

Photon surfaces for robust, unbiased volumetric density estimation

XI DENG*, SHAOJIE JIAO*, BENEDIKT BITTERLI, and WOJCIECH JAROSZ, Dartmouth College, USA



Fig. 1. We compare the equal-time variance of different estimators in a scene containing participating media. We show the full light transport in the scene (left), single scattering (middle/right, top half) and multiple scattering volumetric transport (middle/right, bottom half). Our estimators (middle) provide significant variance reduction compared to prior density estimators (right) at equal render time.

We generalize photon planes to *photon surfaces*: a new family of unbiased volumetric density estimators which we combine using multiple importance sampling. To derive our new estimators, we start with the extended path integral which duplicates the vertex at the end of the camera and photon subpaths and couples them using a blurring kernel. To make our formulation unbiased, however, we use a delta kernel to couple these two end points. Unfortunately, sampling the resulting singular integral using Monte Carlo is impossible since the probability of generating a contributing light path by independently sampling the two subpaths is zero. Our key insight is that we can eliminate the delta kernel and make Monte Carlo estimation practical by integrating any three dimensions analytically, and integrating only the remaining dimensions using Monte Carlo. We demonstrate the practicality of this approach by instantiating a collection of estimators which analytically integrate the distance along the camera ray and two arbitrary sampling dimensions along the photon subpath (e.g., distance, direction, surface area). This generalizes photon planes to curved “photon surfaces”, including new “photon cone”, “photon cylinder”, “photon sphere”, and multiple new “photon plane” estimators. These estimators allow us to handle light paths not supported by photon planes, including single scattering, and surface-to-media transport. More importantly, since our estimators have complementary strengths due to analytically integrating different dimensions of the path integral, we can combine them using multiple importance sampling. This combination mitigates singularities present in individual estimators, substantially reducing variance while remaining fully unbiased. We demonstrate our improved estimators on a number of scenes containing homogeneous media with highly anisotropic phase functions, accelerating both multiple scattering and single scattering compared to prior techniques.

CCS Concepts: • Computing methodologies → Ray tracing.

Additional Key Words and Phrases: global illumination, light transport, participating media, photon density estimation

*Joint first authors.

Authors' address: Xi Deng, xi.deng.gr@dartmouth.edu; Shaojie Jiao, shaojie.jiao.gr@dartmouth.edu; Benedikt Bitterli, benedikt.bitterli.gr@dartmouth.edu; Wojciech Jarosz, wojciech.k.jarosz@dartmouth.edu, Dartmouth College, 9 Maynard Street, Hanover, NH, 03755, USA.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3306346.3323041>.

ACM Reference Format:

Xi Deng, Shaojie Jiao, Benedikt Bitterli, and Wojciech Jarosz. 2019. Photon surfaces for robust, unbiased volumetric density estimation. *ACM Trans. Graph.* 38, 4, Article 46 (July 2019), 12 pages. <https://doi.org/10.1145/3306346.3323041>

1 INTRODUCTION

Accurate and efficient simulation of light transport in participating media is a challenging yet important problem in fields ranging from the movie industry, medical imaging, video games and even nuclear reactor design. Solving the governing rendering equation [Immel et al. 1986; Kajiya 1986] and equation of radiative transfer [Chandrasekhar 1960] has been the subject of substantial research over the past decades [Cerezo et al. 2005; Novák et al. 2018a,b], resulting in a wealth of rendering algorithms that tackle this problem.

Unbiased rendering methods such as path tracing [Kajiya 1986] are among the oldest forms of Monte Carlo (MC) light transport, and they remain popular [Pharr et al. 2016; Christensen and Jarosz 2016; Fassione et al. 2017] due to their simplicity and ability to render images where the only error is noise. Naive path tracing converges poorly in difficult lighting scenarios, so subsequent variants such as bidirectional path tracing (BPT) [Veach and Guibas 1994, 1997; Lafourcade and Willem 1993] augment it with a diverse set of complementary path sampling strategies which can be combined using multiple importance sampling (MIS) [Veach and Guibas 1995].

Photon density estimation using points [Jensen 1996; Jensen and Christensen 1998] or beams [Jarosz et al. 2008, 2011a] gains efficiency partly from the path reuse permitted by the density estimation framework. This efficiency, however, comes at the cost of bias, which manifests as blurring in the image. In the same spirit as BPT, recent work has sought to combine some of these estimators with unbiased methods [Georgiev et al. 2012; Hachisuka et al. 2012] and each other [Křivánek et al. 2014] using MIS to leverage their complementary strengths. Unfortunately, this is made difficult by the fact that photon mapping and unbiased path sampling approaches operate in path spaces of different dimension, and the resulting combination remains biased. Most recently, Bitterli and Jarosz [2017] developed the first unbiased photon density estimators in the form of photon planes and volumes; however, these unbiased estimators

1) suffer from singularities, and 2) require at least two bounces in the medium past any surface, relegating the remaining transport to other (potentially) biased or inferior techniques (Fig. 1 right).

In this paper, we present a general framework for deriving unbiased volumetric density estimators by first elevating the problem to the extended path integral [Hachisuka et al. 2012, 2017] and then analytically integrating out any three of its dimensions. We identify a particularly practical subset of estimators enabled by our theory, which we call *photon surfaces*, representing the surface arising from sweeping out the location of a photon for all possible values of two of its sampling dimensions (e.g., distance, direction, surface area). Each choice of these two analytically integrated dimensions results in a different type of estimator with complementary strengths and singularities, including new “photon cone”, “photon cylinder”, “photon sphere”, and multiple new “photon plane” estimators. Our estimators support a wider range of volumetric transport, including single-scattering, which was previously impossible with unbiased density estimators (Fig. 1 top). Since our estimators operate in standard path space, we can combine them using MIS while avoiding the complexities of bridging different dimensionalities of path space. This results in a more robust rendering algorithm which mitigates singularities in an unbiased way (Fig. 1 center). We also introduce a generalized version of MIS that allows for combining an infinite continuum of strategies arising from our theory, compared to the discrete number of strategies considered by standard MIS. Finally, we show how existing unbiased path sampling techniques such as next-event estimation, BPT and virtual point lights (VPLs) [Dachs-bacher et al. 2014] can be derived from our framework through a particular choice of dimensions for analytic integration. This means that we can interpret our approach more generally as a way to derive unbiased path sampling strategies, opening the possibility of combining our estimators more broadly with existing unbiased path sampling methods for increased efficiency and robustness.

1.1 Related Work

We focus primarily on connections to prior work in volumetric density estimation and analytic integration, and refer to Pharr et al. [2016] and Novák et al. [2018a] for a broader survey of recent (volumetric) rendering techniques.

Density estimation has a long history of using higher-dimensional samples to reduce variance and bias, starting with beams [Havran et al. 2005; Jarosz et al. 2008, 2011a; Sun et al. 2010], their progressive [Hachisuka et al. 2008] variants [Jarosz et al. 2011b], and more recent planes and volumes [Bitterli and Jarosz 2017]. These higher dimensional samples were traditionally understood in terms of the limit process or expected value of free-flight sampling, but we can express them more generally in our framework as techniques that analytically pre-integrate two (for beams) or more (for planes/volumes) consecutive distance dimensions prior to MC integration. Beam estimators collapse only two of the three singular dimensions, and subsequently blur has to be introduced to permit MC integration. In contrast, photon planes allow for unbiased estimation because they collapse all singular dimensions. Our formulation allows us to pick non-consecutive distance dimensions to obtain complementary photon planes, or angular/area dimensions to obtain novel photon

surfaces, all of which can be combined with MIS to robustly mitigate the singularities present in any one estimator alone. Our estimators are consistent with the density estimation framework and can be combined with orthogonal approaches such as gradient domain rendering [Lehtinen et al. 2013; Kettunen et al. 2015], as has previously been shown for beam/plane estimators [Gruson et al. 2018].

