Eurographics Symposium on Rendering 2019
T. Boubekeur and P. Sen
(Guest Editors)

Volume 38 (2019), Number 4

Combining Point and Line Samples for Direct Illumination

Katherine Salesin and Wojciech Jarosz
Dartmouth College

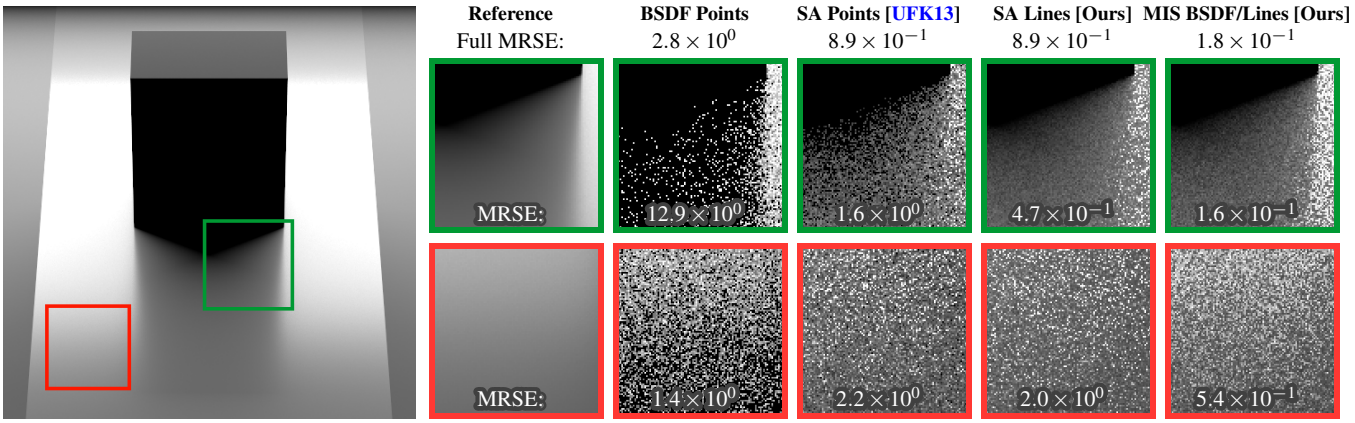

Figure 1: *We compare BSDF and solid-angle (SA) point sampling to our line sampling and its MIS with BSDF sampling at equal time (three seconds). As shown by the mean relative squared error (MRSE), line samples provide a clear benefit in soft shadows, but the glossy material is handled better by MISing additionally with BSDF sampling.*

Abstract

We develop a unified framework for combining point and line samples in direct lighting calculations. While line samples have proven beneficial in a variety of rendering contexts, their application in direct lighting has been limited due to a lack of formulas for evaluating advanced BRDFs along a line and performance tied to the orientation of occluders in the scene. We lift these limitations by elevating line samples to a shared higher-dimensional space with point samples. Our key insight is to separate the probability distribution functions of line samples and points that lie along a line sample. This simple conceptual change allows us to apply multiple importance sampling (MIS) between points and lines, and lines with each other, in order to leverage their respective strengths. We also show how to improve the convergence rate of MIS between points and lines in an unbiased way using a novel discontinuity-smoothing balance heuristic. We verify through a set of rendering experiments that our proposed MISing of points and lines, and lines with each other, reduces variance of the direct lighting estimate while supporting an increased range of BSDFs compared to analytic line integration.

CCS Concepts

• ***Computing methodologies*** → ***Ray tracing; Visibility;*** • ***Mathematics of computing*** → *Stochastic processes;*

1. Introduction

Despite decades of progress, physically accurate direct illumination remains a challenging problem. At any shading point in the scene, direct illumination requires computing a two-dimensional integral of incident light modulated by the material's reflectance properties over the visible regions of the light sources. Since this integral cannot be computed analytically in the general setting, it is typically approximated numerically. Monte Carlo integration has prevailed as the dominant approach because it can approximate integrals while only requiring the ability to *evaluate* the integrand at a number of sample point locations. Traditionally, evaluating such samples correspond to tracing shadow rays. The price to be paid for this flexibility is that Monte Carlo integration suffers from noise and many samples are often necessary to reduce this noise in the presence of complex visibility discontinuities and materials.

Recently, Billen and Dutré [BD16] proposed an alternative approach which poses direct illumination not as a 2D Monte Carlo point sampling problem, but as a 1D Monte Carlo line sampling problem. Such line samples can dramatically improve the quality of soft shadows, and can even provide an improved rate of convergence. Unfortunately, direct lighting with line samples is conceptually and computationally more challenging since evaluating each sample in

fact requires deriving and computing a 1D analytic integral for each type of light/material combination. Since each line sample must be evaluated analytically, the technique currently only supports Lambertian and Phong BSDFs, and remains incompatible with more complex microfacet-based materials. Additionally, since the idea is relatively recent, line sampling does not yet have the same array of powerful importance sampling strategies as point sampling. Unfortunately, while it would be desirable to combine the strengths of line samples with the flexibility and generality of point samples for direct illumination, this combination remains elusive since the two classes are expressed in different mathematical frameworks.

Contributions. In this paper we take the first steps to bridge the gap between these two forms of Monte Carlo integration. One of our main contributions is showing that Monte Carlo line sampling can be viewed equivalently as Monte Carlo point sampling with a joint probability density function that *importance samples the visibility function*. Our theory allows us to leverage lines in the presence of arbitrary materials while retaining their variance reduction in soft shadows (see Fig. 1). Since line sampling is just a form of point sampling in our framework, we adapt recent methods for point sampling [UFK13] to distribute lines proportionally to solid angle, reducing variance even compared to Billen and Dutré's [BD16] analytic line evaluation. Since points and lines of all orientations share a common sampling space, we can combine lines with points, and differently oriented lines with each other, to form more robust estimators using multiple importance sampling (MIS) [VG95b]. Moreover, we show how to inject the additional visibility information gained from line samples into a discontinuity-smoothing MIS heuristic which we found allows us, in many cases, to retain line sampling's improved convergence rate even when combined with point sampling. We implement our approach in PBRT, where we outperform both point sampling and analytic line sampling on scenes with simple materials, and we trivially extend to scenes with arbitrary materials not previously supported by analytic integration.

2. Related Work

Since Cook et al. [CPC84] introduced MC integration to graphics, the vast majority of MC rendering research has focused on point sampling, resulting in a vast arsenal of advanced (importance) sampling techniques for a variety of rendering sub-problems. We refer to standard texts [PJH16] on the subject, and instead focus on the less explored area of line sampling and its combination with points.

MC line sampling and analytic integration in rendering. Recently, researchers have increasingly started employing Monte Carlo-like estimators using line or segment samples for rendering problems as diverse as anti-aliased scan-conversion [JP00; Max86; Max90], motion blur [GDA10], depth of field [TPD*12], hair rendering [BGA12], volumetric [BJ17; GKH*13; HCJ13; JNSJ11; JNT*11; JZJ08; KGH*14; NNDJ12a; NNDJ12b; SZLG10] and transient [JMM*14; MGJ*19; MJGJ17] light transport, and soft shadows from hemispherical [GBA11], area [BD16] or environment [NBMJ14] lights. We build most directly on the work of Billen and Dutré [BD16], who proposed an efficient intersection approach to evaluate shadow line samples, and coupled this with an analytic integration of the shading integrand along visible line segments. We directly reuse their intersection routine, but instead of using their analytic integration, cast line sampling as an MC point sampling problem for increased flexibility.

