\omega}_l)}, \text{ and } C(\bar{t}_l) = \frac{f(\bar{t}_l)}{p(\bar{t}_l)}. \tag{7}$$

Camera subpath. The photon subpath $\bar{\mathbf{x}}_l$ is connected to a camera subpath $\bar{\mathbf{y}}_k = \mathbf{y}_k \dots \mathbf{y}_0$ which is traced from the opposite direction, starting at the camera located at $\mathbf{y}_0$. All the previously defined functions for throughput, pdf, and contribution are defined analogously for camera subpaths, with the only exception that the sensor importance W_e replaces the emitted radiance L_e in the subpath throughput Eq. (3). We use s_i and ω_i' to denote the distance and direction arriving at camera subpath vertex $\mathbf{y}_i$, to distinguish from the t_i and ω_i notation used for light subpaths.

3.2 Density Estimators

We define density estimators in terms of approximations to the complete path throughput,

$$\frac{f(\bar{\mathbf{z}})}{p(\bar{\mathbf{z}})} \approx C(\bar{\omega}_l) C(\bar{t}_{l-1}) \langle D \rangle^{l,k} C(\bar{s}_{k-1}) C(\overline{\omega'}_k), \tag{8}$$

where $\langle D \rangle^{l,k}$ is the density estimator at the vertex pair $\mathbf{x}_l \mathbf{y}_k$.

Photon point–sensor point (0D×0D, 3D blur): The original volumetric density estimator connects a photon point to a sensor point using a 3D blur [Jensen and Christensen 1998]:

$$\langle D \rangle^{l,k}_{\text{P-P3D}} = \underbrace{\frac{f(t_l)}{p(t_l)}}_{\text{photon sampling}} \left\{ K_3(\mathbf{x}_l, \mathbf{y}_k) f_\omega^{l,k} \right\} \underbrace{\frac{f(s_k)}{p(s_k)}}_{\text{query point sampling}}, \tag{9}$$

where K_3 evaluates a 3D blur kernel at the vertices $\mathbf{x}_l \mathbf{y}_k$, and $f_\omega^{l,k} = \rho_\mathrm{p}(\omega_l \to -\omega_k') \sigma_s$ evaluates the scattering function at the connection point using the last photon and camera subpath directions.

Photon point–sensor beam (0D×1D, 2D blur): Jarosz et al. [2008] later extended this estimator to use sensor rays instead of sensor points by replacing the last distance sampling step on the camera subpath $f(s_k)/p(s_k)$ in Eq. (9) with an expected value estimator while simultaneously reducing the kernel to 2D:

$$\langle D \rangle^{l,k}_{\text{P-B2D}} = \frac{f(t_l)}{p(t_l)} \left\{ K_2(\mathbf{x}_l, \mathbf{y}_k) f_\omega^{l,k} \right\} \underbrace{f(s_k^*)}_{\text{expected value}}, \tag{10}$$

where s_k^* is the camera ray × photon disc intersection distance.

Photon beam–sensor beam (1D×1D, 2D blur): Photon beams result from replacing the last distance sampling step on the light subpath $f(t_l)/p(t_l)$ with an expected value estimator, in a fashion similar to sensor beams. If we start with Eq. (9), we obtain [Jarosz et al. 2011a]:

$$\langle D \rangle^{l,k}_{\text{B-B2D}} = \int_{s_{k-}}^{s_{k+}} \underbrace{f(\tilde{t}_l)}_{\text{expected value}} \left\{ K_2(\mathbf{x}_l, \tilde{\mathbf{y}}_k) f_\omega^{l,k} \right\} f(s)\, \mathrm{d}s, \tag{11}$$

where the integral bounds are defined by the intersection between the camera ray and the photon beam, and variables $\tilde{t}_l$, $\tilde{\mathbf{y}}_k$ are functions of the integration variable s.

Photon beam–sensor beam (1D×1D, 1D blur): Similarly, if Eq. (10) is used, we obtain

$$\langle D \rangle^{l,k}_{\text{B-B1D}} = f(t_l^*) \left\{ \frac{K_1(\mathbf{x}_l, \mathbf{y}_k)}{J^{l,k}_{\text{B-B1D}}} f_\omega^{l,k} \right\} f(s_k^*), \tag{12}$$

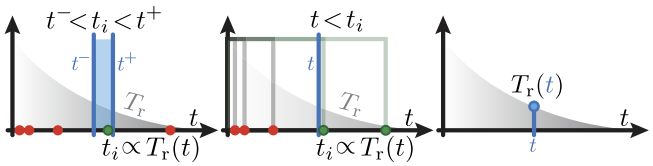

(a) collision (b) track-length (c) expected-value

Fig. 3. Collision estimators (a) check if the sampled t_i distances fall within an interval centered at t; track-length estimators (b) check if the t_i fall past distance t; and expected value estimators evaluate transmittance directly.

where the integral is replaced by the Jacobian for 1D×1D coupling with 1D blur, $J^{l,k}_{\text{B-B1D}} = \|\omega_l \times \omega_k'\|$, and t_l^*, s_k^* are the distances to the beam-beam intersection along the photon and camera segments.

"Long" and "short" beams. The original volumetric photon mapping approach (Eq. (9)) corresponds to estimating transmittance with a "collision" estimator (Fig. 3 (a)) from neutron transport [Spanier and Gelbard 1969], where a photon scores a constant contribution if it falls within a finite interval (the blur kernel). The subsequent beam estimators replace a distance sampling step along the camera and/or light path segment with a direct evaluation of transmittance. This corresponds to an "expected value" estimator (Fig. 3 (c)), so called because it returns the expected value of the propagation sampling process (the transmittance) directly [Spanier 1966]. These estimators result in "long" beams with infinite extent, which reduce variance, but also have a few undesirable properties: Their infinite extent makes them challenging to incorporate into acceleration structures, and the appearance of a factor $f(t_l)$ on its own requires the evaluation of a transmittance term at every estimation point, which can be costly in heterogeneous media.

Jarosz et al. [2011b] replaced this transmittance term with an unbiased estimate $\langle f(t_l) \rangle_T$, which is equivalent to a constant step function that drops to zero after a randomly sampled propagation distance. Such a transmittance estimator results in a "short" beam of finite extent and constant transmittance. This corresponds to estimating transmittance with a "track-length" estimator [Spanier 1966] from neutron transport (Fig. 3 (b)), which scores a constant value if the sampled distance lies beyond the estimation point.

In the next section, we will work with the expected value (long) forms of these estimators, as this will simplify our derivations, but we will subsequently replace the directly evaluated transmittance with its unbiased track-length estimate to form the short versions.

4 RADIANCE ESTIMATION WITH nD PHOTONS

We now generalize the concept of point and beam estimators to arbitrary dimensions, which is one of our primary contributions. Our key insight is we can convert additional propagation distance steps to their expected values by unfolding the photon throughput prefix $C(\bar{t}_l)$. We describe our approach by deriving novel two- and three-dimensional density estimators in detail; the same method can then be applied repeatedly to even higher dimensions.

Photon plane–sensor beam (2D×1D, 1D blur): We begin by inserting the B-B2D density estimator Eq. (11) into Eq. (8) to obtain

$$\frac{f(\overline{\mathbf{z}})}{p(\overline{\mathbf{z}})} \approx C(\overline{\omega}_l)C(\overline{t}_{l-1})\langle D\rangle^{l,k}_{\text{B-B2D}}C(\overline{s}_{k-1})C(\overline{\omega'}_k) \tag{13}$$
$$= C(\overline{\omega}_l)C(\overline{t}_{l-2})\left\{\frac{f(t_{l-1})}{p(t_{l-1})}\langle D\rangle^{l,k}_{\text{B-B2D}}\right\}C(\overline{s}_{k-1})C(\overline{\omega'}_k).$$

The last step was achieved by assuming $l \geq 2$ and expanding $C(\overline{t}_{l-1})$ by one term. We will name the quantity inside the braces $\langle D\rangle^{l-1,k}_{\text{B-B2D}}$, which is a B-B2D estimator that performs one additional distance sampling step. Expanding this quantity yields

$$\langle D\rangle^{l-1,k}_{\text{B-B2D}} = \frac{f(t_{l-1})}{p(t_{l-1})}\int_{s_{k-}}^{s_{k+}} f(\tilde{t}_l)\left\{K_2(\mathbf{x}_l,\tilde{\mathbf{y}}_k)f^{l,k}_\omega\right\}f(s)\,\mathrm{d}s. \tag{14}$$

The first term on the right-hand side is the result of distance sampling, which is used to obtain t_{l-1}. We now replace this distance sampling step with a deterministic "beam marching" procedure (right). Instead of sampling the location of a single beam, we place a series of beams at regular intervals along the ray $\mathbf{x}_{l-2} + \omega_{l-1}t^{(i)}_{l-1}$. We set the ray offset of each beam to $t^{(i)}_{l-1} = i\Delta t$, where Δt is the step size.