Analytic integration of parts of the path integral has previously been proposed for a wide range of applications such as anti-aliasing [Jones and Perry 2000], visibility [Gribel et al. 2010, 2011; Barringer et al. 2012; Nowrouzezahrai et al. 2014; Billen and Dutré 2016], depth of field [Tzeng et al. 2012], single scattering [Sun et al. 2005; Pegoraro and Parker 2009], and area lighting [Arvo 1995b,a; Chen and Arvo 2000, 2001; Belcour et al. 2018]. These have also been combined with more general Monte Carlo integration using ratio [Heitz et al. 2018] or control variates estimators [Belcour et al. 2018]. Recent theoretical analyses [Sun et al. 2013; Singh et al. 2017; Singh and Jarosz 2017] have explained how such techniques reduce variance and improve convergence rate by smoothing the integrand prior to Monte Carlo integration. Similar to prior work, we also rely on analytic integration to derive our estimators. In contrast however, our formulation begins with the *extended* path space [Hachisuka et al. 2012, 2017], which dramatically simplifies the problem by ensuring that the integrand is singular and all required integrals can be computed easily in closed form. The resulting expressions are remarkably simple and resemble generalized geometry terms rather than cumbersome antiderivatives.

Surprisingly, unbiased path sampling techniques such as BPT, VPLs and next-event estimation can also be formulated in our framework as the result of analytic integration of parts of the extended path integral. Related techniques such as virtual ray lights [Novák et al. 2012b,a] and joint importance sampling of three-segment connections [Georgiev et al. 2013] are also the result of analytic integration (i.e. marginalization), although they do not map directly to our theory. More generally, we can view our approach as a way to generate new parametrizations of paths, and the weight of our estimators then directly corresponds to the geometry factor of our parametrization. This is comparable to previous work on path manifolds [Jakob 2013], which introduces a reparametrization of specular chains with corresponding generalized geometry factor.

2 BACKGROUND

We begin by defining our notation and briefly reviewing volumetric light transport in the extended path integral framework [Veach 1997; Hachisuka et al. 2017].


2.1 Path Integral Framework

In the path integral framework [Veach 1997; Pauly et al. 2000], the intensity of a pixel I is defined as the integral

$$I = \int_{\Xi} f(\bar{z}) d\mu(\bar{z}), \quad (1)$$

which considers the measurement contribution function $f(\bar{z})$ of all paths \bar{z} over the space Ξ of all possible light transport paths connecting light sources to the sensor.

Path. Since we derive our estimators from photon mapping, we express them initially in an extended path space [Hachisuka et al. 2017] where a complete light transport path $\bar{z} = \bar{x} \bar{y}$ consists of two disconnected subpaths: a photon subpath $\bar{x} = (x_l \dots x_0)$ where vertex x_l resides on a light source, and a camera subpath $\bar{y} = (y_0 \dots y_k)$ where y_k resides on the sensor (see figure below).



Note that we number the subpath quantities starting at the duplicated connection vertex ($y_0 x_0$). Furthermore, we use ω_i and t_i to denote the direction and distance *leaving* x_i towards x_{i-1} on a photon subpath, and similarly, ψ_i and s_i for the direction and distance from y_i towards y_{i-1} on a camera subpath.

Measurement Contribution Function. The measurement contribution function $f(\bar{z})$ along a full path \bar{z} is the product of the contribution along each subpath, coupled by the scattering phase function $f_\omega^{1,1}$ and normalized 3D blurring kernel K at the connection:

$$f(\bar{z}) = f(\bar{x}) K(g) f_\omega^{1,1} f(\bar{y}) \quad (2)$$

where $g = x_0 - y_0$ is the offset vector between the endpoints of the two subpaths. The contribution of the light subpath \bar{x} is

$$f(\bar{x}) = \prod_{i=1}^l f_t(t_i) f_\omega(\omega_i), \quad (3)$$

The propagation contribution $f_t(t_i)$ is the product of transmittance T_r and binary visibility V between x_i and x_{i-1} :

$$f_t(t_i) = T_r(t_i) V(x_i, x_{i-1}), \quad \text{with} \quad T_r(t) = e^{-\sigma_t t} \quad (4)$$

and extinction coefficient σ_t , assuming homogeneous media. The directional contribution $f_\omega(\omega_i)$ is the scattering coefficient-weighted phase function $\sigma_s \rho_p$, BSDF ρ_s or emission L_e at vertex x_i :

$$f_\omega(\omega_i) = \begin{cases} L_e(x_i, \omega_i) \cos \theta_i & i = l \\ \rho_s(\omega_{i+1}, \omega_i) \cos \theta_i & x_i \in \text{surface and } i < l \\ \sigma_s \rho_p(\omega_{i+1}, \omega_i) & x_i \in \text{medium and } i < l. \end{cases} \quad (5)$$

where θ_i is the angle between outgoing direction and surface normal. The camera subpath contribution is defined analogously but with s_i and ψ_i taking the place of t_i and ω_i , and the emitted importance $W_e(y_k, \psi_k)$ taking the place of emitted radiance in Eq. (5).

Path Space and Measure. The integration measure μ is the product measure over all integration dimensions $\bar{\xi}$ that control the location of the path vertices in \bar{z} :

$$d\mu(\bar{z}) = d\bar{\xi} = \prod d\xi_i \quad (6)$$

where ξ_i denotes an integration dimension.

Since we chose to define our measurement contribution function (2) in distance-direction space¹, the integration dimensions

¹Note that it is possible to convert from distance-direction to the more typical vertex product measure by incorporating a geometry term for each path segment in \bar{x} and \bar{y} .

consist of the starting vertices and collection of distances and directions along each subpath:

$$\bar{\xi} = \{x_l, \bar{t}, \bar{\omega}; \bar{\psi}, \bar{s}, y_k\}, \quad (7)$$

where $\bar{\omega} = \omega_l \dots \omega_1$, $\bar{t} = t_l \dots t_1$ and $\bar{s} = s_k \dots s_1$, $\bar{\psi} = \psi_k \dots \psi_1$ denote the sequence of directions and distances along both subpaths.

When ξ_i is in $\{\bar{t}, \bar{s}\}$, then it is a one-dimensional variable over the positive real line, with corresponding measure. When it is one of $\{x_l, y_k, \bar{\omega}, \bar{\psi}\}$, then ξ_i is a two-dimensional variable corresponding to surface area or solid angle measure, respectively.

For convenience, we can rewrite our measurement contribution function in terms of the abstract integration variables $\bar{\xi}$:

$$f(\bar{z}(\bar{\xi})) = f(\bar{\xi}) K(g(\bar{\xi})) f_\omega^{1,1} \quad (8)$$

where $f(\bar{\xi}) = f(\bar{x}) f(\bar{y}) = \prod f(\xi_i)$, with

$$f(\xi_i) = \begin{cases} f_\omega(\xi_i) & \text{if } \xi_i \in \{x_l, y_k, \bar{\omega}, \bar{\psi}\} \\ f_t(\xi_i) & \text{if } \xi_i \in \{\bar{t}, \bar{s}\} \end{cases}. \quad (9)$$

The path integral then takes the following form:

$$I = \int_{\Xi} f(\bar{\xi}) K(g(\bar{\xi})) f_\omega^{1,1} d\bar{\xi}. \quad (10)$$

Monte Carlo Estimation in Photon Mapping. Photon mapping uses a normalized 3D kernel and approximates Eq. (10) using Monte Carlo:

$$\langle I \rangle \approx \frac{\int f(\bar{\xi}) K(g(\bar{\xi})) f_\omega^{1,1} p(\bar{\xi})}{p(\bar{\xi})}, \quad (11)$$

where $p(\bar{\xi})$ is the joint PDF of sampling all dimensions in $\bar{\xi}$. Unfortunately, the blurring kernel in the integrand leads to bias.

3 A GENERAL FRAMEWORK

In this paper, we are interested in unbiased Monte Carlo estimators. Theoretically these are possible to express in the extended path integral framework by shrinking the blurring kernel to a Dirac delta so a path contributes only when $x_0 = y_0$. The kernel becomes:

$$K(g) = \delta^3(g) = \delta(x(g)) \delta(y(g)) \delta(z(g)) \quad (12)$$

where $x(g)$, $y(g)$, and $z(g)$ return the Cartesian x -, y - and z -coordinates of the vector $g = x_0 - y_0$, respectively. From the incremental construction of x_0 and y_0 , we know that g is influenced by all the elements in variable set $\bar{\xi}$:

$$g(\bar{\xi}) = x_0(\bar{\xi}) - y_0(\bar{\xi}) = \left(x_l + \sum_{i=l}^1 t_i \omega_i \right) - \left(y_k + \sum_{i=k}^1 s_i \psi_i \right). \quad (13)$$

Unfortunately, in practice it is impossible to estimate Eq. (11) in the presence of this delta kernel since the probability of generating a valid light path by independently sampling the two subpaths is zero.