While such higher-dimensional (nD) samples often provide an impressive reduction in integration error, they are typically more expensive since *evaluating* each such sample ultimately requires actually *integrating* these n dimensions. nD sampling can therefore leverage and provide a theoretical bridge to analytic solutions to rendering sub-problems, such as the so-called airlight integral [PP09; SRNN05], or analytic lighting on diffuse and glossy surfaces [Arv95b; Arv95c; BP93; BXH*18; CA00; CA01; NON85; Pic92; WR18]. Unfortunately, nD sampling is limited since it can only be applied to problems that admit analytic evaluations, and, despite recent progress [Hei18], it does not currently offer the same arsenal of importance sampling techniques amassed for points. We offer a practical way to address these problems by reinterpreting line samples in direct illumination as equivalent point samples.

Combining with point sampling. Multiple importance sampling (MIS) [VG95b] has proven to be one of the single most important tools for robustly leveraging several competing point sampling strategies in MC integration. Unfortunately, since nD samples and point samples operate in different dimensional spaces, it is not possible to directly combine them using MIS. Researchers have therefore resorted to combining such analytic nD samples with points either via unbiased control variates [BXH*18] or biased ratio estimators [HHM18]. Prior work [BD16; SMJ17] has shown that different orientations of nD samples (e.g. horizontal vs. vertical lines on an area light source) can also have dramatically different variances based on how the samples align with features of the integrand. It would therefore be desirable to leverage multiple line sampling strategies for increased robustness. Unfortunately, such differently oriented lines operate on different effective integrands so they cannot be directly combined with MIS. We solve these problems by elevating all line samples to equivalent point samples in a shared higher-dimensional space. This allows robustly combining lines with each other and with arbitrary point sampling strategies using MIS. This is conceptually similar to (though much simpler than) extending [GKDS12; HGJ*17; HPJ12] standard path space to enable MISing (bidirectional [LW93; VG95a]) path tracing [ICG86; Kaj86] and photon mapping [Jen01; Jen96] techniques.

Variance and convergence analysis. Recent work [Dur11; PSC*15; SJ17; SK13; SNJ*14; SSC*19] has established a firm mathematical connection between the properties of MC point sampling and the *magnitude* and *convergence rate* of MC integration error. The error depends on the dimensionality and smoothness of the integrand, and the spectral properties of the MC sampling pattern, with various flavors of stratified sampling [Coo86; CSW94; Ken13] provably leading to asymptotically faster convergence rates. Singh et al. [SMJ17] established a theoretical connection between these point sampling results and the observed behavior of nD sampling by noting that nD samples 1) reduce the dimensionality of the integration problem, and 2) produce a smoothed effective integrand, both of which can contribute to an improved convergence rate. Most recently, Singh et al. [SSC*19] noted that seemingly innocuous changes like BSDF vs. solid-angle or surface-area sampling can lead to different

convergence rates due to the introduction or inhibition of discontinuities. This becomes important when combining strategies since MIS inherits the worst of the constituent convergence rates. To mitigate this, they propose a biased discontinuity-smoothing approach to retain a good convergence rate, though only for unoccluded integrands. We propose an unbiased way to improve convergence rates by manipulating the MIS weights, and additionally show how leveraging line samples allows us to extend these improvements to a limited extent to previously unimproved penumbra regions.

3. Background & Problem Statement

Direct Lighting. The outgoing radiance L_o at a shadepoint $\mathbf{x}$ towards the eye $\mathbf{e}$ is the integral of the radiance emitted L_e from all points $\mathbf{l}$ on the surface area $\mathcal{A}$ of a light modulated by the BSDF f_r, geometry term G, and binary visibility function V (see inset below):

$$L_o(\mathbf{e},\mathbf{x}) = \int_{\mathcal{A}} f_r(\mathbf{e},\mathbf{x},\mathbf{l})\, G(\mathbf{x},\mathbf{l})\, V(\mathbf{x},\mathbf{l})\, L_e(\mathbf{x},\mathbf{l})\, \mathrm{d}A(\mathbf{l}), \quad \text{where} \quad (1)$$

$$G(\mathbf{x},\mathbf{l}) = \frac{(\vec{\omega}_{\mathbf{xl}} \cdot \vec{n}_{\mathbf{x}})^+ (\vec{\omega}_{\mathbf{lx}} \cdot \vec{n}_{\mathbf{l}})^+}{\|\mathbf{x}-\mathbf{l}\|^2}. \quad (2)$$

We use bold symbols ($\mathbf{x}$) to denote points, arrows ($\vec{\omega}$) to decorate unit direction vectors, and $()^+$ clamps the enclosed quantity to zero. In the general case, Eq. (1) is impractical to compute in closed form, so we resort to approximating it using Monte Carlo.

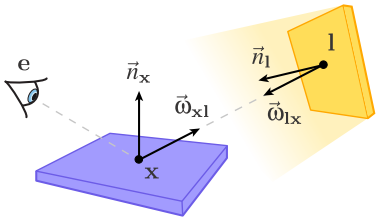

Abstracting away the details, we can write Eq. (1) more compactly as a simple two-dimensional integration problem:

$$L_o = \int_{\mathcal{U}} \int_{\mathcal{V}} f(u,v)\, \mathrm{d}v\, \mathrm{d}u, \quad (3)$$

$$\text{with} \quad f(u,v) = f_r(\mathbf{e},\mathbf{x},\mathbf{l}_{uv})\, G(\mathbf{x},\mathbf{l}_{uv})\, V(\mathbf{x},\mathbf{l}_{uv})\, L_e(\mathbf{x},\mathbf{l}_{uv}), \quad (4)$$

where for conceptual simplicity we also assume a coordinate system aligned with the light so $\mathbf{l}_{uv} = (u,v,0)$ and $\vec{n}_{\mathbf{l}} = (0,0,1)$, see Fig. 2.

Monte Carlo Integration. In this simplified form, an MC estimator of direct illumination chooses N random points u_i, v_i according to some joint probability density function (PDF) $p(u_i,v_i)$ and averages the integrand evaluated at these locations weighted by the inverse PDF:

$$L_o = \int_{\mathcal{U}} \int_{\mathcal{V}} f(u,v)\, \mathrm{d}v\, \mathrm{d}u \quad \approx \quad \langle L_o \rangle_{uv} = \frac{1}{N} \sum_{i=1}^{N} \frac{f(u_i,v_i)}{p(u_i,v_i)}. \quad (5)$$

The efficiency of such MC estimators depends critically on the sampling PDF. Ideally we want to choose a PDF that is exactly proportional to the integrand, but this is generally not available. A common strategy is therefore to choose a PDF that is proportional to some portion of the integrand, for instance proportional to the BSDF, the geometry term (solid-angle sampling), or the emitted radiance (surface-area sampling). Each of these choices provides variance reduction in complementary settings. To obtain the benefits of S different PDF strategies with N_s samples each, it is tempting to simply average their estimators,

$$\langle L_o \rangle_{uv}^{\text{avg}} = \frac{1}{S} \sum_{s=1}^{S} \left(\frac{1}{N_s} \sum_{i=1}^{N_s} \frac{f(u_i,v_i)}{p_s(u_i,v_i)} \right), \quad (6)$$

but since variance is additive, any sample with high variance will pollute the averaged result.