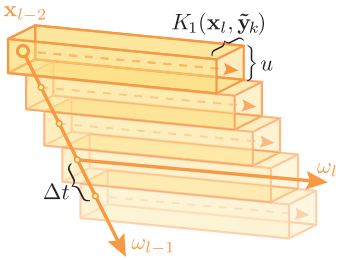

We select a blurring kernel which is uniform along one dimension, $K_2(\mathbf{x}_l,\mathbf{y}_k) = u^{-1}K_1(\mathbf{x}_l,\mathbf{y}_k)$, where u defines the uniform blur extent, and the direction of the uniform blurring is as in the figure above. The contribution of this estimator then becomes a sum,

$$\sum_{i=0} f(t^{(i)}_{l-1})\Delta t\int_{s^{(i)}_{k-}}^{s^{(i)}_{k+}} f(\tilde{t}_l)\left\{\frac{K_1(\mathbf{x}_l,\tilde{\mathbf{y}}_k)}{u}f^{l,k}_\omega\right\}f(s)\,\mathrm{d}s. \tag{15}$$

Because of the deterministic marching procedure, the inverse sampling density $p(t_{l-1})^{-1}$ becomes Δt. We now choose the uniform blur extent such that kernels of adjacent beams touch exactly, making $s^{(i)}_{k+} = s^{(i+1)}_{k-}$. This is achieved with $u = \Delta t\|\omega_{l-1}\times\omega_l\|$. Substituting into Eq. (15) and rearranging yields

$$\sum_{i=0}\int_{s^{(i)}_{k-}}^{s^{(i+1)}_{k-}} f(t^{(i)}_{l-1})f(\tilde{t}_l)\Delta t\left\{\frac{K_1(\mathbf{x}_l,\tilde{\mathbf{y}}_k)}{\Delta t J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_\omega\right\}f(s)\,\mathrm{d}s, \tag{16}$$

with $J^{l-1,l}_{\text{Q-B1D}} = \|\omega_{l-1}\times\omega_l\|$. The constant Δt can be moved into the braces and cancels. Taking the limit as $\Delta t \to 0$ merges the beams into a continuous *photon plane* with contribution

$$\langle D\rangle^{l-1,k}_{\text{Q-B1D}} = \int_{s_{k-}}^{s_{k+}} f(\tilde{t}_{l-1})f(\tilde{t}_l)\left\{\frac{K_1(\mathbf{x}_l,\tilde{\mathbf{y}}_k)}{J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_\omega\right\}f(s)\,\mathrm{d}s. \tag{17}$$

Photon plane–sensor beam (2D×1D, 0D blur): In a similar fashion, we now insert the B-B1D estimator (Eq. (12)) into Eq. (8) and expand the distance throughput term to obtain the quantity

$$\langle D\rangle^{l-1,k}_{\text{B-B1D}} = \frac{f(t_{l-1})}{p(t_{l-1})}f(t^*_l)\left\{\frac{K_1(\mathbf{x}_l,\mathbf{y}_k)}{J^{l,k}_{\text{B-B1D}}}f^{l,k}_\omega\right\}f(s^*_k). \tag{18}$$

Again, we replace distance sampling along t_{l-1} with a deterministic beam marching procedure. We choose a uniform blurring kernel $K_1(\mathbf{x}_l,\mathbf{y}_k) = u^{-1}$ with blur extent u. The direction of the blur $\vec{u} = (\omega_l\times\omega'_k)/\|\omega_l\times\omega'_k\|$ is oriented orthogonal to the last photon and camera subpath directions (see figure below).

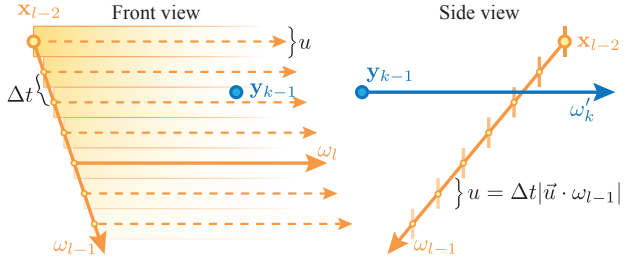

The contribution then becomes

$$\sum_{i=0} f(t^{(i)}_{l-1})\Delta t f(t^{*(i)}_l)\left\{\frac{K_1(\mathbf{x}_l,\mathbf{y}_k)}{J^{l,k}_{\text{B-B1D}}}f^{l,k}_\omega\right\}f(s^{*(i)}_k). \tag{19}$$

We choose u such that kernels of adjacent beams touch exactly when viewed from ω'_k. This can be achieved by projecting the spacing between beams onto the blur direction, yielding $u = \Delta t|\vec{u}\cdot\omega_{l-1}|$. Since only one kernel overlaps the camera ray, the summation disappears

$$f(t^*_{l-1})f(t^*_l)\Delta t\left\{\frac{f^{l,k}_\omega}{\Delta t\,|\vec{u}\cdot\omega_{l-1}|\,J^{l,k}_{\text{B-B1D}}}\right\}f(s^*_k). \tag{20}$$

The constant Δt can be moved into the braces and cancels. Additionally, the term $\|\omega_l\times\omega'_k\|$ occurs both in $J^{l,k}_{\text{B-B1D}}$ and the denominator of $\vec{u}$, and can be cancelled. Taking the limit and simplifying yields

$$\langle D\rangle^{l-1,k}_{\text{Q-B0D}} = f(t^*_{l-1})f(t^*_l)\left\{\frac{f^{l,k}_\omega}{J^{l-1,l,k}_{\text{Q-B0D}}}\right\}f(s^*_k), \tag{21}$$

where $J^{l-1,l,k}_{\text{Q-B0D}} = \left|\omega_{l-1}\cdot(\omega_l\times\omega'_k)\right|$ is the Jacobian for 2D×1D coupling with 0D blur, yielding a continuous photon plane.

Photon volume–sensor beam (3D×1D, 0D blur): We insert and expand the Q-B1D estimator (17) into Eq. (8) to obtain $\langle D\rangle^{l-2,k}_{\text{Q-B1D}}$:

$$\frac{f(t_{l-2})}{p(t_{l-2})}\int_{s_{k-}}^{s_{k+}} f(\tilde{t}_{l-1})f(\tilde{t}_l)\left\{\frac{K_1(\mathbf{x}_l,\tilde{\mathbf{y}}_k)}{J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_\omega\right\}f(s)\,\mathrm{d}s. \tag{22}$$

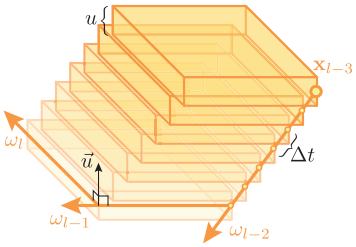

We replace distance sampling along t_{l-2} with deterministic "plane marching" (left) and select a uniform blurring kernel $K_1(\mathbf{x}_l,\mathbf{y}_k) = u^{-1}$ with blur direction $\vec{u} = (\omega_{l-1}\times\omega_l)/\|\omega_{l-1}\times\omega_l\|$ normal to the plane.