Our key insight is that we can eliminate the delta function and make Monte Carlo practical by splitting the extended path integral into two parts: three dimensions of which we integrate analytically, and the remaining we estimate numerically with Monte Carlo. Critically, the analytic integration eliminates the delta function from the integrand, allowing subsequent integration with Monte Carlo.

To accomplish this, we introduce a set $\xi_a \subset \bar{\xi}$ which consists of three dimensions which we will integrate analytically, and we

denote the remaining (numerically integrated) dimensions $\overline{\xi}_n = \overline{\xi} \setminus \overline{\xi}_a$. We can now write the measurement contribution function as the product of a contribution $f(\overline{\xi}_a)$ depending on $\overline{\xi}_a$, and a contribution $f(\overline{\xi}_n)$ depending on $\overline{\xi}_n$:

$$f(\overline{\xi}) = f(\overline{\xi}_a)f(\overline{\xi}_n). \quad (14)$$

The path integral in Eq. (10) can therefore be rewritten as:

$$I = \int_{\Xi_n} f(\overline{\xi}_n) \underbrace{\int_{\Xi_a(\overline{\xi}_n)} f(\overline{\xi}_a) \delta^3(g(\overline{\xi}_a)) f_\omega^{1,1} d\overline{\xi}_a d\overline{\xi}_n}, \quad (15)$$

where we use the shorthand $g(\overline{\xi}_a)$ for $g(\overline{\xi}_a, \overline{\xi}_n)$ since we only focus on the analytic part.

With this decomposition, we can construct a virtually unbounded collection of practical, unbiased Monte Carlo estimators of the form:

$$\langle I \rangle_{a,n} \approx \frac{f(\overline{\xi}_n) I_a(\overline{\xi}_n)}{p(\overline{\xi}_n)}, \quad (16)$$

where $p(\overline{\xi}_n)$ is the joint PDF of sampling only the dimensions in $\overline{\xi}_n$, and $I_a(\overline{\xi}_n)$ is the analytically preintegrated path contribution underbraced in Eq. (15) for a particular choice of three analytic dimensions denoted by a . Importantly, different choices of these dimensions result in distinct hybrid Monte Carlo-analytic estimators.

In the next section we will first model this procedure to rederive Bitterli and Jarosz's unbiased photon planes [2017] which correspond to analytically integrating three distance sampling dimensions. We then (Sec. 5) show how to derive a number of new "photon surface" estimators arising from different choices of analytic integration dimensions. With a plethora of new unbiased estimators at our disposal, we then show how to combine their relative strengths using multiple importance sampling (Sec. 6).


4 PHOTON PLANES

The photon plane estimator employs edges $\mathbf{x}_2\mathbf{x}_1$ and $\mathbf{x}_1\mathbf{x}_0$ to form a photon plane and a ray along edge $\mathbf{y}_1\mathbf{y}_0$ to query the plane. In our framework, by setting $\overline{\xi}_a = \{t_2, t_1, s_1\}$, I_a in Eq. (15) becomes:

$$I_a = \int_{\Xi_a} f_t(t_2) f_t(t_1) f_t(s_1) \delta^3(g(\overline{\xi}_a)) f_\omega^{1,1} dt_2 dt_1 ds_1. \quad (17)$$

The delta kernel is taken with respect to the offset vector \mathbf{g} , but to integrate it out, we need to express it directly in terms of the integration variables $\overline{\xi}_a$. Composing the delta function with a function with one root gives us:

$$\delta^3(g(\overline{\xi}_a)) = \frac{\delta^3(\overline{\xi}_a - \overline{\xi}_a^*)}{\left| \frac{\partial \mathbf{g}}{\partial \overline{\xi}_a}(\overline{\xi}_a^*) \right|} = \frac{\delta^3(\overline{\xi}_a - \overline{\xi}_a^*)}{J_{\overline{\xi}_a}^g(\overline{\xi}_a^*)}, \quad (18)$$



where $\overline{\xi}_a^* = \{t_2^*, t_1^*, s_1^*\}$ are the values of $\overline{\xi}_a$ for which $g(\overline{\xi}_a) = \mathbf{0}$ (i.e. the values of $\overline{\xi}_a$ at the intersection between the camera ray and the photon plane, as illustrated to the left), and the denominator evaluates the change-of-variable Jacobian J from $\overline{\xi}_a$ to \mathbf{g} at this root.

For this specific choice of $\overline{\xi}_a$, we have:

$$g(\overline{\xi}_a) = \underbrace{(\mathbf{x}_2 + \omega_2 t_2 + t_1 \omega_1)}_{\mathbf{x}_0(t_2, t_1)} - \underbrace{(\mathbf{y}_1 + s_1 \psi_1)}_{\mathbf{y}_0(s_1)} = \mathbf{x}_2 - \mathbf{y}_1 + \mathbf{A} \overline{\xi}_a \quad (19)$$

where $\mathbf{A} = [\omega_2, \omega_1, -\psi_1]$. Then the Jacobian is

$$J_{t_2, t_1, s_1}^g(\overline{\xi}_a^*) = \det(\mathbf{A}) = |(\omega_2 \times \omega_1) \cdot \psi_1|. \quad (20)$$

Inserting Eq. (18) with Eq. (20) into Eq. (17) allows us to preintegrate over $\overline{\xi}_a$, eliminating the delta function and yielding:

$$I_a = \frac{f_\omega^{1,1} f_t(t_2^*) f_t(t_1^*) f_t(s_1^*)}{|(\omega_2 \times \omega_1) \cdot \psi_1|}. \quad (21)$$

After inserting Eq. (21) into Eq. (16) and accounting for the different notation, this estimator matches Bitterli and Jarosz's 0D photon plane estimator exactly.

"Short" vs. "Long" variants. Eq. (21) corresponds to a "long" or "expected value" [Spanier 1966] estimator which requires evaluating the distance throughput term (containing the transmittance) at each hitpoint of a semi-infinite plane. The alternative "short" or "track-length" [Spanier 1966] variety estimates transmittance with a constant step function that drops to zero beyond the next sampled propagation distance, effectively turning the semi-infinite plane into a finite parallelogram. Since the derivation of our Jacobians is orthogonal to which of these strategies is used, we omit this extra complexity from our equations. Even though we restrict ourselves to homogeneous media where the transmittance can be evaluated analytically, we always use the short variety in our diagrams (and the final implementation) since their finite extent eases illustration (and enables more efficient ray intersection).

5 PHOTON SURFACES

In Sec. 3, we introduced a general theory for deriving unbiased density estimators by picking any three dimensions of the extended path integral and integrating them analytically. This leads to an infinite family of new and unusual unbiased estimators. In this section, we focus on a subset of this family in which exactly two elements of $\overline{\xi}_a$ lie on the photon path, and the remainder is always s_1 , i.e. the distance along the last segment of the camera path.

Analytically integrating two dimensions of the photon path while keeping all other dimensions fixed corresponds to "sweeping" the last vertex \mathbf{x}_0 of the photon subpath into a two-dimensional parametric surface. Depending on which dimensions are picked, this surface may be planar, cylindrical, spherical and more (Fig. 2). At the same time, analytic integration of s_1 instantiates a ray query along the last segment of the camera path.

Interpreting estimators in this manner highlights the advantages of restricting ourselves to this subset of estimators: Finding the root of \mathbf{g} for these estimators corresponds directly to computing the intersection between a ray and an analytic primitive. This makes root finding comparatively easy, and it maps well to current approaches implementing photon beams or planes. Additionally, it allows us to take full advantage of existing acceleration structures for ray-primitive queries.