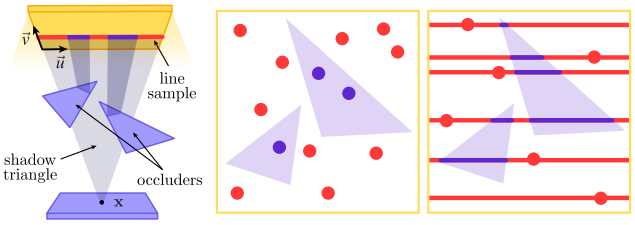

Figure 2: Left: *Setup of the general line sampling algorithm. First, a line sample is chosen along the $\vec{v}$ axis. Next, a shadow triangle connecting the line sample and shading point is intersected with the scene and shadowed regions are culled from the line sample.* Middle to right: *Visualization of point and horizontal line sampling strategies over the light source when there are occluders, whose shadow projections are shown. Traditional point samples may land in occluded regions, whereas with line samples we can ensure points are generated only in unoccluded regions.*

Multiple importance sampling. A more robust strategy is to downweight each potentially high-variance sample *before* it has a chance to pollute the average, using MIS [VG95b]:

$$\langle L_o \rangle_{uv}^{\text{mis}} = \frac{1}{S} \sum_{s=1}^{S} \left(\frac{1}{N_s} \sum_{i=1}^{N_s} w_s(u_i,v_i) \frac{f(u_i,v_i)}{p_s(u_i,v_i)} \right). \quad (7)$$

A popular and provably good choice of weight function is the *balance heuristic*:

$$w_s(u_i,v_i) = \frac{N_s p_s(u_i,v_i)}{\sum_j N_j p_j(u_i,v_i)}. \quad (8)$$

Effectively, MIS assumes a sampling strategy is poor at locations where its PDF is low compared to other strategies. Crucially, MIS requires that the estimators all compute the *same integral*, and the PDFs are expressed wrt a *common measure* before comparison.

Line sampling. Instead of sampling points, Billen and Dutré [BD16] recently proposed to sample lines along the light and integrate the contribution along their lengths analytically. There is a continuum of possible line directions. For instance, using the notation of Eq. (3), horizontal line sampling computes

$$L_o = \int_{\mathcal{V}} f_u(v)\, \mathrm{d}v \quad \approx \quad \langle L_o \rangle_v = \frac{1}{N} \sum_{i=1}^{N} \frac{f_u(v_i)}{p(v_i)}, \quad (9)$$

while vertical line sampling computes

$$L_o = \int_{\mathcal{U}} f_v(u)\, \mathrm{d}u \quad \approx \quad \langle L_o \rangle_u = \frac{1}{N} \sum_{i=1}^{N} \frac{f_v(u_i)}{p(u_i)}, \quad (10)$$

$$\text{where} \quad f_u(v) = \int_{\mathcal{U}} f(u,v)\, \mathrm{d}u, \quad \text{and} \quad f_v(u) = \int_{\mathcal{V}} f(u,v)\, \mathrm{d}v \quad (11)$$

are one-dimensional integrands which arise from analytically pre-integrating $f(u,v)$ along u or v respectively.

Billen and Dutré's approach first samples a line offset, e.g. according to $p(v_i)$ for horizontal lines. It then traces a "shadow triangle" through the scene (see Fig. 2) to determine all intervals along the line sample where $V = 1$, and integrates the remaining portions of $f_u(v_i)$ analytically along each interval.

Billen and Dutré [BD16] explained the improved performance of line samples by virtue of their reduced MC integration dimensionality, and Singh et al. [SMJ17] later explained this from a frequency perspective by interpreting lines as point sampling a lower-dimensional, smoothed integrand. In either view, lines provide impressive variance reduction (and even convergence rate improvement), but they are limited to material and lighting properties that admit efficient analytic line integration. Unfortunately, neither interpretation allows us to combine with MIS the benefits of line sampling in different orientations or point importance sampling strategies. This is because the integrands for different line orientations (11) are different from one another, and are different than the 2D point sampling integrand (4). This is what we address.

4. Method

Our main contribution is unifying point sampling and line sampling into a common sampling space so that their relative strengths can be combined using MIS. We use a horizontal line as an example below, but the same ideas apply to vertical lines by swapping u and v.

The 2D point-sampling integrand (5) is clearly different from the 1D line-sampling integrands (11) since the former requires two random numbers, while the latter only require one. Conceptually, however, we know that these estimators are actually computing the same quantity. Our key insight is that we can interpret an analytic 1D horizontal line sample as a 2D point sample with joint PDF $p(u_i,v_i) = p(v_i)p(u_i \mid v_i)$ that uses some marginal PDF $p(v_i)$ to choose a line offset v_i and then the *ideal conditional PDF*

$$p(u_i \mid v_i) = \frac{f(u_i,v_i)}{\int_{\mathcal{U}} f(u,v_i)\,\mathrm{d}u} = \frac{f(u_i,v_i)}{f_u(v_i)} \tag{12}$$

to choose a location u_i along the line. Inserting this into Eq. (5), reduces it to Eq. (9):

$$\langle L_\mathrm{o}\rangle_{uv} = \frac{f(u_i,v_i)}{p(v_i)p(u_i \mid v_i)} = \frac{f(u_i,v_i)}{p(v_i)\frac{f(u_i,v_i)}{f_u(v_i)}} = \frac{f_u(v_i)}{p(v_i)} = \langle L_\mathrm{o}\rangle_v, \tag{13}$$

where we have omitted the averaging across N samples for brevity.

By definition, if an MC sample is drawn from an ideal conditional PDF, it would return the same (exact) result for the conditional integration along $\mathcal{U}$ regardless of the actual sampled location u_i; hence, in some sense this collapses the dimensionality of the integration as the u component of the 2D sample is never needed. While Eq. (13) shows that using $p(v_i)$ in Eq. (9) or $p(v_i)p(u_i \mid v_i)$ in Eq. (5) is mathematically equivalent, in practice the former only requires analytically integrating along the line, but the latter additionally requires generating a point sample along the line proportional to the integrand. Importantly, only the latter technique provides an effective 2D PDF which we can use to MIS with point sampling or with vertical lines.

While Billen and Dutré [BD16] derived expressions for analytically integrating the geometry term multiplied by a diffuse or Phong [Pho75] BSDF along the line, the expressions are quite involved and producing a point sample proportional to these functions would further require inverting them. Instead, we can still benefit from a line sample by coupling it with simpler, non-ideal conditional PDFs, which do not perfectly cancel out all terms of the integrand along the line, but do leverage a line's core strengths.