The contribution from all planes is

$$\sum_{i=0} f(t^{(i)}_{l-2})\Delta t\int_{s^{(i)}_{k-}}^{s^{(i)}_{k+}} f(\tilde{t}_{l-1})f(\tilde{t}_l)\left\{\frac{u^{-1}}{J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_\omega\right\}f(s)\,\mathrm{d}s. \tag{23}$$

To ensure that adjacent planes touch exactly, we project the plane spacing onto the blur direction to obtain $u = \Delta t\,|\omega_{l-2} \cdot \vec{u}|$. Substituting into Eq. (23) and rearranging yields

$$\sum_{i=0} \int_{s_{k-}^{(i)}}^{s_{k-}^{(i+1)}} f(t_{l-2}^{(i)}) f(\tilde{t}_{l-1}) f(\tilde{t}_l) \Delta t \left\{ \frac{|\omega_{l-2} \cdot \vec{u}|^{-1}}{\Delta t\, J_{\text{Q-B1D}}^{l-1,l}} f_\omega^{l,k} \right\} f(s)\,\mathrm{d}s.$$

The term $\|\omega_{l-1} \times \omega_l\|$ occurs both in the denominator of $\vec{u}$ and the Jacobian $J_{\text{Q-B1D}}^{l-1,l}$, and can be cancelled. Taking the limit and simplifying yields a continuous *photon volume* with contribution

$$\langle D \rangle_{\text{C-B0D}}^{l-2,k} = \int_{s_{k-}}^{s_{k+}} f(\tilde{t}_{l-2}) f(\tilde{t}_{l-1}) f(\tilde{t}_l) \left\{ \frac{f_\omega^{l,k}}{J_{\text{C-B0D}}^{l-2,l-1,l}} \right\} f(s)\,\mathrm{d}s, \quad (24)$$

where $J_{\text{C-B0D}}^{l-2,l-1,l} = |\omega_{l-2} \cdot (\omega_{l-1} \times \omega_l)|$ is the Jacobian for 3D×1D coupling.

Summary. Our theory generalizes previous work by demonstrating how arbitrary numbers of distance sampling steps along photon paths can be replaced by expected value estimators, which ultimately results in higher-dimensional photon samples. While we explicitly derived new density estimators in 2D (photon planes, blurred and unblurred) as well as 3D (photon volumes), we can apply this process repeatedly to obtain estimators of arbitrary dimension.

While our derivations focused on photons, the exact same approach can also be applied to camera paths, leading to novel sensor planes, sensor volumes and beyond. Additionally, these sensor primitives can be combined freely with photon planes, photon volumes and so forth. The expressions for the resulting density estimators only depend on the combined dimensionality of the photon, sensor sample, and blur kernel. It is worth noting that this combined dimensionality has to be at least 3 for density estimation to succeed. This explains why unbiased photon planes are possible, whereas e.g. Beam–Beam estimates require at least a 1D blur.

Finally, our theory is not restricted to expected value estimators, and any estimator in our derivations can be replaced by a corresponding track-length estimator. For planes alone, this quadruples the number of estimators (long–long, short–long, long–short, short–short), depending on which segments use track-length estimation.

5 THEORETICAL ERROR ANALYSIS

Since our theory predicts a growing number of possible estimators, we perform an error analysis to better understand their behavior compared to prior approaches. Our analysis is inspired by the one of Křivánek et al. [2014], who examined *variance* in a single-scattering scenario with a point or beam estimator along a camera ray and light ray. Since our theory allows unfolding the distance propagation term for additional segments, we generalize their analysis by considering a 2-scattering scenario involving three path segments and include *bias* in our analysis, leading to new insights.

5.1 Analysis Setup

We consider a canonical configuration for 2-scattering as depicted in Fig. 4. Since our theory only changes how the propagation distances are evaluated, we leave the directions of the light and camera subpaths fixed, with the two directions of the photon path ω_1, ω_2 aligned with the x- and y-axes, and the camera ray direction ω_1' set to the z-axis. The site of density estimation corresponds to the intersection of the camera ray with the xy-plane. To simplify notation, we denote the distances to this point along each of the path segments as x, y, and z. Since the shape of the kernel is an arbitrary choice, we use a d-dimensional cube kernel of side-width u as the resulting separability will simplify our analysis. This kernel will be the product of 1D kernels along the three path segments, producing a 1D step, 2D square, or 3D cube kernel.

An abstract double-scattering estimator in this setup becomes

$$\langle D \rangle_*^{x,y,z} \sim \langle A(x) \rangle \langle A(y) \rangle \langle A(z) \rangle u^{-d} \quad (25)$$

where u is the kernel width and d the kernel dimension. We omit the directional scattering terms as they are fixed and influence the estimators by a shared constant factor. The three estimators $\langle A(x) \rangle$, $\langle A(y) \rangle$, $\langle A(z) \rangle$ determine the way transmittance is computed on the three segments, for which we consider three approaches:

$$\langle A(t) \rangle = \begin{cases} \langle \hat{T}_\mathrm{r}(t) \rangle_C & \text{collision estimator,} \\ \langle T_\mathrm{r}(t) \rangle_T & \text{track-length estimator,} \\ \langle T_\mathrm{r}(t) \rangle_E & \text{expected value estimator.} \end{cases} \quad (26)$$

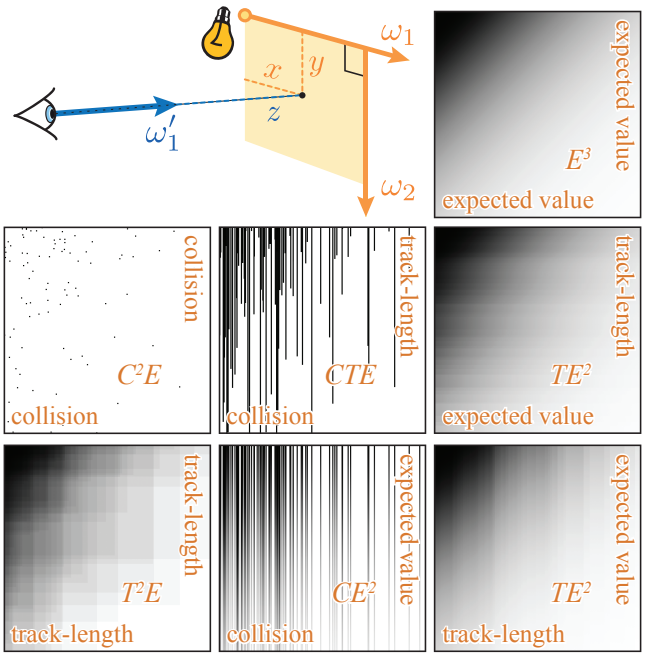

Fig. 4. Our canonical error analysis setup (top left) considers a double-scattering scenario where the two light path segments and the camera segment are mutually orthogonal. We analyze the error for different choices of collision (C), track-length (T), or expected value (E) estimators for transmittance along these three directions. The images visualize results from a simulation using 100 photons for a grid of camera rays spanning the xy-plane. The seven images here correspond to the estimators that always use (E) along the camera ray. The images are inverted to aid in visual comparison.

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

$$\hat{T}_r(t) \equiv \int_{t-u}^{t+u} T_r(t')\,dt' = \frac{e^{-\sigma_t(t-u)} - e^{-\sigma_t(t+u)}}{\sigma_t}. \quad (27)$$

5.2 Relation to Volumetric Density Estimators

The different combinations of these three estimators along the three segments yields $3^3 = 27$ possible definitions of $\langle D\rangle_*^{x,y,z}$. This allows us to examine, for our canonical configuration, the beam and/or point estimators developed by prior work, and also the estimators we derived in Sec. 4. All estimators from prior work simulate 2-scattering by restricting at least one of the axes, typically x, to a collision estimator, but allowing the other dimensions, y and z, to be any of the three estimators in Eq. (26). For instance, using a collision estimator along all three axes corresponds to traditional volumetric photon mapping with a 3D kernel; using a collision estimator along x, and track-length or expected value estimators along y and z result in the variety of short and long beam estimators from prior work. Our theory allows replacing the collision estimator along x with one of the other two estimators. For instance, using track-length estimators along x and y, and an expected value estimator along z results in the "short–short" variant of estimator (Eq. (21)).