Further, the Jacobian of such estimators has an intuitive interpretation. Let $\overline{\xi}_{a_1}$ and $\overline{\xi}_{a_2}$ represent the two analytic dimensions on the




Fig. 2. Our framework allows us to derive new density estimators depending on which dimensions are analytically integrated. Integrating non-consecutive distance dimensions lead to generalized photon planes (left two); integrating azimuth and distance leads to photon cones and cylinders (middle, middle-right); integrating azimuth and inclination leads to photon spheres (right).

photon subpath. The Jacobian is then

$$\frac{\mathbf{J}\mathbf{g}}{\xi_a}(\bar{\xi_a}) = \det \left[\frac{\partial \mathbf{x}_0}{\partial \xi_{a_1}}(\bar{\xi_a}), \frac{\partial \mathbf{x}_0}{\partial \xi_{a_2}}(\bar{\xi_a}), -\frac{\partial \mathbf{y}_0}{\partial s_1}(\bar{\xi_a}) \right] = \left| \mathbf{n}(\bar{\xi_a}) \cdot \psi_1 \right|, \quad (22)$$

where $\mathbf{n} = \partial \mathbf{x}_0 / \partial \xi_{a_1} \times \partial \mathbf{x}_0 / \partial \xi_{a_2}$ can be interpreted as a scaled *surface normal* of the photon primitive that arises from “sweeping” \mathbf{x}_0 for all values of ξ_{a_1} and ξ_{a_2} . The Jacobian of such an estimator is then simply the dot product of the surface normal and the query ray direction. Intuitively, this means that photon surfaces will become brighter when viewed at grazing angles.

It is worth noting that the normal \mathbf{n} is not necessarily of unit length. Indeed, its length encodes how a differential 2D element stretches and squishes as we map from parametric space (ξ_{a_1}, ξ_{a_2}) to a differential area on the photon surface. For example, in the case of a photon plane, a sheared plane will be brighter as it leads to a shorter normal and thus smaller Jacobian.

By interpreting these estimators as ray-surface intersections, we can easily see that it is possible for \mathbf{g} to have multiple roots. For example, in the case of a spherical photon surface (Fig. 2, right), a ray can hit both the front- and back side of the sphere. This requires a more general form of Eq. (18) that holds for multiple roots:

$$\delta^3(\mathbf{g}(\bar{\xi_a})) = \sum_r \frac{\delta^3(\bar{\xi_a} - \bar{\xi_a^r})}{\left| \frac{\partial \mathbf{g}}{\partial \xi_a}(\bar{\xi_a^r}) \right|}, \quad (23)$$

where the summation is over all roots $\bar{\xi_a^r}$. In practical terms, this means that the total contribution of such an estimator is simply the sum of contributions over all hitpoints. This results in a remarkably simple formula for the estimator contribution, generalizing the photon plane estimator (21) to photon surfaces:

$$I_a = \sum_r \frac{f_\omega^{1,1} f(\bar{\xi_a^r})}{\left| \mathbf{n}(\bar{\xi_a^r}) \cdot \psi_1 \right|}. \quad (24)$$

In the following subsection, we will instantiate this estimator for specific choices of ξ_{a_1} and ξ_{a_2} , listing the required scaled surface normals for Eq. (24) to obtain new photon cone, sphere, cylinder estimators and more.

5.1 Multiple-Scattering Photon Surfaces

Generalized Photon Planes. Previously we integrated out the last two propagation distance dimensions $\xi_{a_1} = t_1$ and $\xi_{a_2} = t_2$ to obtain an unbiased photon plane (Eq. (21)). We can create a larger family of generalized, multi-scattering photon plane estimators by simply choosing *any* two distance dimensions from \bar{t} . Let us denote these $\xi_{a_1} = t_i$ and $\xi_{a_2} = t_j$ where $l \geq i > j \geq 1$. Intuitively, integrating

Table 1. Each choice of two dimensions ξ_{a_1}, ξ_{a_2} leads to a different multiple-scattering photon surface with corresponding scaled surface normal.

Primitive	$[\xi_{a_1}, \xi_{a_2}]$	Scaled surface normal $\mathbf{n}(\bar{\xi_a})$
Plane	$[t_i, t_j]$	$(\omega_i \times \omega_j)$
Sphere	$[\cos \theta_1, \phi_1]$	$\omega_1^* t_1^2$
Cone	$[t_1, \phi_1]$	$(\omega_1^* \times (\omega_1^* \times \omega_2)) t_1^*$
Cylinder	$[t_2, \phi_1]$	$(\omega_2 \times (\omega_1^* \times \omega_2)) t_1$
Disk	$[t_1, \cos \theta_1]$	$\frac{(\omega_1^* \times \omega_2) t_1^*}{\sin^2 \theta_1}$
Toroid	$[\phi_i, \phi_j]$	$\frac{(\omega_{i+1} \times \omega_i^*) \times (\omega_{j+1}^* \times \omega_j^*) ((\mathbf{x}_0^* - \mathbf{x}_i) \times \omega_{i+1}^*) (\mathbf{x}_0^* - \mathbf{x}_j^*) \times \omega_{j+1}^* }{\sin \theta_i \sin \theta_j}$
Hyperboloid	$[\phi_i, t_1]$	$\frac{(\omega_{i+1} \times \omega_i^*) \times \omega_i^* (\mathbf{x}_0^* - \mathbf{x}_i) \times \omega_{i+1}^* }{\sin \theta_i}$

out these two dimensions sweeps the photon position \mathbf{x}_0 over the plane spanned by the corresponding directions ω_i and ω_j (see Fig. 2, left). We provide a full derivation in the supplemental document, but the resulting scaled surface normal is simply $\omega_i \times \omega_j$ (instead of $\omega_2 \times \omega_1$), as listed in Table 1. Inserting this surface normal into Eq. (24) gives us a generalized photon plane estimator.

Photon Spheres. It is also possible to integrate out directional dimensions. For instance, analytically integrating ω_1 can be thought of as sweeping the photon position \mathbf{x}_0 over all 4π steradians centered at \mathbf{x}_1 , resulting in a spherical photon surface (see Fig. 2, right). The surface normal (see Table 1) is simply the normal of the sphere ω_1^* , but scaled by its squared radius t_1^2 to account for the change from solid angle (unit sphere) to the surface area of a non-unit sphere.² Just as with generalized photon planes, we are not limited to choosing the last directional domain, but can choose an arbitrary ω_i along the path to obtain different photon sphere estimators. We include a full derivation in the supplemental document and list the resulting normal that needs to be inserted into Eq. (24) in Table 1.

Photon Cones, Cylinders, and beyond. By decomposing a direction ω_i into spherical angles $\phi_i, \cos \theta_i$, we can mix and match an analytic distance dimension with a polar or azimuthal angle. Simultaneously spinning the photon \mathbf{x}_0 along ϕ_i and sweeping along distance t_i results in a photon cone, while choosing ϕ_i and the distance along the previous segment t_{i+1} leads to a photon cylinder (see Fig. 2, middle). In general, it is possible to mix and match the angles and distances from different bounces, producing more general photon toroids and hyperboloids. We list the scaled surface normals for several options in Table 1.

²Integrating both direction and distance at a vertex leads to an infinite “spherical photon volume”, whose Jacobian is exactly equivalent to the standard geometry term for next-event estimation, VPLs, or shadow rays in BPT. This means all of these techniques can be derived from our framework, for a particular choice of integration dimensions.

5.2 Single-Scattering Photon Surfaces

Since a photon plane estimator Eq. (21) needs at least two propagation distances along the photon subpath, it cannot handle single scattering. Our more general photon surface estimator (24) can, however, support single scattering since it allows us to choose other analytic integration dimensions. More concretely, any photon surface that integrates at most one distance sampling dimension analytically can be used for single scattering.

Photon Cones and Spheres. Previously, multiple scattering photon cones and spheres resulted from integrating the last polar angle $\cos \theta_1$ and distance t_1 or the last sampled solid angle ω_1 on the photon path. When $l = 1$, these photon surfaces simulate single scattering. For photon cones, the apex will be a position on the light source, and spheres would generalize to photon spherical caps for spot lights and photon hemispheres for area light sources, where L_e replaces the phase function in the path contribution.