Interpreted as a 2D importance sampling strategy, the reason line sampling performs well is because it *importance samples the binary visibility function* along the line sample (see Fig. 2). For instance, for a given v_i, $p(u_i \mid v_i)$ only generates a u_i in regions where $V = 1$ (except in cases where the entire line sample is occluded). This has the benefit of focusing samples in the visible regions (reducing variance), but also eliminating the discontinuities in the conditional integrand (improving smoothness and convergence). By explicitly sampling points along a line using simple conditional PDFs, we retain this key strength. We propose two such PDFs in the following subsection, which we show are effective despite their imperfection.

4.1. Joint line importance sampling.

Uniform surface-area-based line sampling. Billen and Dutré [BD16] proposed sampling a line's offset proportional to its length. On a quad light, this corresponds to sampling v such that:

$$p(v_i) = \frac{\overline{u_i} - \underline{u_i}}{\text{area of light source}}, \tag{14}$$

where $[\underline{u_i}, \overline{u_i}]$ are the u-bounds of line sample v_i. A simple approach would be to couple this with uniform sampling of u within the visible intervals, which results in the following PDF:

$$p(u_i \mid v_i) = \begin{cases} \dfrac{1}{\sum\limits_{r \in \mathcal{R}_i} \overline{u_r} - \underline{u_r}} & \text{if } u_i \text{ is inside a visible region,} \\ 0 & \text{otherwise,} \end{cases} \tag{15}$$

where $\mathcal{R}_i$ is the set of all visible regions along line sample v_i and $[\underline{u_r}, \overline{u_r}]$ are the u-bounds of visible region r. For unoccluded shade points this corresponds to uniform surface area point sampling. Occluded regions, however, benefit from the line sample only generating corresponding points within the visible region of the integrand. Such penumbra regions improve considerably with line samples compared to surface-area point sampling, but neither approach importance samples the BSDF or geometry terms, which means they both suffer from high variance for glossy BSDFs or shade points close to a light source (see Fig. 3).

Uniform solid-angle-based line sampling. Uniformly sampling the solid angle of the light source typically performs much better than surface area sampling since it accounts not only for L_e, but also the inverse-squared distances and one of the cosines in the geometry term (2). We adapt Ureña et al.'s [UFK13] approach to line samples, allowing us to additionally importance sample the visibility.

Ureña et al.'s approach projects a quad light onto a unit sphere surrounding the shade point $\mathbf{x}$ and then samples the resulting spherical rectangle $\mathcal{Q}$ in a two-step process similar to ours (see inset figure).

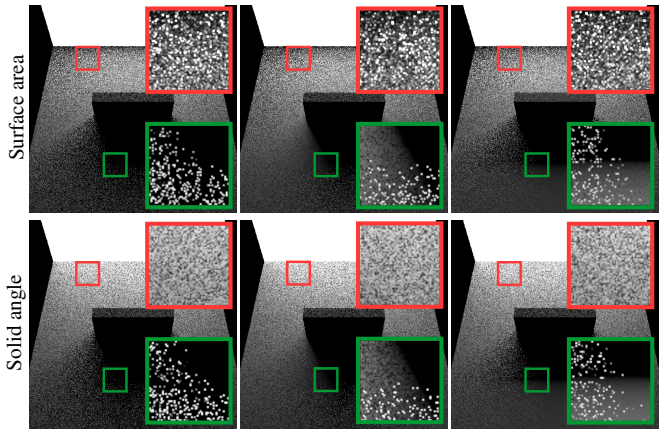

(a) Points (b) Horizontal lines (c) Vertical lines

Figure 3: *Even at one sample per pixel, our extension of solid angle sampling to line samples dramatically reduces noise in soft shadows (red) while simultaneously suppressing variance from the geometry term near the light source (blue).*

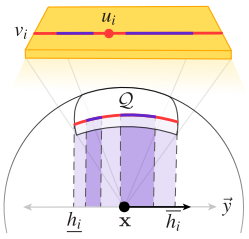

First, they choose line offset v_i, yielding an arc on $\mathcal{Q}$. Then, they choose u_i within line v_i by projecting the endpoints of the arc onto $\vec{y}$ (an axis aligned with v_i) to yield bounds $[\underline{h_i}, \overline{h_i}]$ and linearly interpolating between them.

We modify this approach to incorporate our line sampling routine and account for occlusions. Given a line offset v_i chosen by their algorithm, we intersect the shadow triangle formed by v_i and $\mathbf{x}$ with the scene to find occluded regions. We then project the endpoints $[\underline{u_r}, \overline{u_r}]$ of all visible regions onto $\mathcal{Q}$ to form a set of arcs. Projecting the endpoints of each arc onto $\vec{y}$ yields a set of height bounds $[\underline{h_r}, \overline{h_r}]$ for all visible regions, within which we sample h_i uniformly.

Analogously to surface-area-based line sampling, we can separate the joint PDF of this strategy into a marginal and conditional PDF. While Ureña et al. provide sampling routines from which we could derive both the conditional and marginal PDF, we already know the joint PDF $p(u_i, v_i) = 1/\Omega(\mathcal{Q})$, where $\Omega(\mathcal{Q})$ is the solid angle of $\mathcal{Q}$, in the unoccluded case. The marginal PDF is therefore

$$p(v_i) = \frac{p(u_i, v_i)}{p_{\text{vis}}(u_i \mid v_i)} = \frac{1}{\Omega(\mathcal{Q}) p_{\text{vis}}(u_i \mid v_i)}, \tag{16}$$

where $p_{\text{vis}}(u_i \mid v_i)$ denotes a conditional PDF assuming full visibility. Given a v_i, Ureña et al. uniformly sample an h_i over the entire $[\underline{h_i}, \overline{h_i}]$ interval and map this h_i to a u_i, leading to

$$p_{\text{vis}}(u_i \mid v_i) = \frac{p_{\text{vis}}(h_i \mid v_i)}{J}, \tag{17}$$

where $J = \frac{\mathrm{d}u_i}{\mathrm{d}h_i}$ is the h-to-u Jacobian, and $p_{\text{vis}}(h_i \mid v_i) = \frac{1}{\overline{h_i} - \underline{h_i}}$.

To account for the occluders, we instead sample h_i uniformly within the *visible* h-intervals and then map h_i to u_i as before. The conditional PDF in the occluded case is

$$p_{\text{occ}}(u_i \mid v_i) = \frac{p_{\text{occ}}(h_i \mid v_i)}{J}, \tag{18}$$

where $p_{\text{occ}}(h_i \mid v_i)$ sums over all visible interval lengths as in Eq. (15), but with h substituted for u. While J is trivial to find, it is unnecessary to compute in practice due to the problem's construction. Multiplying Eq. (18) by Eq. (16) and expanding gives us our joint PDF in the occluded scenario,

$$p(u_i, v_i) = \frac{J(\overline{h_i} - \underline{h_i})}{\Omega(\mathcal{Q}) J \sum_{r \in \mathcal{R}_i} \overline{h_r} - \underline{h_r}} = \frac{\overline{h_i} - \underline{h_i}}{\Omega(\mathcal{Q}) \sum_{r \in \mathcal{R}_i} \overline{h_r} - \underline{h_r}}, \tag{19}$$

where we have omitted the zero case for brevity. For unoccluded regions, $p(u_i, v_i)$ simply reduces to $1/\Omega(\mathcal{Q})$ just like point sampling [UFK13]. Note that this joint PDF is exactly like that of surface-area sampling, except with solid angle substituted for area and h substituted for u. It can be converted to surface area measure by the usual geometry Jacobian, $\frac{\vec{\omega}_{\mathbf{lx}} \cdot \vec{n}_{\mathbf{l}}}{\|\mathbf{x}-\mathbf{l}\|^2}$.