In Fig. 4 we visualize the results of a photon simulation in this canonical setup using several of these possibilities (all combinations are provided in the supplemental). The simulation generates 100 photon paths, and evaluates the estimators for a 2D grid of camera rays parallel to the z axis.

5.3 Error Derivation

Given this generic setup, we wish to determine the relative root-mean-squared error (rRMSE) of Eq. (25):

$$\mathrm{rRMSE}[\langle D\rangle_*^{x,y,z}] = \sqrt{\mathbb{V}[\langle D\rangle_*^{x,y,z}] + \mathbb{B}^2[\langle D\rangle_*^{x,y,z}]}/D_E, \quad (28)$$

where $\mathbb{E}$, $\mathbb{V}$, and $\mathbb{B}^2$ are the mean, variance and squared bias of the estimator, and $D_E = T_r(x)T_r(y)T_r(z)$ is the unbiased reference value corresponding to using an expected value estimator along each axis.

Under our orthogonality assumption, the estimators along x, y, and z are statistically independent, so the expected value and variance of $\langle D\rangle_*^{x,y,z}$ can be calculated from the first and second moments of $\langle A(t)\rangle$ using the standard relations:

$$\mathbb{E}[\langle D\rangle_*^{x,y,z}] = u^{-d}\,\mathbb{E}[\langle A(x)\rangle]\,\mathbb{E}[\langle A(y)\rangle]\,\mathbb{E}[\langle A(z)\rangle] \quad (29)$$

$$\mathbb{V}[\langle D\rangle_*^{x,y,z}] = u^{-2d}\Big(\mathbb{E}[\langle A(x)\rangle^2]\,\mathbb{E}[\langle A(y)\rangle^2]\,\mathbb{E}[\langle A(z)\rangle^2] - \mathbb{E}[\langle A(x)\rangle]^2\mathbb{E}[\langle A(y)\rangle]^2\mathbb{E}[\langle A(z)\rangle]^2\Big), \quad (30)$$

with squared bias:

$$\mathbb{B}^2[\langle D\rangle_*^{x,y,z}] = \Big(\mathbb{E}[\langle D\rangle_*^{x,y,z}] - D_E\Big)^2. \quad (31)$$

The first and second moments for each choice of $\langle A(t)\rangle$ (collision, track-length, or expected value) were previously derived by Křivánek et al. [2014], which we restate here for completeness:

$$\mathbb{E}[\langle \hat{T}_r(t)\rangle_C] = \hat{T}_r(t), \qquad \mathbb{E}[\langle \hat{T}_r(t)\rangle_C^2] = \hat{T}_r(t)/\sigma_t, \quad (32)$$

$$\mathbb{E}[\langle T_r(t)\rangle_T] = T_r(t), \qquad \mathbb{E}[\langle T_r(t)\rangle_T^2] = T_r(t), \quad (33)$$

$$\mathbb{E}[\langle T_r(t)\rangle_E] = T_r(t), \qquad \mathbb{E}[\langle T_r(t)\rangle_E^2] = T_r(t)^2. \quad (34)$$

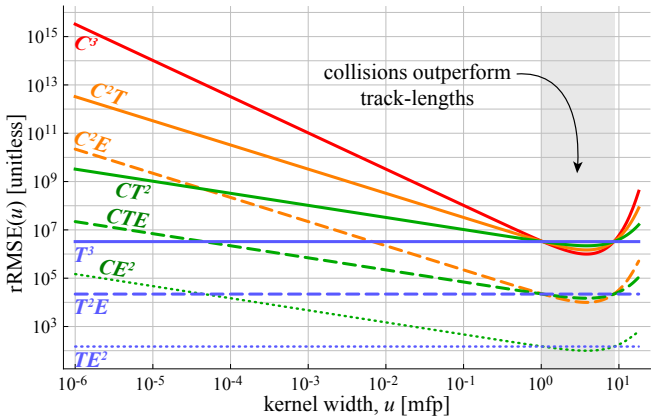

Fig. 5. LogLog visualization of relative root-mean-squared error (rRMSE) as a function of the kernel width u for the nine possible ways of choosing a collision (C), track-length (T), or expected value (E) estimator for 3-scattering. The graph labels denote the number of dimensions each type of estimator is used for. The E^3 estimator does not appear because its rRMSE is zero. Approaches with the same number of collision estimators use the same colors, and approaches with the same number of expected value estimators use the same line style (solid, dashed, dotted).

5.4 Error Comparison and Discussion

We can now analyze the estimators' errors by plugging the first and second moments (Eq. (32–34)) into the definitions of mean, variance and squared bias (Eq. (29–31)), and finally into Eq. (28) for relative error. To reduce the number of possible estimators down from the available 27, we set the distances along each of the axes to be equal: $x = y = z = 10$ mfp, where mfp $\equiv 1/\sigma_t$ is a mean free path. This leaves ten remaining unique possibilities, differing only by the *number* of times each type of estimator is used, and not their relative order. We denote these possibilities with the shorthand "C" for collision estimators, "T" for track-length, and "E" for expected value, and will express repetition of an estimator as an exponent. For instance, "C^2T" means two collision estimators and one track-length estimator. Fig. 5 compares the relative error of these ten estimators as we change the width u of the blur kernel.

Impact of bias. The intersection of the curves at $u = 1$ is due to the fact that collision estimators have lower variance than track-length estimators for blur widths that are greater than 1 mfp [Křivánek et al. 2014]. By examining error instead of just variance, however, we see that as the kernel width for collision estimation continues to grow, the introduced bias dominates, and the error rapidly increases once again. This shows that collision estimators are not always better for large blur widths, but only if the blur width is within a sweet-spot range highlighted in Fig. 5 (approximately 1–9 mfps in our setup).

Impact of expected value estimators. Since the expected value estimator contributes no variance, employing it will always reduce

variance compared to using the other estimators. This leads to three distinct groups of intersecting curves at u=1, which are determined by the number of times the expected value estimator is employed (decreasing/increasing our choice of $x = y = z$ would furthermore bring these groups closer/further apart, respectively). The graphs also show that the expected value and track-length estimators always lead to a constant, finite variance, irrespective of the blur width, since they perform no blurring.

Impact of collision estimators for small widths. Unfortunately, the presence of just one collision estimator will lead to infinite variance as $u \longrightarrow 0$. More interestingly, the *rate* at which the estimators shoot off to infinite variance differs based on the number of collision estimators used: the estimators with the same number of collision estimations (depicted in the same colors) share similar slopes. From a practical perspective this means that each additional replacement of a collision estimator with a track-length or expected value estimator will make variance asymptotically lower as the blur width diminishes. Current approaches *use collision estimation for all but the two connecting segments*, providing considerable opportunity for improvement especially for longer paths and small blur kernels.

5.5 Applicability to other Estimators

We only included estimators with the minimum required blur to make the number of possibilities manageable, but it would be trivial to also account for higher dimensions of blur by including track-length $\langle \hat{T}_\mathrm{r} \rangle_T$ and expected value estimators $\langle \hat{T}_\mathrm{r} \rangle_E$ of integrated transmittance in Eq. (26). We tried this for some common estimators, but since the total number of possible estimators for 2-scattering would grow to $3^5 = 243$ (or 56 if $x = y = z$), we have omitted these choices since we did not find them to provide further insights.

Since Eq. (25) is agnostic to which segments are on the camera vs. light subpath, the results from our analysis also apply to other choices of camera and light subpath lengths, as long as the total number of segments remains 3. For instance, the analysis is valid also for photon beam×sensor plane, or photon point×sensor volume.

While our analysis only considers up to double scattering, we anticipate that the insights gained here apply similarly to higher scattering orders. However, since it is not possible to create more than three mutually orthogonal directions in 3D, a formal variance analysis for 3+ scattering would need to account for the statistical dependence of the estimators along the additional path segments.