Photon Surface Area Lights. So far we have considered directional and propagation distance dimensions for analytic integration, but an as-of-yet unexploited possibility is the surface area of the light source itself. Just like solid angle ω , the starting position \mathbf{x}_l of the photon path is a two-dimensional variable. Analytically integrating these two dimensions corresponds to sweeping the starting location \mathbf{x}_l over the entire surface area of the light source. Since the location of photon \mathbf{x}_0 depends on the initial emission location, this results in a remarkably simple photon surface: an exact copy of the light source (with the same surface normal), but positioned at the photon. This allows creating single-scattering photon surfaces from arbitrarily complicated mesh or parametric surface lights (see Fig. 4 for an illustrative example using an emissive Stanford bunny).

Photon u or v Surfaces. Just as with the two-dimensional solid angle domain, given some parametrization $(u, v) \mapsto \mathbf{x}_l$, we can decompose the 2D surface area of the light into individual parametric dimensions. This allows us to create additional types of photon surfaces which mix and match an analytically integrated u or v with a propagation distance or emission angle $\xi_{a_2} \in \{t_l, \phi_l, \cos \theta_l\}$. This results in a photon surface with scaled normal $\mathbf{n}(\bar{\xi}_a^*) = \partial \mathbf{x}_0 / \partial u \times \partial \mathbf{x}_0 / \partial \xi_{a_2} = \partial \mathbf{x}_l / \partial u \times \partial \mathbf{x}_0 / \partial \xi_{a_2}$ for u , and analogously for v . While this is theoretically possible for any type of parametrization and light source shape, for simplicity we illustrate and implement this idea for a rectangular light source and an orthonormal surface parametrization in Fig. 4, where we combine u or v with the first propagation distance t_l . This produces a single-scattering photon plane with scaled normal identical to that of the generalized plane in Table 1, but with ω_j set to ω_l , and the unit vectors defining the u, v axes ($\mathbf{u} = \partial \mathbf{x}_l / \partial u$ or $\mathbf{v} = \partial \mathbf{x}_l / \partial v$) taking the place of ω_i . Note that by rotating the u, v parametrization (Fig. 4d), we can arrive at a continuum of differently oriented single-scattering photon planes.

6 MULTIPLE IMPORTANCE SAMPLING

We could now render volumetric transport using any one of the many photon surfaces estimators. Unfortunately, each of the photon surface estimators has a singularity where the Jacobian in the denominator may approach zero. This means that certain (portions of) photon surfaces can become arbitrarily bright, producing distracting




Fig. 3. We compare MIS (right column) to straight averaging (left column) of generalized photon planes (top) and continuous single scattering planes (bottom). MIS provides significant variance reduction at equal render time.

artifacts and high variance in the rendered image. By examining the form of the Jacobian (22), such singularities can occur whenever a photon surface is viewed at a grazing angle and also when the scaled surface normal approaches zero length (e.g. at the apex of a photon cone, or if the segments defining a photon plane are coplanar).

While we could remove the singularities by blurring the photon surfaces along their normals (as previously done for photon planes [Bitterli and Jarosz 2017]), this would introduce bias. Instead, since we now have many estimators at our disposal, each with singularities arising in complementary configurations, we show how we can combine their respective strengths in an unbiased way using multiple importance sampling [Veach and Guibas 1995]. This leads to a significantly more robust combination than if we had simply averaged the estimator contributions (Fig. 3).

6.1 Combining a Discrete Collection of Strategies

We first observe that each intersection between a camera ray and photon surface corresponds to a standard light transport path (Fig. 2). This allows us to interpret our photon surface estimators as unbiased Monte Carlo path sampling strategies in standard path space.

Traditionally, the MIS weight of a sampling strategy a is defined as the ratio of its PDF, $p_a(\bar{z})$, to the sum of the PDFs of other strategies that can produce the same path \bar{z} : $w_a(\bar{z}) = p_a^\beta(\bar{z}) / \sum_{k=1}^m p_k^\beta(\bar{z})$, where $\beta = 1$ gives the balance heuristic. This computation requires bringing each sampling PDF into a common space (e.g., solid-angle-distance, or surface area product measure) for proper comparison. While this is theoretically possible to do for our estimators, it is cumbersome since each one was derived by marginalizing away different integration dimensions, making their remaining domains different. Instead, we can bypass the need to compute path PDFs entirely by using the inverse estimator score as a proxy for the PDF, as proposed by Jendersie [2018]. The power heuristic for a photon surface which analytically integrates away dimensions $\bar{\xi}_a$ becomes:

$$w_a(\bar{z}) = \langle I \rangle_a^{-\beta}(\bar{z}) / \sum_{k=1}^m \langle I \rangle_k^{-\beta}(\bar{z}), \quad (25)$$

where $\langle I \rangle_k$ is the score (16) returned by the k -th estimator for the path \bar{z} . Note that the sum over roots for curved photon surfaces (24) actually considers multiple full paths, each requiring their own contribution, Jacobian, and MIS weight computation. Photon surfaces




Fig. 4. We show illustrations (top row) and renderings (bottom row; left/right half are low/high sample count) of single-scattering photon surfaces. Analytically integrating the surface parametrization of a light source (left two columns) leads to (u, v) -photon surfaces that are shaped duplicates of the light source itself (e.g. spherical, left top; bunny-shaped, left bottom; square, second column). Alternatively, we can integrate the distance along the first segment of the photon subpath and only one dimension of the surface parametrization, leading to (u, t) and (v, t) -planes that form lines at their intersection with the light source (middle, middle right). For a square, there are infinitely many orthogonal parametrizations, and there exists a continuum of such (u_α, t) planes (right).

that could intersect the camera ray multiple times must be handled with care, since each such intersection corresponds to a different light path. We find the light paths corresponding to each intersecting point and apply Eq. (25) for each one separately.

6.2 Combining a Continuum of Strategies

As mentioned in Sec. 5.2, there is a continuum of possible single-scattering u - or v -planes, one for every possible rotation of the u, v parametrization on the light source. This means that instead of a large but countable number of strategies, we have a continuum of strategies corresponding to the same light transport. Each one of these possible strategies has a different photon surface normal and Jacobian, and therefore singularity. While we could simply choose a random u, v orientation for each emitted single-scattering plane and use its corresponding scaled normal to compute the Jacobian, each such Jacobian has the potential to go singular.

To mitigate these singularities, we generalize the balance heuristic to allow combining a continuum of strategies. We first interpret the balance heuristic as a standard MC estimator which draws a sample from the average PDF of the available strategies:

$$\langle I \rangle_{\text{MIS}} = \frac{\bar{q}_a p_a(\bar{z})}{\sum_{k=1}^m q_k p_k(\bar{z})} \frac{f(\bar{z})}{\bar{q}_a p_a(\bar{z})} = \frac{f(\bar{z})}{\sum_{k=1}^m q_k p_k(\bar{z})}, \quad (26)$$

where $q_k = 1/m$ is the probability of selecting strategy k . By taking the limit as the number of strategies goes to infinity, the denominator becomes the average PDF of all possible randomly oriented single-scattering planes. This average PDF leads to a remarkably simple effective Jacobian (derivation in the supplemental document):

$$J_{\xi_a}^g(\xi_a^*) = \frac{2}{\pi} \sqrt{((\mathbf{u} \times \omega_l) \cdot \psi_1)^2 + ((\mathbf{v} \times \omega_l) \cdot \psi_1)^2}. \quad (27)$$

In contrast to a single arbitrarily oriented photon plane, which can go singular at any grazing angle, this Jacobian can only go singular when both the terms in the square root go to zero, which happens only if the camera direction ψ_1 is parallel to the emission direction ω_l . We can then treat this combined estimator as a single strategy, denoted (u_α, t) -planes, which we combine with (u, v) -planes (or cones or spheres) for additional robustness.

7 IMPLEMENTATION

We implemented our new estimators in the open-source Tungsten [Bitterli 2018] renderer to leverage its reference implementation of unbiased photon planes [Bitterli and Jarosz 2017] for comparison. We implemented a wide assortment of the photon surfaces predicted by our theory. For 3+ scattering we include $(\cos \theta_1, \phi_1)$ -spheres, (t_1, ϕ_1) -cones, and (t_2, ϕ_1) -cylinders, as well as the three possible ways of forming multi-scattering photon planes using the last three propagation distances: (t_1, t_2) -, (t_1, t_3) -, and (t_2, t_3) -planes. For brevity, we will simply refer to these as photon spheres, cones, cylinders and “3-planes”. For 2+ scattering we can also use spheres, cones, and cylinders, but only the original (t_1, t_2) -plane. For single scattering we can use spheres and cones, and additionally include four different single-scattering photon planes: (u, v) -, (u, t_1) -, (v, t_1) -, as well as randomly rotated (u, t_1) -planes.