4.2. MIS between lines and points

Solid-angle line sampling works reasonably well for diffuse scenes. However, two issues remain: 1) the relative alignment of the line samples with the occluder dramatically influences the benefits in penumbra regions, and 2) with any orientation of lines, the variance increases as the BSDF becomes more glossy since it is not importance sampled. To mitigate both of these issues, and arrive at a more robust estimator, we leverage our point-sampling interpretation of lines to MIS differently oriented lines with each other and with BSDF point importance sampling.

In order to MIS between point and line strategies, we need the ability to evaluate their PDFs given a point $\mathbf{l}_{uv}$ on the light source. For point sampling, such formulas are well-known. For line samples, we evaluate $p(v_i)p(u_i \mid v_i)$ (Eqs. (14)–(18)). A drawback to this strategy is that computing MIS weights requires evaluating $p(u_i \mid v_i)$, which requires a line sample-scene intersection to find occluded regions of v_i; however, in practice, we have found the variance reduction achieved by MIS favorable even in equal-time scenarios.

4.3. Discontinuity-smoothing MIS heuristic

Though MIS allows us to combine line and point sampling while mitigating the high variance of any individual strategy, the resulting estimator will take on the worst of the constituent estimators' convergence rates. When using uncorrelated random sampling, each estimator has a convergence rate of $\mathcal{O}(N^{-1})$, so this is not an issue. However, stratified sampling can improve the convergence rate of 2D integration to $\mathcal{O}(N^{-1.5})$ for discontinuous integrands and $\mathcal{O}(N^{-2})$ for smooth integrands [PSC*15]. Singh et al. [SSC*19] recently showed that BSDF vs. surface-area sampling achieve different convergence rates because the former treats the light boundary as an integrand discontinuity, while the latter pushes this discontinuity to the boundary of the integration domain, allowing unoccluded shade points to "see" a smooth integrand. They proposed a biased solution to artificially smooth the light source boundary for BSDF-based sampling, allowing it, and the MIS of both strategies, to retain an improved $\mathcal{O}(N^{-2})$ convergence, but only for unoccluded shade points. We propose to solve the same problem in an unbiased way, and can use the additional information provided by line sample to extend a limited improvement even for partially occluded integrands.

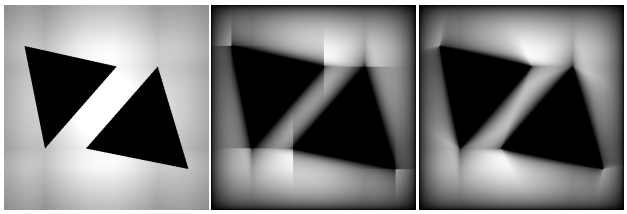

(a) no smoothing (b) fixed-width smoothing (c) variable-width smoothing

Figure 4: *MIS weights over a quad light source with two triangle occluders as seen by a BSDF sampling routine without smoothing (left), with fixed-width discontinuity smoothing (middle), and with variable-width discontinuity smoothing (right). With fixed-width smoothing, the discontinuities at the light and occluder boundaries are smoothed away in the MIS weights. With variable-width smoothing, discontinuities at occluder vertices are also smoothed.*

We begin by rewriting Eq. (7) as:

$$\langle L_o \rangle_{uv}^{\text{mis}} = \frac{1}{S}\sum_{s=1}^{S}\left(\frac{1}{N_s}\sum_{i=1}^{N_s}\frac{f_s^{\text{eff}}(u_i,v_i)}{p_s(u_i,v_i)}\right), \quad \text{where} \tag{20}$$

$$f_s^{\text{eff}}(u_i,v_i) = w_s(u_i,v_i) f(u_i,v_i) \tag{21}$$

is the effective integrand "seen" by strategy s during MIS. Our key insight is that if we could design the MIS weights to ensure the effective integrand seen by each Monte Carlo integration strategy is smooth, then we could obtain better convergence rates (as long as the PDFs p_s are themselves not discontinuous).

Let us consider three sampling strategies: horizontal (V) and vertical (U) lines, as well as BSDF sampling B, with corresponding MIS weighting functions w_s for $s \in \{U,V,B\}$. While strategies U and V do not suffer from a discontinuity due to the light source boundary, we'd like to remove the discontinuity from the BSDF strategy B by having w_B smoothly transition to zero for points $\mathbf{l}_{uv}$ that approach the light boundary. Doing so, the effective integrand would appear to be a light source with a smoothly blurred boundary. We can accomplish this by rewriting the balance heuristic (8) as

$$w_s(u_i,v_i) = \frac{N_s p'_s(u_i,v_i)}{\sum_j N_j p'_j(u_i,v_i)} \tag{22}$$

where $p'_s = p_s$ for $s \in \{U,V\}$, but for the BSDF strategy we will multiply it by a function that smoothly decreases to zero as u and v approach the edge of the light: $p'_B(u_i,v_i) = \hat{m}(u_i)\hat{m}(v_i)p_B(u_i,v_i)$. Normalizing by some distance parameter d_m, we have:

$$\hat{m}(u_i) = m\left(\frac{u_i - \underline{u_i}}{d_m}\right)\cdot m\left(\frac{\overline{u_i} - u_i}{d_m}\right) \tag{23}$$

and likewise for v_i, where $[\underline{u_i}, \overline{u_i}]$ are the u-bounds of the light source and m is a quadratic ease-out curve: $m(t) = -\langle t\rangle^2 + 2\langle t\rangle$ where $\langle t\rangle$ is t clamped to $[0,1]$. We found using d_m equal to half the smallest dimension of the light to be effective.

By intersecting a line sample with the scene, we can also detect how close a point sample is to an occluder. In this case, we replace $\underline{u_i}$ and $\overline{u_i}$ with $\underline{u_r}$ and $\overline{u_r}$ in Eq. (23) for the visible region r it lies in. While this accounts for most discontinuities, there may also be discontinuities lurking at occluder vertices where a line slips on or off the occluder altogether (see Fig. 4b). To account for this, we shrink distance d_m as the line approaches an occluder vertex. We use the length of the occluded regions neighboring each visible region boundary, $\underline{d_{\text{occ}}}$ and $\overline{d_{\text{occ}}}$, to determine the amount by which to shrink the smoothstep, such that:

$$\hat{m}(u_i) = m\left(\frac{u_i - \underline{u_r}}{m\left(\underline{d_{\text{occ}}}/d_m\right)}\right)\cdot m\left(\frac{\overline{u_r} - u_i}{m\left(\overline{d_{\text{occ}}}/d_m\right)}\right) \tag{24}$$

for occluder edges. We multiply the smoothsteps of the light boundaries and the occluder edges to achieve all-around smoothness. Fig. 4 visualizes how this influences the MIS weights for a partially occluded shade point.