5.6 Singularities

All of our estimators contain an inverse Jacobian in their contribution, but our canonical setup side-steps their influence by choosing orthogonal scattering directions. In practice, however, the relative directions of successive path segments could cause these inverse Jacobians to become singular. To understand the potential impact on variance, we note that the Jacobian for blurred photon planes is in fact identical to the Jacobian of the Beam×Beam (1D blur) estimator, so potential singularities should be of the same order. Further analysis is required to quantify the variance behavior of unblurred photon planes and photon volumes, which share a similar Jacobian. One important advantage of our estimators over prior work is that the Jacobians for blurred photon planes and photon volumes involve only the photon directions (and not directions on both the camera and light subpaths), and could conceivably be cancelled using specially crafted importance sampling techniques [Georgiev et al. 2013].

6 IMPLEMENTATION

We validate our theoretical error analysis with implementations of several of the estimators predicted by our theory. Implementing all combinations of collision, track-length and expected value estimators across multiple photon path segments would be prohibitively time consuming, and instead we focus on a subset of 2D and 3D photon samples. Our most general implementation adds photon planes with track-length estimation to an open source ray tracing renderer. We select short–short planes due to their finite extent, which makes them more amenable to efficient acceleration compared to their long counterparts. We also demonstrate a hybrid CPU-GPU implementation, which traces photon paths on the CPU and rasterizes the resulting photons using the GPU.

6.1 General Ray Tracing Renderer

Our most general implementation proceeds similar to a traditional two-stage photon mapping algorithm. In the photon shooting stage, we deposit photon points on surfaces and both beams and planes inside the medium. For every segment of the photon path that lies inside the medium, we insert a photon plane by sampling an additional scattering direction and propagation distance. If the segment originated on a surface, we additionally insert a photon beam.

Special care needs to be taken at the intersection of photon paths with solid objects. To ensure energy conservation, the length of the segments used for a short plane must always be the sampled propagation distance, not the distance to the nearest surface. Planes can therefore potentially extend beyond the surfaces they intersect. Visibility tests will prevent light leaks across surface boundaries.

Acceleration structure. Due to the size and distribution of the photon planes, ray-plane intersection tests can be challenging to accelerate using traditional bounding volume hierarchies. We instead use a uniform grid to store both planes and beams in our implementation. We additionally use a specialized frustum-aligned grid for intersection tests with primary rays, which allows for efficient culling. We provide detailed performance comparison of these optimizations applied to beams and planes in Table 1 with respect to a baseline BVH implementation. In order to provide a fair comparison, we use the same optimized data structures for both beams and planes in our implementation.

Visibility caching. In contrast to photon beams, photon planes require a visibility test along the last segment of the photon path for each evaluation. The performance impact of these visibility tests can be significant, and we employ a caching strategy to accelerate rendering with blurred photon planes. Conceptually, we blur the visibility term along the first segment of the plane by reusing visibility tests for samples in close proximity. For a first intersection point $\mathbf{y}^{(1)}$ (see inset figure), we look up into the visibility cache consisting of regularly spaced bins of width u (the blur extent) along segment ω_{l-1}. A cache miss results in a ray being traced along ω_l

Table 1. Absolute render times and speedup (in parantheses) of our performance optimizations applied to photon beams and planes across all seven of our test scenes, with respect to a baseline BVH implementation. Optimizations are enabled incrementally from left to right. Please see Sec. 6 for details.

Scene	Baseline BVH			Uniform Grid			Frustum Grid			Vis. Cache
	B-B1D	Q-B0D	Q-B1D	B-B1D	Q-B0D	Q-B1D	B-B1D	Q-B0D	Q-B1D	Q-B1D
Bathroom	763s	3284s	5201s	750s (1.02×)	1440s (2.28×)	3249s (1.60×)	154s (4.96×)	1394s (2.36×)	2679s (1.94×)	515s (10.10×)
Bedroom	905s	3012s	4671s	973s (0.93×)	1403s (2.15×)	2797s (1.67×)	180s (5.03×)	1287s (2.34×)	2312s (2.02×)	807s (5.79×)
Kitchen	759s	2319s	3145s	814s (0.93×)	1095s (2.12×)	1896s (1.66×)	165s (4.61×)	1059s (2.19×)	1635s (1.92×)	566s (5.56×)
Living Room	614s	2065s	3076s	599s (1.03×)	876s (2.36×)	1840s (1.67×)	144s (4.25×)	826s (2.50×)	1402s (2.19×)	517s (5.94×)
Red Room	524s	2498s	3692s	356s (1.47×)	1022s (2.44×)	2080s (1.77×)	98s (5.32×)	1035s (2.41×)	1698s (2.17×)	645s (5.73×)
Doorway	545s	2334s	3190s	549s (0.99×)	1479s (1.58×)	2351s (1.36×)	194s (2.81×)	1433s (1.63×)	1803s (1.77×)	446s (7.15×)
Staircase	491s	2186s	2917s	568s (0.86×)	1019s (2.15×)	1819s (1.60×)	134s (3.67×)	1902s (1.15×)	1631s (1.79×)	506s (5.76×)

toward $\mathbf{y}^{(1)}$, and the bin corresponding to $t_{l-1}^{(1)}$ is populated with the distance to the nearest occluder. For a next intersection $\mathbf{y}^{(2)}$ with similar offset $t_{l-1}^{(2)}$ along ω_{l-1}, the stored hit distance is reused to check for occlusion of $\mathbf{y}^{(2)}$, which avoids one ray tracing operation. Conversely, the cache lookup of an intersection $\mathbf{y}^{(3)}$ with too dissimilar hit distance $t_{l-1}^{(3)}$ will fail, and a new ray is traced to populate the corresponding bin. To achieve a bounded memory footprint and avoid large up-front cost of the cache, we implement it as a per-thread hash map keyed by the index of the bin and the photon index. Performance savings depend on the blur extent of the plane, but are significant even on our test scenes with near imperceptible blur. We observe that the small additional blur of the visibility introduced by the caching is insignificant, but we cannot use this caching strategy for unblurred planes if we wish for them to remain unbiased. For this reason, we only enable visibility caching for blurred photon planes. We provide concrete performance numbers of this optimization in Table 1 (right column).

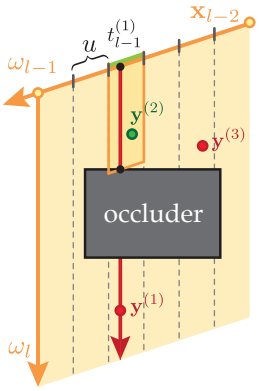

Control variates. Computing the contribution of a blurred photon plane involves the evaluation of an integral along the segment of overlap on the camera ray. Due to the presence of binary visibility in the integrand, this integral generally cannot be evaluated analytically, even in homogeneous media. Although the integral is straightforward to estimate with MC sampling, this process may introduce additional variance for long camera segments. In practice, we notice that only a small percentage of visibility tests end up occluded, which motivates us to use a *control variate* [Glasserman 2003] as a variance reduction technique. We provide full technical details of our approach in the supplemental material. Although this procedure adds to the rendering cost (10%-15% in our scenes), for planes with large blur radii it can reduce variance significantly compared to naïve MC sampling. Its usefulness is reduced in our test scenes with small blur radii, but we still incorporate it to remain robust for all plane configurations.

Although our implementation uses a number of performance optimizations, we stress that these do not diminish the theoretical improvements of our estimators. Optimizing ray-photon intersections is not specific to our method, and indeed, uniform grids have been employed in previous work [Křivánek et al. 2014] to accelerate rendering with photon beams. Additionally, we use the same data structures for both beams and planes to provide a fair comparison; in fact, photon beams benefit significantly more from these optimizations than photon planes (Table 1).

Despite these optimization techniques, our implementation does not require parameter fine-tuning to achieve reliable performance. The only user-exposed parameter is the grid resolution, which does not noticeably affect performance over a reasonable range of values.