After tracing a photon subpath, we iterate over all bounces, generate all supported photon surfaces, and store them in a uniform grid acceleration structure. During rendering, we compute the intersection between the camera ray and all stored photon surfaces. To compute the contribution of a photon surface, we first compute the full light path corresponding to each hitpoint, evaluate the corresponding estimator (16), and compute its MIS weight (25) by evaluating the scores of all other possible estimators for this same light path. For the purposes of MIS, we only need to evaluate the subset of the light path score affected by the differing Jacobians.

Handling visibility. Care must be taken to properly handle visibility during both the rendering and photon tracing stages. An intersection between a camera ray and a photon surface determines a complete light transport path. However, certain segments of the full light path \bar{z} differ from the originally constructed light and camera subpaths, and need to be retraced to evaluate the visibility in the path contribution. The number of needed visibility queries depends directly on how far the last analytic integration dimension is up the photon subpath, e.g., for a (t_1, t_2) -plane we only need to trace one shadow ray, but (t_1, t_3) - and (t_2, t_3) -planes both require two shadow queries (see the light subpath overlays in Fig. 2 and Fig. 4). For analogous reasons, photon surfaces need to be constructed using the sampled free-flight distance in the medium (and not terminated




Fig. 5. We compare six of our single-scattering estimators individually (top row), their MIS combinations (bottom row), and path tracing and photon beams baselines at equal time. We overlay variance relative to path tracing. MIS(All) excludes (u_α, t) -planes, as these are redundant with (u, t) and (v, t) . Although estimators that use spinning may be visually converged, they suffer from singularities/excessive variance when used by themselves. This diminishes with MIS, but the added cost does not always lead to an overall improvement. See supplemental for a table of overhead of individual estimators and MIS combinations.

at the surface intersection). Photon surfaces that use analytic dimensions further up the photon subpath require this precaution for more free-flight steps. Single-scattering surfaces are constructed without considering visibility at all, since the one additional visibility query is performed during camera tracing once the hitpoint is known. We only included photon surfaces that integrate analytic dimensions as close as possible to x_0 to reduce the cost of visibility queries.

8 RESULTS

We compare our new estimators against prior work in a variety of test scenes containing participating media. We render all images using the same renderer and compare them at equal time. We list the variance of each approach as an objective comparison metric which we estimate by rendering 100 images with different random seeds and computing the sample variance across all images. For unbiased approaches, the variance is simultaneously the mean squared error, but for photon beams, the variance underestimates the error.

The first row of Fig. 7 compares photon beams against our single-scattering planes (Fig. 4 middle three columns) in an outdoor scene. We show both a combination using MIS (right half) and a straight average of estimator contributions (left half). We compare more estimators and their MIS combinations in Fig. 5. Using MIS offers significant variance reduction over both the averaged combination and photon beams, demonstrating the benefit of our diverse surfaces and their ability to use MIS. The benefit of single-scattering planes over beams depends on the size of the light source. For scenes with huge area lights like the CLASSRoom in Fig. 1 they provide significant improvements, but for the relatively small lights of GASSTATION in Fig. 7, they are outperformed by photon beams or our photon cones and spheres. Fig. 6 evaluates how the light source size affects the relative performance of single-scattering planes.

In Fig. 7 (rows 2–3), we compare Bitterli and Jarosz’s photon planes to our 3-planes as well as a combination of 3-planes and photon cones and cylinders (bottom half). These estimators require at least 3 medium segments, and we only show medium transport handled by all estimators. Using 3-planes by themselves offers significant variance reduction over previous work due to the added robustness of MIS. In the presence of anisotropy (third row), cones and cylinders provide additional robustness and variance improvement, although they are of limited usefulness in the isotropic case. We refer the reader to the supplemental material for additional results comparing our estimators with path tracing.




Fig. 6. We examine the effect of light source size (rows) on the relative performance of continuously rotated (u_α, t) -planes (first column) and (u, v) -planes (second column). For each estimator we show its contribution individually (left split, AVG), as well as weighted by MIS (right split). The third column shows the weighted and unweighted combination of these two estimators, as well as their variance (relative to beams). The (u, v) -planes perform well for large light sources, but for smaller lights their weight diminishes compared to (u_α, t) -planes, which start to resemble the photon beams baseline, though without bias. Expanded comparison is in the supplemental.

Finally, we compare both single-scattering and multiple-scattering in Fig. 7 (rows 3–5) for both our single-scattering planes (top half) and our 3-planes (bottom half) against standard photon beams and planes. Because our combination of planes requires at least three photon segments, it cannot be used for double scattering. We instead combine photon cones, cylinders and standard planes to handle this transport robustly. This combination of estimators handles all transport supported by prior work at significantly reduced variance.

9 CONCLUSIONS, LIMITATIONS, & FUTURE WORK

We presented a new framework for deriving unbiased density estimators. By coupling the light- and camera- subpath endpoints using a delta distribution, but pre-integrating three of its dimensions analytically, we can construct an infinite collection of new unbiased estimators of the standard volumetric path integral. Specific choices of the three analytic dimensions correspond to existing unbiased estimators, but other choices allow estimating previously unsupported transport. Moreover, each choice leads to an estimator with different strengths and weaknesses, which we show how to combine for increased robustness using MIS. While our method




Fig. 7. We compare different estimators in volumetric scenes. We show the full transport (left), volumetric transport estimated by our estimators (middle) and by previous photon beams/planes (right). The first row compares single scattering produced by our single scattering planes, without MIS (left half) and with MIS (right half) to photon beams. The second and third rows compare different combinations of our photon surfaces to photon planes for an isotropic (second row) and anisotropic phase function (third row) for triply scattered volumetric light transport. The remaining rows compare single scattering (top half) and doubly scattered volumetric light transport for photon surfaces and photon beams/planes.



Fig. 8. The “long” forms of some of our estimators can be applied to heterogeneous media, here comparing single scattering at equal photon count.

exhibits significantly lower variance compared to prior unbiased density estimators, limitations and avenues for future work remain.

Heterogeneity. Our *theory* is not limited to homogeneous media; however, several practical issues need to be overcome to translate this theory into an efficient solution for general heterogeneous media. Generally speaking, applying our approach to heterogeneous media requires 1) evaluating transmittances; and (potentially) 2) knowing and evaluating the free-flight distance PDFs.

For (t_1, t_2) -planes, $(u, t)/(v, t)$ -planes, and cones, we can trivially accommodate heterogeneity by simply using their “long” or “expected value” forms (c.f. Sec. 4) and incurring requirement 1). Fig. 8 shows a proof-of-concept implementation of this. However, evaluating the contribution of these estimators involves evaluating visibility *and transmittance* along one or two segments of the light path, which adds non-negligible computational cost in heterogeneous media. In addition, accelerations structures become less effective with semi-infinite primitives, further degrading performance.

For estimators which use a numerical free-flight sampling step after the last analytic dimension (e.g. photon spheres, cylinders, and generalized photon planes), supporting heterogeneity becomes more challenging, as the free-flight PDF of these photon surfaces no longer perfectly cancels the transmittance. This will generally increase variance and also requires the ability to evaluate free-flight PDFs within the medium. The latter prohibits the use of null-collision-based sampling methods (which do not currently provide evaluable PDFs), so methods such as raymarching (biased) or regular tracking (computationally expensive) would need to be used instead.