5. Results

We implemented our novel algorithms as an integrator in PBRT [PJH16] and perform variance analysis using the empirical error analysis framework of Subr et al. [SSJ16]. We compare our results in a variety of scenes to state-of-the-art line sampling [BD16] and solid angle point sampling [UFK13].

MIS between line directions. We begin by testing the efficacy of multiple importance sampling between line directions. We illustrate a simple best- and worst-case scenario with a test scene of a quad floating in a Cornell box beneath a quad light (Fig. 5). For both the MIS and average estimators we use a multijittered sampler [CSW94] and allocate half the total samples to each component strategy by splitting and rescaling. Each pixel we analyze has a completely continuous integrand in one line direction (best-case) and a sharply discontinuous one in the other (worst-case).

We can see that the convergence rate of the best-case line direction, $\mathcal{O}(N^{-3})$, is steeper than the worst case, $\mathcal{O}(N^{-2})$, and that both the MIS and average estimators inherit the worse convergence rate of the two, as expected. However, while the variance of the average estimator is tied closely to the worse line direction, we can see that MIS is able to act as a moderating influence and find a middle ground between the two.

MIS between lines and points. Next, we test the efficacy of multiple importance sampling between line sampling and microfacet BSDF point sampling. The scene in Fig. 6 shows a row of occluders sitting on glossy strips with microfacet BSDFs of varying roughness. We again use a multijittered sampler and split and rescale the samples evenly between the strategies. We analyze the convergence of several pixels, one with a rough BSDF and one with a very glossy BSDF. Regardless of material glossiness, our MIS approach is more robust than BSDF point or solid-angle line sampling alone, and consistently achieves a lower variance than the average estimator.

Discontinuity-smoothing MIS heuristic. We verify that our discontinuity-smoothing MIS heuristic improves the convergence rate of MIS using stratified sampling in Fig. 7. For the unoccluded pixel (A), which benefits only from smoothed edges of the light, the convergence rate improves from $\mathcal{O}(N^{-1.4})$ without to $\mathcal{O}(N^{-2.04})$ with smoothing, almost achieving the same convergence rate as lines alone. For the occluded pixel (B), which benefits from smoothed

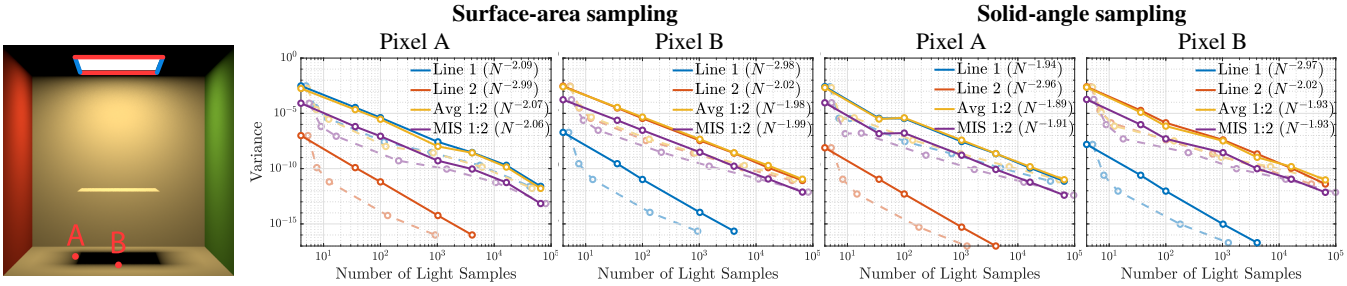

Figure 5: *We illustrate the benefits of MIS between best- and worst-case line directions for both surface-area lines (left) and solid-angle lines (right). Line 1 and Line 2 are two perpendicular line directions on the light source, drawn in blue and red, respectively, in the image. MIS between the two directions consistently finds a middle ground and outperforms the average strategy. Dashed lines show average time over all trials to evaluate* n *light samples.*

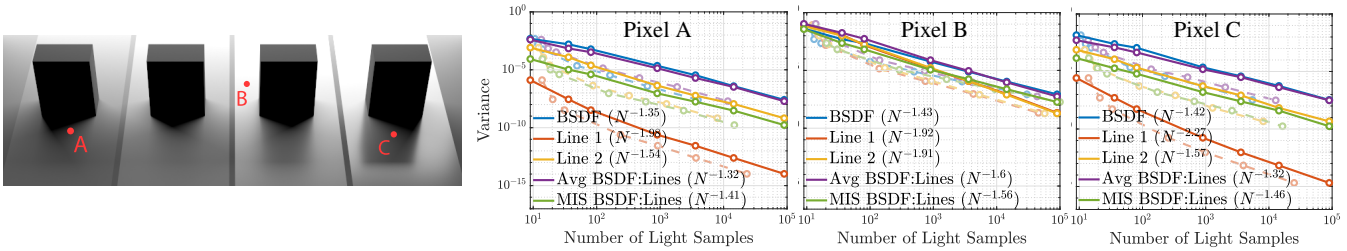

Figure 6: *We illustrate the benefits of MIS between a microfacet BSDF point strategy and solid-angle line strategy, where Line 1 and Line 2 are two perpendicular line directions on the light source. MIS consistently outperforms BSDF sampling or the worse of the two line directions alone, as well as the straight average of the strategies. Dashed lines show average time over all trials to evaluate* n *light samples.*

edges of both the light and occluder, the improvement is more modest, from $\mathcal{O}(N^{-1.53})$ without to $\mathcal{O}(N^{-1.68})$ with smoothing.

Equal-time comparisons. Lastly, we compare our novel line sampling framework to existing line and point sampling schemes at equal time. Our primary metric for comparison is the mean relative square error (MRSE), defined as $\frac{1}{n}\sum_{i=1}^{n}(p_i-\hat{p}_i)^2/(\hat{p}_i^2+\varepsilon)$, where p_i and $\hat{p}_i$ are the i-th pixels in the approximate and reference images, respectively, and $\varepsilon = 1\times10^{-6}$ avoids division by zero.

For these comparisons, we use a random sampler with one light sample per scene intersection. We use an equal-render time of two

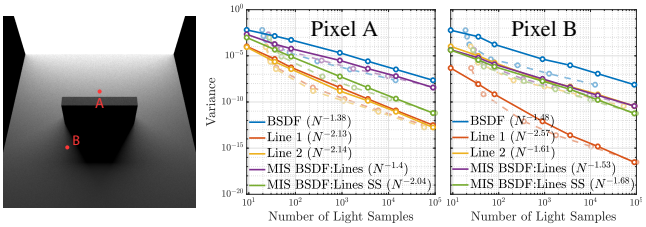

Figure 7: *Our discontinuity-smoothing balance heuristic allows us to retain the variance reduction from line sampling and the improved convergence rate from stratification [CSW94] for unoccluded (A) pixels. Partially occluded pixels (B) also see a slightly improved convergence rate. Dashed lines show average time over all trials to evaluate* n *light samples.*

seconds for the Cornell Box scene (Fig. 8), three seconds for the Monoliths scene (Fig. 1), and four seconds for the Veach-inspired MIS scene (Fig. 10). In Fig. 8, we also compare to Billen and Dutré's analytic line integration scheme, and find that at equal time our method outperforms their approach in the zoomed regions. Our method also consistently achieves lowest MRSE in shadowed regions, though in unshadowed regions, faster strategies can often evaluate more samples in the allotted time.