6.2 Hybrid CPU-GPU Renderer

To generate the results in Fig. 2, we also implemented our approach using a hybrid CPU-GPU renderer specialized to the *searchlight problem*. Our implementation traces photon paths originating from a collimated beam through a semi-infinite homogeneous medium and rasterizes the resulting photon samples using the GPU. While this implementation only handles occlusion with the medium boundary, it demonstrates the relative performance of photon samples with up to 3 dimensions.

7 RESULTS

We demonstrate our implementation on seven indoor scenes containing scattering media, and compare the effectiveness of two of our estimators (Q-B1D and Q-B0D) against photon beams. All render times were measured on a Linux cluster with 8 core 2.7GHz E5-2680 CPUs and 64GB RAM, using 16 threads.

We use our uniform and frustum-aligned grids to accelerate both beams and planes, with the grid resolution adjusted for best performance in each scene. In all scenes, the grid resolution is relatively low (less than 50 cells on the largest dimension). Additionally, we make use of visibility caching and control variates for 1D blurred planes. We use the same blur radius for beams and blurred planes, which is set to cover roughly two pixels at the scene's focal point. Since photon planes are more costly to evaluate than photon beams, we perform an equal-time comparison to include both the quality of the estimators and their cost in the comparison metric. Each estimator is run for 10 minutes on each scene.

Additionally, we compute the variance of each estimator as an objective comparison metric. We run 100 instances of each estimator with a different random seed, and compute the variance between all

Table 2. Variance and effective speedup (in parantheses) of two of our estimators (Q-B0D and Q-B1D) and photon beams (B-B1D) on all seven of our test scenes. Please see Sec. 7 for details.

Scene	B-B1D	Q-B0D	Q-B1D
Bathroom	1.127 (1.00×)	0.462 (2.44×)	0.106 (10.68×)
Bedroom	0.901 (1.00×)	0.156 (5.79×)	0.071 (12.66×)
Kitchen	0.777 (1.00×)	0.163 (4.77×)	0.088 (8.84×)
Living Room	0.182 (1.00×)	0.032 (5.65×)	0.011 (16.41×)
Red Room	0.144 (1.00×)	0.010 (14.77×)	0.004 (40.03×)
Doorway	0.171 (1.00×)	0.060 (2.85×)	0.008 (20.70×)
Staircase	0.276 (1.00×)	0.073 (3.77×)	0.020 (14.14×)

100 renderings after five minutes of render time. For unblurred photon planes, the variance is simultaneously the error of the estimator, whereas photon beams and blurred photon planes contain an additional bias term that we do not compute. Since photon planes require a minimum of two or more scattering events in the medium, we exclude surface- and low-order medium scattering in our comparison and visualize only light paths with 2+ bounces in the medium. We stress that this is only for comparison reasons – our implementation can compute the full light transport in the scene.

We show renderings of four of our test scenes (Bathroom, Doorway, Kitchen, Living Room) at equal render time in Fig. 6, and full-resolution comparisons for all seven scenes in the supplemental material. The leftmost column shows the full light transport in the scene, whereas the images on the right demonstrate estimates of medium scattering obtained from each technique. For better comparison, we slightly boost the exposure of the medium-only images. We additionally show metrics across all seven of our test scenes in Table 2, both in terms of absolute variance and speedup (computed as the inverse of relative variance). The variance numbers were multiplied by a factor of 10^3 for formatting reasons.

In all scenes, both of our plane estimators provide a substantial improvement in image quality at equal render time. Additionally, both of our plane estimators provide a significant decrease in variance, corresponding to a decrease in render time of 2.4×–40× compared to photon beams at equal quality.

Despite the fact that unblurred photon planes are unbiased, they provide less of an improvement than blurred photon planes in all of our scenes. We believe this to be for two reasons: Since we only measure variance, not error, the absence of bias in unblurred photon planes is not captured by our metrics. Additionally, we make use of visibility caching for blurred photon planes, which leads to a decrease in the number of visibility tests by a factor of 3 on average, significantly improving render times. The same technique cannot be applied to unblurred photon planes without introducing bias.

The scenes in Fig. 6 demonstrate media with isotropic scattering. We provide additional results for media with anisotropic scattering in our supplemental material.

8 LIMITATIONS, DISCUSSION & FUTURE WORK

Heterogeneity. Although neither of our implementations currently support heterogeneous media, we stress that this is not a fundamental limitation of our theory. Evaluating expected value estimators is straightforward (but more costly) even in the presence of heterogeneity, and our 2+D estimators can be easily used to estimate light transport in such media. However, when *track-length estimation* is used on more than one path segment, the probability of constructing the path no longer cancels with the transmittance, and special care would need to be taken when using the short variants of our estimators. Since our implementations use track-length estimation exclusively, additional work is needed to support heterogeneity.

Visibility. Evaluating the contribution of an nD photon requires $n-1$ visibility tests, which could potentially be costly in practice. We use a visibility cache in our implementation to significantly reduce the number of visibility tests, but unfortunately this approach introduces additional bias and cannot be applied to unbiased photon samples. We believe an approach similar to unbiased deep shadow maps [Jarosz et al. 2011b] could lead to unbiased visibility caching, with performance implications for unbiased nD photon samples.

Path construction with nD samples. Previous work on combining points and beams with path tracing [Křivánek et al. 2014] could be extended to include the family of estimators introduced in this paper. Since the majority of our higher-dimensional estimators are unbiased, they operate natively in the same path space as regular path tracing, and their combination with other techniques is mathematically simpler than it is with points and beams. Additionally, our unblurred nD samples are not limited to the field of density estimation. The absence of bias in these estimators makes them suitable as standalone path construction techniques, and could be used to extend previous work on joint sampling of length–3 connection paths [Georgiev et al. 2013] to arbitrary lengths. Conceivably, these path construction techniques could even be used as mutation strategies in the context of Metropolis Light Transport [Pauly et al. 2000; Veach and Guibas 1997].

Surface rendering. Our estimators have potential applications even beyond density estimation within the medium. While the intersection of photon beams with surfaces results in photon points, surface intersections of photon planes induce line samples on surfaces, and photon volumes result in photon polygons. In general, nD photons induce $(n-1)$D samples when intersected with a surface, and including such higher-dimensional surface photons could lead to significant improvements in media-to-surface transport.

Extensions beyond distance sampling. While our derivations focus on replacing distance samples with other estimators, the same idea can be applied to directional and even emission sampling on the photon and camera subpath. Analogous to "photon marching", a whole new family of estimators can be obtained by "photon spinning", in which angular dimensions of directional sampling are collapsed into a higher dimensional photon. Dimensions of emission sampling can be collapsed to obtain line [Billen and Dutré 2016] and polygonal samples on the light source, and combining such samples with track-length or expected value estimation results in photon planes and photon volumes that can be used even for single scattering, as opposed to their current restriction to 2+ or 3+ scattering. Applying the same idea to the camera path results in line or polygonal samples on the aperture [Tzeng et al. 2012] or across time [Gribel et al. 2011, 2010], with implications for depth of field or motion blur.