Surface-to-medium transport. For the special case of direct emission, we derived photon surfaces for single scattering (Sec. 5.2), but we currently lack a robust solution for medium interactions immediately following scattering from a surface. Photon spinning could be used for such transport, but care is needed to ensure the BSDF is sampled well. This is trivial for the sphere estimator and the cone estimator with symmetric BSDFs, but efficiently sampling a cone from arbitrary BSDFs would require more work. This is less of a problem in media, since phase functions are typically radially symmetric. One potential solution is to perform analytic integration in the space of random numbers rather than spherical coordinates, akin to primary sample space methods [Kelemen et al. 2002]. This could leverage all importance sampling already present in the rendering system, but the resulting photon surfaces may be difficult to intersect due to their complex shape. Fig. 9 demonstrates surface-to-media transport using photon spheres and cones in a scene with purely Lambertian surfaces, allowing our method to handle full light transport in an unbiased way. Unfortunately, in the absence of




Fig. 9. While using photon cones and sphere for surface-to-media transport (right) enables our approach to handle all light transport in the medium (left), remaining singularities in the surface-to-media transport dominate variance and allow path tracing and photon beams to outperform our approach.

other surface-to-media estimators, singularities remain in the MIS combination, so this transport dominates the variance.

None of these techniques provide a benefit for scattering from specular surfaces, and different techniques are required to render effects such as volume caustics. A potential solution is to integrate the dimensions immediately preceding specular interactions, but the resulting photon surface would depend on the geometry of the specular object and would be difficult to determine in general.

Medium-to-surface and surface-to-surface transport. The focus of this paper was on participating media, and our method currently does not provide a benefit to surface rendering. However, our framework can readily derive practical “photon volumes” that induce transport on surfaces. Indeed, virtual point lights are just one special case from this family of estimators, and through choosing different integration dimensions, we could derive estimators that could remove the singularities encountered in VPLs in an unbiased way.

“Short” spheres. For highly anisotropic media, our photon spheres become computationally inefficient since the path contribution is small for a majority of the sphere even though the entire sphere is stored and intersected during rendering. It may be possible to create “spherical sector” estimators that integrate only parts of the directional domain – akin to “short” transmittance estimators – to concentrate effort where the phase function is high. This is comparable to sampling with line segments rather than lines, and would likely lead to similar variance tradeoffs [Singh et al. 2017].

MIS(u_α, t) vs. MIS($(u, t), (v, t)$) planes. We use (u, t) - and (v, t) -planes in Figs. 1 and 7 because we found (see Fig. 5 in supplemental) this works slightly better than (u_α, t) -planes. We suspect this is because for quad lights, (u_α, t) -planes can be narrow and make worse angles with the camera compared to (u, t) - and (v, t) -planes in our scenes.

Unbiased path sampling. Although we focused on deriving density estimators in this paper, our framework can be interpreted more broadly as a method for deriving unbiased path sampling techniques. These techniques operate in the standard path space, and they could be readily combined with existing unbiased rendering techniques such as bidirectional path tracing to improve robustness. In addition, it is trivial to convert our photon surfaces into equivalent “sensor surfaces” which could be used as alternative shadow connection strategies in unbiased rendering algorithms.

ACKNOWLEDGMENTS

We would like to thank members of the Dartmouth Visual Computing Lab for fruitful discussions and the anonymous reviewers for suggestions on improving the paper. We use scenes provided by the following *Blend Swap* artists: **Mareck** (BATHROOM), **Jay-Artist** (KITCHEN, LIVINGROOM), **NovaAshbell** (CLASSROOM), and **Wig42** (HORSEROOM, DININGROOM). The bunny model was provided by the *Stanford 3D Scanning Repository* and the cloud data by *Walt Disney Animation Studios*. This work was supported by National Science Foundation Grants IIS-1812796 and CNS-1205521.

REFERENCES

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

James Richard Arvo. 1995b. Applications of irradiance tensors to the simulation of non-Lambertian phenomena. In *Annual Conference Series (Proceedings of SIGGRAPH)*. ACM, 335–342. <https://doi.org/10/c2fss9>

Rasmus Barringer, Carl Johan Gribel, and Tomas Akenine-Möller. 2012. High-quality curve rendering using line sampled visibility. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 31, 6 (Nov. 2012), 162:1–162:10. <https://doi.org/10/f25qxk>

Laurent Belcour, Guofu Xie, Christophe Hery, Mark Meyer, Wojciech Jarosz, and Derek Nowrouzezahrai. 2018. Integrating clipped spherical harmonics expansions. *ACM Transactions on Graphics* 37, 2 (March 2018). <https://doi.org/10/gd52p2>

Niels Billen and Philip Dutré. 2016. Line sampling for direct illumination. *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 35, 4 (June 2016), 93–102. <https://doi.org/10/84z2h>

Benedikt Bitterli. 2018. Tungsten Renderer. <https://github.com/tunabrain/tungsten>.

Benedikt Bitterli and Wojciech Jarosz. 2017. Beyond points and beams: Higher-dimensional photon samples for volumetric light transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 36, 4 (July 2017), 1–12. <https://doi.org/10/gfznbr>

Eva Cerezo, Frederic Pérez, Xavier Pueyo, Francisco J. Serón, and François X. Sillion. 2005. A survey on participating media rendering techniques. *The Visual Computer* 21, 5 (2005), 303–328. <https://doi.org/10/cjxqd1>

Subrahmanyam Chandrasekhar. 1960. *Radiative transfer*. Dover Publications, NY.

Min Chen and James Richard Arvo. 2000. A closed-form solution for the irradiance due to linearly-varying luminaires. In *Rendering Techniques (Proceedings of the Eurographics Workshop on Rendering)*, 137–148. <https://doi.org/10/gfz9gv>

Min Chen and James Richard Arvo. 2001. Simulating non-Lambertian phenomena involving linearly-varying luminaires. In *Rendering Techniques (Proceedings of the Eurographics Workshop on Rendering)*, 25–38. <https://doi.org/10/chb4qg>

Per H. Christensen and Wojciech Jarosz. 2016. The path to path-traced movies. *Foundations and Trends in Computer Graphics and Vision* 10, 2 (Oct. 2016), 103–175. <https://doi.org/10/gfjwjc>

Carsten Dachsbacher, Jaroslav Krivánek, Miloš Hašan, Adam Arbree, Bruce Walter, and Jan Novák. 2014. Scalable realistic rendering with many-light methods. *Computer Graphics Forum* 33, 1 (2014), 88–104. <https://doi.org/10/f5twgd>

Luca Fascione, Johannes Hanika, Marcos Fajardo, Per Christensen, Brent Burley, Brian Green, Rob Pieké, Christopher Kulla, Christophe Hery, Ryusuke Villemain, Daniel Heckenberg, and André Mazzone. 2017. Path tracing in production (Parts 1 and 2). In *ACM SIGGRAPH Courses*. <https://doi.org/10/gfz2ck>

Iliyan Georgiev, Jaroslav Krivánek, Tomas Davidovic, and Philipp Slusallek. 2012. Light transport simulation with vertex connection and merging. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 31, 5 (2012), 192:1–192:10. <https://doi.org/10/gbb6q7>

Iliyan Georgiev, Jaroslav Krivánek, Toshiya Hachisuka, Derek Nowrouzezahrai, and Wojciech Jarosz. 2013. Joint importance sampling of low-order volumetric scattering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 32, 6 (Nov. 2013), 164:1–164:14. <https://doi.org/10/gbd5qs>

Carl Johan Gribel, Rasmus Barringer, and Tomas Akenine-Möller. 2011. High-quality spatio-temporal rendering using semi-analytical visibility. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 30, 4 (Aug. 2011), 54:1–54:11. <https://doi.org/10/fqg7b9>

Carl Johan Gribel, Michael Doggett, and Tomas Akenine-Möller. 2010. Analytical motion blur rasterization with compression. In *Proceedings of High Performance Graphics*, 163–172. <https://doi.org/10/fzz5ds>

Adrien Gruson, Binh-Son Hua, Nicolas Vibert, Derek Nowrouzezahrai, and Toshiya Hachisuka. 2018. Gradient-domain volumetric photon density estimation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 37, 4 (July 2018), 82:1–82:13. <https://doi.org/10/gd52p6>

Toshiya Hachisuka, Iliyan Georgiev, Wojciech Jarosz, Jaroslav Krivánek, and Derek Nowrouzezahrai. 2017. Extended path integral formulation for volumetric transport. In *Proceedings of the Eurographics Symposium on Rendering (Experimental Ideas & Implementations)*. <https://doi.org/10/gfznbd3>

Toshiya Hachisuka, Shinji Ogaki, and Henrik Wann Jensen. 2008. Progressive photon mapping. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 27, 5 (2008), 130:1–130:8. <https://doi.org/10/d8xxn3>