In Fig. 9, we compare MISing point and line samples with different ratios in equal-time renders (150s) of a kitchen scene with a lowering window shade. In several cases, we see that tweaking the ratio slightly leads to a clear tradeoff between reducing variance in the shiny table region and reducing variance in the towel's shadow region, even though the overall MRSE of the image changes little. This illustrates that the relative effectiveness of different sampling strategies may vary over a diverse scene, yet MIS is successful in bridging the difference in the general case. The ratio 1 points : 3 lines, in particular, consistently achieves the lowest MRSE over the entire image, which suggests that line samples are key players even in an equal-time scenario with a complex scene.

Additional results. While our current implementation does not support line-sampling textured emittered, our MIS formulation allows us to combine our line samples with point samples that do importance sample the texture (Fig. 11). The supplemental material contains additional results, including equal-sample and five-row versions of Fig. 9, and an interactive comparison of the results.

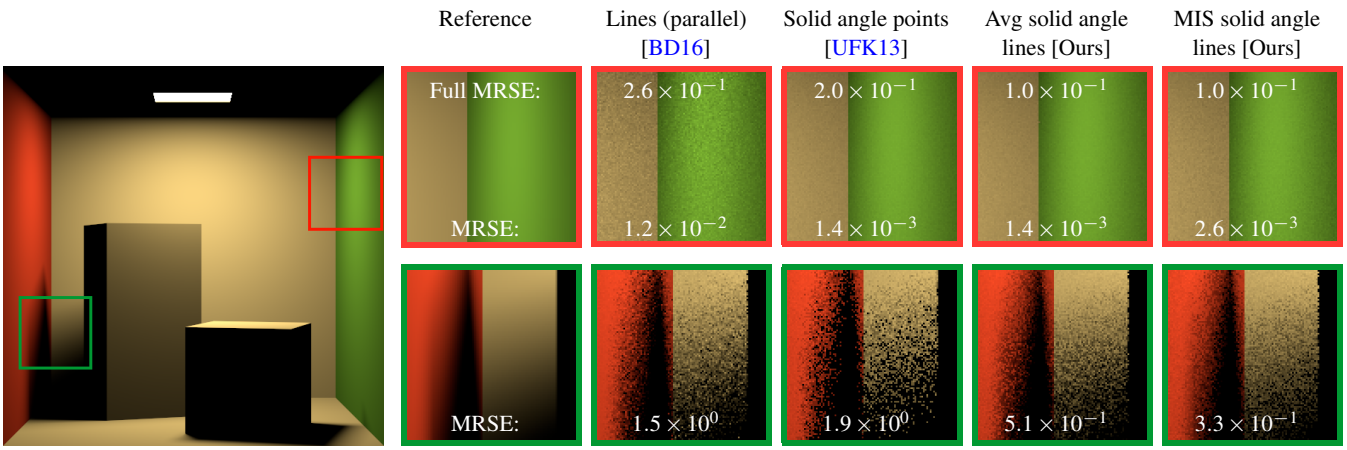

Figure 8: *MRSE of various sampling strategies in a Cornell box at equal time (two seconds). MRSE for the full image is given above the two rows and for each zoom-in beneath the corresponding row. Our average and MIS routines use two perpendicular line directions on the light source. The average routine (column 4) acts as a baseline for comparison against MIS, where line samples are simply averaged instead of weighted. We achieve the best MRSE in the shadowed region next to the red wall, while in an unshadowed region faster methods have an edge.*

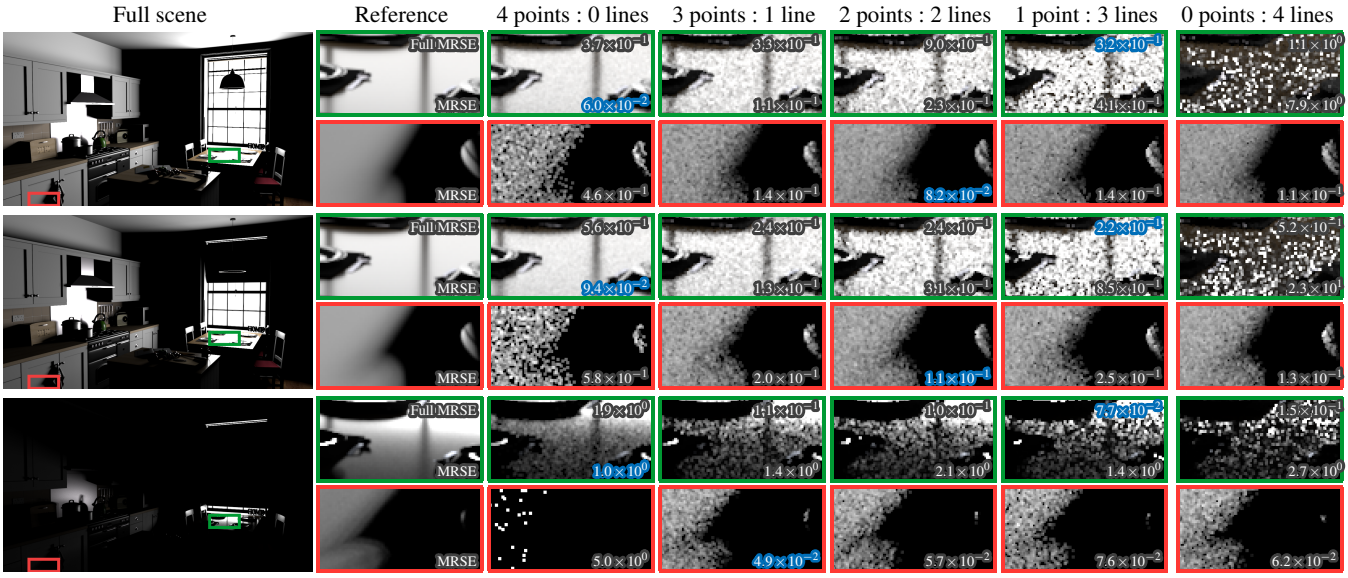

Figure 9: *At equal time (150s), we compare the effect of MISing between BSDF point samples and vertical solid-angle line samples using different ratios. At each surface intersection, we use four light samples divvied according to the ratio above each column. Best MRSE for each row is highlighted in blue. The BSDF strategy clearly dominates on the shiny table near the light source, but struggles in the penumbra region behind the towel. At equal time, MIS (specifically, the split 1 point : 3 lines) consistently achieves the best MRSE across the full image.*

6. Limitations and Future Work

Line sample-scene intersections. We share the same primary limitation as Billen and Dutré in that the line sample-scene intersections required to determine visible line regions dominate the overall runtime, especially since we must perform extra intersections in order to evaluate line sampling PDFs. We found that, on average, tracing a line sample took 1.2–55× as long as tracing a shadow ray for the scenes in this paper, depending on their complexity (see supplemental for a full comparison). We have shown that despite this limitation, our framework has both an improved convergence rate over points and better performance in shadowed regions of quick renders; however, for scenes with many occluders or little image space occupied by shadows, point strategies may be able to achieve a better result in equal time. While ray-scene intersections build on 40+ years of optimization research, we believe there is ample room for optimizing the line sample-scene intersection routine, further amplifying the benefits of line samples.