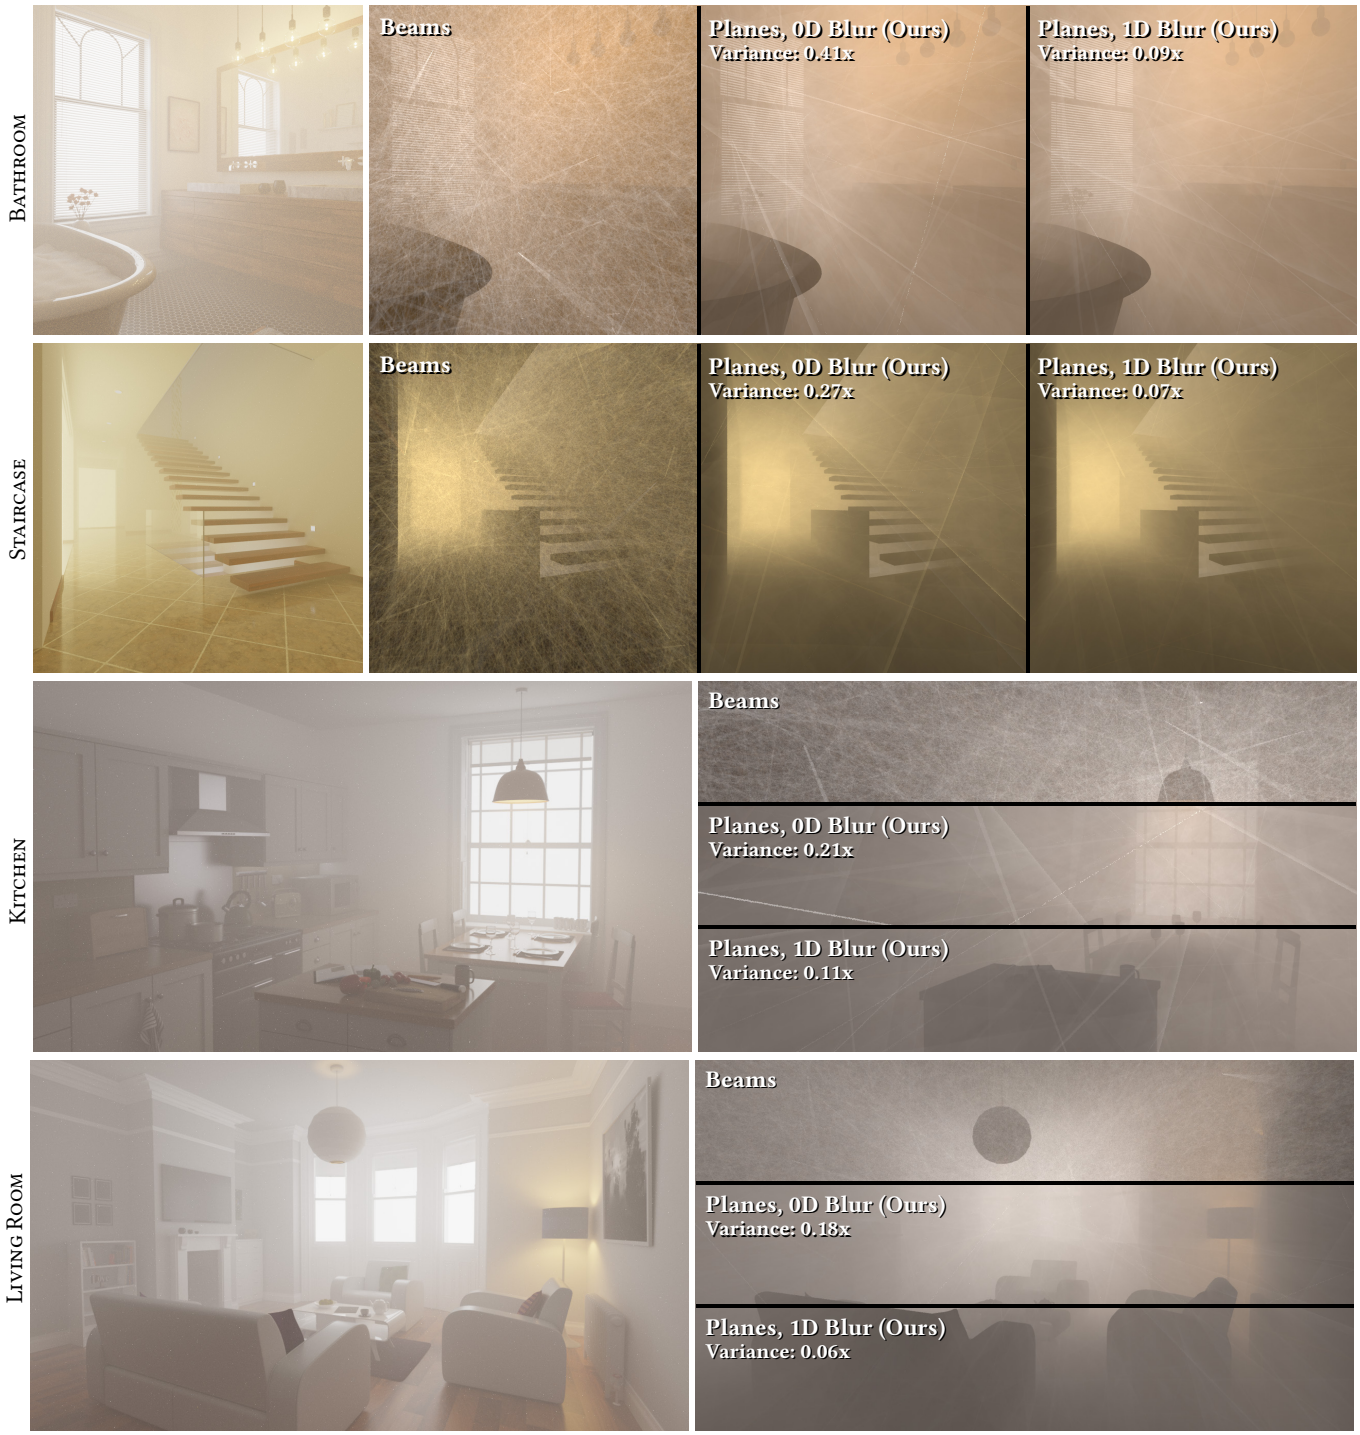

Fig. 6. We show renderings of four volumetric scenes produced by three different estimators at equal render time. We show the full light transport (left image) and multiply scattered volumetric transport (right images; exposure manually increased). Our estimators (photon planes, 0D and 1D blur) provide significant variance reduction and improved image quality compared to photon beams in all scenes. Please refer to the supplemental material for full-resolution comparisons across all seven of our test scenes.

9 CONCLUSION

We presented a new theory of density estimation that generalizes prior work on point and beam samples to arbitrary dimensions. The key idea was to replace an increasing number of distance sampling steps on the photon subpath with track-length or expected value estimators. Our theory predicts an unbounded number of new estimators, both on the photon- and camera subpath. Through careful error analysis, we reveal that in cases where photon beams provide error reduction over photon points, higher-dimensional photons compound this reduction even further. We extend prior variance analysis of these estimators to include their bias, and reveal that collision estimators only outperform track-length estimators in a narrow region of blur radii. We demonstrate practical implementations of some of the estimators predicted by our theory, and show that they provide significant variance reduction and image quality improvement across a variety of scenes, both at equal sample count and equal render cost.

ACKNOWLEDGEMENTS

We thank Jaroslav Křivánek, Iliyan Georgiev, Derek Nowrouzezahrai, and Toshiya Hachisuka for helpful discussions during early stages of this project. We thank the following blendswap.com artists for providing the scenes used in this paper: Mareck (Bathroom), Slyk-Drako (Bedroom), Jay-Artist (Kitchen, Living Room), Wig42 (Red Room, Doorway), NewSee2l035 (Staircase). This work was partially supported by a generous gift from Activision, as well as the National Science Foundation (Grant CNS-1205521).

REFERENCES

James Arvo. 1993. Transfer Functions in Global Illumination. In *ACM SIGGRAPH '93 Course Notes - Global Illumination.*

James Arvo and David Kirk. 1990. Particle transport and image synthesis. In *Proc. SIGGRAPH.* ACM, New York, NY.

James Richard Arvo. 1995a. *Analytic methods for simulated light transport.* Ph.D. Dissertation. Yale University.

James R. Arvo. 1995b. Applications of Irradiance Tensors to the Simulation of Non-Lambertian Phenomena. In *Proc. SIGGRAPH.*

Rasmus Barringer, Carl Johan Gribel, and Tomas Akenine-Möller. 2012. High-quality Curve Rendering Using Line Sampled Visibility. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 31, 6 (Nov. 2012).

Niels Billen and Philip Dutré. 2016. Line Sampling for Direct Illumination. *Computer Graphics Forum (Proc. EGSR)* 35, 4 (June 2016).