Toshiya Hachisuka, Jacopo Pantaleoni, and Henrik Wann Jensen. 2012. A path space extension for robust light transport simulation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 31, 5 (2012), 191:1–191:10. <https://doi.org/10/gbb6n3>

Vlastimil Havran, Jiri Bittner, Robert Herzog, and Hans-Peter Seidel. 2005. Ray maps for global illumination. In *Rendering Techniques (Proceedings of the Eurographics Symposium on Rendering)*, 43–54. <https://doi.org/10/c2xphk>

Eric Heitz, Stephen Hill, and Morgan McGuire. 2018. Combining analytic direct illumination and stochastic shadows. In *Proceedings of the Symposium on Interactive 3D Graphics and Games (i3D)*. ACM, 2:1–2:11. <https://doi.org/10/gfznb7>

David Immel, Michael Cohen, and Donald Greenberg. 1986. A radiosity method for non-diffuse environments. *Computer Graphics (Proceedings of SIGGRAPH)* 20, 4 (1986), 133–142. <https://doi.org/10/dmjn9t>

Wenzel Jakob. 2013. *Light transport on path-space manifolds*. Ph.D. Dissertation. Cornell University.

Wojciech Jarosz, Derek Nowrouzezahrai, Iman Sadeghi, and Henrik Wann Jensen. 2011a. A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM Transactions on Graphics* 30, 1 (Feb. 2011), 5:1–5:19. <https://doi.org/10/fcdh2f>

Wojciech Jarosz, Derek Nowrouzezahrai, Robert Thomas, Peter-Pike Sloan, and Matthias Zwicker. 2011b. Progressive photon beams. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 30, 6 (Dec. 2011). <https://doi.org/10/fn5xzj>

Wojciech Jarosz, Matthias Zwicker, and Henrik Wann Jensen. 2008. The beam radiance estimate for volumetric photon mapping. *Computer Graphics Forum (Proceedings of Eurographics)* 27, 2 (April 2008), 557–566. <https://doi.org/10/bjfsfx>

Johannes Jendersie. 2018. Path throughput importance weights. [arXiv:1806.01005](https://arxiv.org/abs/1806.01005)

Henrik Wann Jensen. 1996. Global illumination using photon maps. In *Rendering Techniques (Proceedings of the Eurographics Workshop on Rendering)*, 21–30. <https://doi.org/10/fzc619>

Henrik Wann Jensen and Per H. Christensen. 1998. Efficient simulation of light transport in scenes with participating media using photon maps. In *Annual Conference Series (Proceedings of SIGGRAPH)*, 311–320. <https://doi.org/10/b64p36>

Thouis R. Jones and Ronald N. Perry. 2000. Antialiasing with line samples. In *Rendering Techniques (Proceedings of the Eurographics Workshop on Rendering)*, 197–206. <https://doi.org/10/gfznb9>

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

Csaba Kelemen, László Szirmay-Kalos, György Antal, and Ferenc Csonka. 2002. A simple and robust mutation strategy for the Metropolis light transport algorithm. *Computer Graphics Forum* 21, 3 (2002), 531–540. <https://doi.org/10/bfrsqn>

Markus Kettunen, Marco Manzi, Miika Aittala, Jaakko Lehtinen, Frédo Durand, and Matthias Zwicker. 2015. Gradient-domain path tracing. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 34, 4 (July 2015), 123:1–123:13. <https://doi.org/10/10/gfzrhn>

Jaroslav Krivánek, Iliyan Georgiev, Toshiya Hachisuka, Petr Vévoda, Martin Šík, Derek Nowrouzezahrai, and Wojciech Jarosz. 2014. Unifying points, beams, and paths in volumetric light transport simulation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 33, 4 (July 2014), 103:1–103:13. <https://doi.org/10/f6cz72>

Eric Lafortune and Yves Willems. 1993. Bi-directional path tracing. In *Proceedings of Compugraphics*, Vol. 93. 145–153.

Jaakko Lehtinen, Tero Karras, Samuli Laine, Miika Aittala, Frédo Durand, and Timo Aila. 2013. Gradient-domain Metropolis light transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 32, 4 (July 2013), 95:1–95:12. <https://doi.org/10/gbdihd>

Jan Novák, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz. 2018a. Monte Carlo methods for volumetric light transport simulation. *Computer Graphics Forum (Eurographics State of the Art Reports)* 37, 2 (May 2018), 551–576. <https://doi.org/10/gd2jqq>

Jan Novák, Iliyan Georgiev, Johannes Hanika, Jaroslav Krivánek, and Wojciech Jarosz. 2018b. Monte Carlo methods for physically based volume rendering. In *ACM SIGGRAPH Courses*. <https://doi.org/10/c5fj>

Jan Novák, Derek Nowrouzezahrai, Carsten Dachsbacher, and Wojciech Jarosz. 2012a. Progressive virtual beam lights. *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 31, 4 (June 2012), 1407–1413. <https://doi.org/10/gfzndw>

Jan Novák, Derek Nowrouzezahrai, Carsten Dachsbacher, and Wojciech Jarosz. 2012b. Virtual ray lights for rendering scenes with participating media. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 31, 4 (July 2012), 60:1–60:11. <https://doi.org/10/gbbw2k>

Derek Nowrouzezahrai, Ilya Baran, Kenny Mitchell, and Wojciech Jarosz. 2014. Visibility silhouettes for semi-analytic spherical integration. *Computer Graphics Forum* 33, 1 (Feb. 2014), 105–117. <https://doi.org/10/f5t6tf>

Mark Pauly, Thomas Kollig, and Alexander Keller. 2000. Metropolis light transport for participating media. In *Rendering Techniques (Proceedings of the Eurographics*

Workshop on Rendering), 11–22. <https://doi.org/10/gfzm93>

Vincent Pego and Steven G. Parker. 2009. An analytical solution to single scattering in homogeneous participating media. *Computer Graphics Forum (Proceedings of Eurographics)* 28, 2 (2009), 329–335. <https://doi.org/10/c9zhxn>

Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically based rendering: From theory to implementation* (3rd ed.).

Gurprit Singh and Wojciech Jarosz. 2017. Convergence analysis for anisotropic Monte Carlo sampling spectra. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 36, 4 (July 2017), 137:1–137:14. <https://doi.org/10/gbxhfj>

Gurprit Singh, Bailey Miller, and Wojciech Jarosz. 2017. Variance and convergence analysis of Monte Carlo line and segment sampling. *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 36, 4 (June 2017), 79–89. <https://doi.org/10/gfzncj>

Jerome Spanier. 1966. Two pairs of families of estimators for transport problems. *SIAM J. Appl. Math.* 14, 4 (1966), 702–713. <https://doi.org/10/dg35nt>

Bo Sun, Ravi Ramamoorthi, Srinivasa G. Narasimhan, and Shree K. Nayar. 2005. A practical analytic single scattering model for real time rendering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 24, 3 (2005), 1040–1049. <https://doi.org/10/fgnbqt>

Xin Sun, Kun Zhou, Jie Guo, Guofu Xie, Jingui Pan, Wencheng Wang, and Baining Guo. 2013. Line segment sampling with blue-noise properties. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 32, 4 (July 2013), 127:1–127:14. <https://doi.org/10/gbdg4r>

Xin Sun, Kun Zhou, Stephen Lin, and Baining Guo. 2010. Line space gathering for single scattering in large scenes. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 29, 4 (July 2010), 54:1–54:8. <https://doi.org/10/dzvxvr>

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 High Performance Graphics*, 23–31. <https://doi.org/10/gfznqc>

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 (Proceedings of the Eurographics Workshop on Rendering)*, 147–162. <https://doi.org/10/gfznbh>

Eric Veach and Leonidas Guibas. 1995. Optimally combining sampling techniques for Monte Carlo rendering. *Annual Conference Series (Proceedings of SIGGRAPH)* 29 (1995), 419–428. <https://doi.org/10/d7b6n4>

Eric Veach and Leonidas Guibas. 1997. Metropolis light transport. *Annual Conference Series (Proceedings of SIGGRAPH)* 31 (1997), 65–76. <https://doi.org/10/bkjqq4>