Subrahmanyan Chandrasekhar. 1960. *Radiative Transfer.* Dover Publications.

Min Chen and James Arvo. 2001. Simulating Non-Lambertian Phenomena Involving Linearly-Varying Luminaires. In *Rendering Techniques (Proc. EGWR).*

Min Chen and James R. Arvo. 2000. A Closed-Form Solution for the Irradiance due to Linearly-Varying Luminaires. In *Rendering Techniques (Proc. EGWR).*

Carsten Dachsbacher, Jaroslav Křivánek, Miloš Hašan, Adam Arbree, Bruce Walter, and Jan Novák. 2014. Scalable Realistic Rendering with Many-Light Methods. *Computer Graphics Forum* 33, 1 (2014).

Eugene D'Eon and Geoffrey Irving. 2011. A quantized-diffusion model for rendering translucent materials. *ACM Trans. Graph. (Proc. SIGGRAPH)* 30, 4 (2011).

Iliyan Georgiev, Jaroslav Křivànek, Tomas Davidovic, and Philipp Slusallek. 2012. Light transport simulation with vertex connection and merging. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 31, 5 (2012).

Iliyan Georgiev, Jaroslav Křivánek, Toshiya Hachisuka, Derek Nowrouzezahrai, and Wojciech Jarosz. 2013. Joint Importance Sampling of Low-Order Volumetric Scattering. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 32, 6 (Nov. 2013).

Paul Glasserman. 2003. *Monte Carlo Methods in Financial Engineering.* Stochastic Modelling and Applied Probability, Vol. 53. Springer New York, Chapter 4.

Carl Johan Gribel, Rasmus Barringer, and Tomas Akenine-Möller. 2011. High-Quality Spatio-Temporal Rendering using Semi-Analytical Visibility. *ACM Trans. Graph. (Proc. SIGGRAPH)* 30, 4 (Aug. 2011).

Carl Johan Gribel, Michael Doggett, and Tomas Akenine-Möller. 2010. Analytical Motion Blur Rasterization with Compression. In *Proceedings of HPG.*

Toshiya Hachisuka and Henrik Wann Jensen. 2009. Stochastic progressive photon mapping. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 28, 5 (2009).

Toshiya Hachisuka, Shinji Ogaki, and Henrik Wann Jensen. 2008. Progressive Photon Mapping. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 27, 5 (2008).

Toshiya Hachisuka, Jacopo Pantaleoni, and Henrik Wann Jensen. 2012. A path space extension for robust light transport simulation. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 31, 5 (2012).

Vlastimil Havran, Jiri Bittner, Robert Herzog, and Hans-Peter Seidel. 2005. Ray Maps for Global Illumination. In *Rendering Techniques (Proc. EGSR).*

David Immel, Michael Cohen, and Donald Greenberg. 1986. A radiosity method for non-diffuse environments. *Proc. SIGGRAPH* 20, 4 (1986).

Wojciech Jarosz, Derek Nowrouzezahrai, Iman Sadeghi, and Henrik Wann Jensen. 2011a. A Comprehensive Theory of Volumetric Radiance Estimation Using Photon Points and Beams. *ACM Trans. Graph.* 30, 1 (Feb. 2011).

Wojciech Jarosz, Derek Nowrouzezahrai, Robert Thomas, Peter-Pike Sloan, and Matthias Zwicker. 2011b. Progressive Photon Beams. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 30, 6 (Dec. 2011).

Wojciech Jarosz, Matthias Zwicker, and Henrik Wann Jensen. 2008. The Beam Radiance Estimate for Volumetric Photon Mapping. *Computer Graphics Forum (Proc. Eurographics)* 27, 2 (April 2008).

Henrik Wann Jensen. 2001. *Realistic Image Synthesis Using Photon Mapping.* A. K. Peters, Ltd., Natick, MA, USA.

Henrik Wann Jensen and Per H. Christensen. 1998. Efficient Simulation of Light Transport in Scenes With Participating Media Using Photon Maps. In *Proc. SIGGRAPH.*

Thouis R. Jones and Ronald N. Perry. 2000. Antialiasing with Line Samples. In *Rendering Techniques (Proc. EGWR).* Springer-Verlag, London, UK.

James T. Kajiya. 1986. The Rendering Equation. *Proc. SIGGRAPH* 20, 4 (Aug. 1986).

Claude Knaus and Matthias Zwicker. 2011. Progressive Photon Mapping: A Probabilistic Approach. *ACM Trans. Graph.* 30, 3 (2011).

Jaroslav Křivánek, 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 Trans. Graph. (Proc. SIGGRAPH)* 33, 4 (July 2014).

Eric Lafortune and Yves Willems. 1993. Bi-directional path tracing. In *Proc. Compugraphics.*

Eric Lafortune and Yves Willems. 1996. Rendering participating media with bidirectional path tracing. *Photorealistic Rendering Techniques (Proc. EGWR)* (1996).

Iván Lux. 1978. Unified Definition of a Class of Monte Carlo Estimators. *Nuclear Science and Engineering* 67, 1 (July 1978).

Jan Novák, Derek Nowrouzezahrai, Carsten Dachsbacher, and Wojciech Jarosz. 2012a. Progressive Virtual Beam Lights. *Computer Graphics Forum (Proc. EGSR)* 31, 4 (2012).

Jan Novák, Derek Nowrouzezahrai, Carsten Dachsbacher, and Wojciech Jarosz. 2012b. Virtual Ray Lights for Rendering Scenes with Participating Media. *ACM Trans. Graph. (Proc. SIGGRAPH)* 31, 4 (July 2012).

Derek Nowrouzezahrai, Ilya Baran, Kenny Mitchell, and Wojciech Jarosz. 2014. Visibility Silhouettes for Semi-Analytic Spherical Integration. *Computer Graphics Forum* 33, 1 (Feb. 2014).

Mark Pauly, Thomas Kollig, and Alexander Keller. 2000. Metropolis Light Transport for Participating media. In *Rendering Techniques (Proc. EGWR).*

Vincent Pegoraro and Steven G. Parker. 2009. An Analytical Solution to Single Scattering in Homogeneous Participating Media. *Computer Graphics Forum (Proc. Eurographics)* 28, 2 (2009).

Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically Based Rendering: From Theory To Implementation* (3rd ed.).

Peter Schröder and Pat Hanrahan. 1993. On the Form Factor Between Two Polygons. In *Proc. SIGGRAPH.* 163–164.

Jerome Spanier. 1966. Two Pairs of Families of Estimators for Transport Problems. *SIAM J. Appl. Math.* 14, 4 (1966).

Jerome Spanier and Ely Meyer Gelbard. 1969. *Monte Carlo principles and neutron transport problems.* Addison-Wesley.

Bo Sun, Ravi Ramamoorthi, Srinivasa G. Narasimhan, and Shree K. Nayar. 2005. A practical analytic single scattering model for real time rendering. *ACM Trans. Graph. (Proc. SIGGRAPH)* 24, 3 (2005).

Xin Sun, Kun Zhou, Stephen Lin, and Baining Guo. 2010. Line space gathering for single scattering in large scenes. *ACM Trans. Graph. (Proc. SIGGRAPH)* 29, 4 (2010).

Stanley Tzeng, Anjul Patney, Andrew Davidson, Mohamed S. Ebeida, Scott A. Mitchell, and John D. Owens. 2012. High-quality Parallel Depth-of-field Using Line Samples. In *Proceedings of HPG.*

Eric Veach. 1997. *Robust Monte Carlo methods for light transport simulation.* Ph.D. Dissertation. Stanford, CA, USA.

Eric Veach and Leonidas Guibas. 1994. Bidirectional estimators for light transport. In *Photorealistic Rendering Techniques (Proc. EGWR).*

Eric Veach and Leonidas Guibas. 1995. Optimally combining sampling techniques for Monte Carlo rendering. *Proc. SIGGRAPH* 29 (1995).

Eric Veach and Leonidas Guibas. 1997. Metropolis light transport. *Proc. SIGGRAPH* 31 (1997).