Solid-angle sampling. Due to the nature of Ureña et al.'s algorithm, our solid-angle line sampling is currently limited to sampling line directions aligned with edges of the light source. It should be possible to overcome this limitation by breaking the spherical rectangle into spherical triangles and using Arvo's [Arv95a] algorithm instead. It may also be possible to generalize our use of Ureña et al.'s approach to support arbitrarily oriented lines and polygonal emitters by line

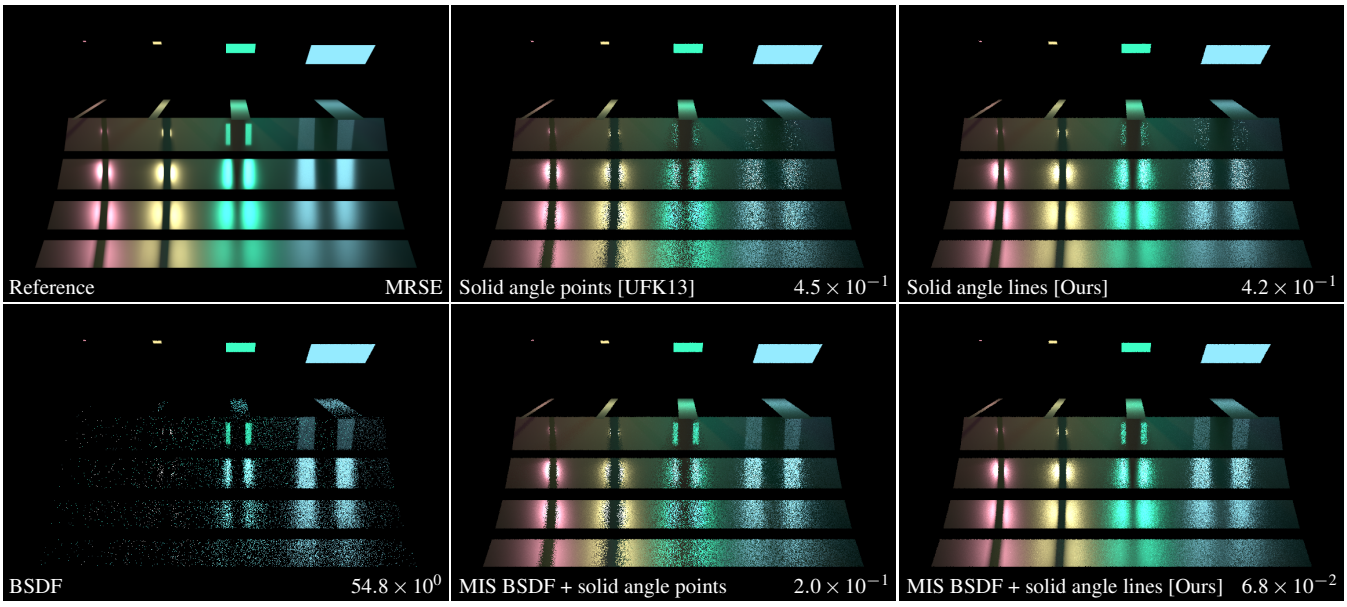

Figure 10: *We remodel Veach's classic MIS scene to include occluders. The plates use microfacet distributions which increase in roughness from top to bottom. As expected, BSDF sampling performs best in the upper right region of the outer product and solid-angle methods perform best in the lower left region. For our line strategies, we choose to use horizontal lines on the light source. At equal time (four seconds), our MIS between BSDF and lines has the lowest MRSE across the entire image. Our line methods are clearly most effective at reducing noise around shadow edges, while MIS between BSDF and solid-angle points reduces noise further than ours in some areas (see glossiest plate with largest light source) due to its speed.*

sampling the light's bounding quad and trimming the line samples to the extent of the actual emitter. Our surface-area line sampling handles any arbitrary line direction on the light source.

Smoothing MIS heuristic. Our smoothing MIS heuristic is effective at eliminating discontinuities in simple scenarios (see Fig. 4); however, the improvement in the occluded pixel's variance and convergence rate is more modest than the unoccluded pixel in Fig. 7. Designing a more robust smoothing MIS heuristic that works well in the general case would be an interesting direction for future work.

7. Conclusion

We developed a unified framework for combining point and line samples for efficient MC approximation of direct illumination. We demonstrated that the simple conceptual change of separating a 2D point PDF into marginal and conditional terms creates a general shared space for both point and line samples, and opens up new possibilities for combining the strengths of each. We extended an existing solid-angle point sampling scheme to lines, resulting in a new hybrid strategy that importance samples more terms in the direct lighting integral than either surface-area lines or solid-angle points. Lastly, we showed how to leverage the information found during a line sample-scene visibility intersection to smooth the integration domain of less-aware point strategies, resulting in an improved convergence rate for some scenarios.

8. Acknowledgements

We are grateful to the anonymous reviews for their suggestions on improving the paper. We also thank past and present members of the Dartmouth Visual Computing Lab. This work was partially supported by NSF grant IIS-181279.

References

[Arv95a] Arvo, J. R. "Stratified sampling of spherical triangles". *Annual Conference Series (Proceedings of SIGGRAPH)*. ACM Press, Aug. 1995, 437–438. ISBN: 978-0-89791-701-8. DOI: 10/bdrqbf 8.

[Arv95b] Arvo, J. R. "Analytic Methods for Simulated Light Transport". Ph.D. Thesis. New Haven, Connecticut: Yale University, Dec. 1995 2.

[Arv95c] Arvo, J. R. "Applications of irradiance tensors to the simulation of non-Lambertian phenomena". *Annual Conference Series (Proceedings of SIGGRAPH)*. ACM Press, Aug. 1995, 335–342. ISBN: 978-0-89791-701-8. DOI: 10/c2fss9 2.

[BD16] Billen, N. and Dutré, P. "Line sampling for direct illumination". *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 35.4 (July 1, 2016), 93–102. ISSN: 0167-7055. DOI: 10/f84z2h 1–4, 6–8.

[BGA12] Barringer, R., Gribel, C. J., and Akenine-Möller, T. "High-quality curve rendering using line sampled visibility". *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 31.6 (Nov. 2012), 162:1–162:10. ISSN: 0730-0301. DOI: 10/f25qxk 2.

[BJ17] Bitterli, B. and Jarosz, W. "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. ISSN: 0730-0301. DOI: 10/gfznbr 2.

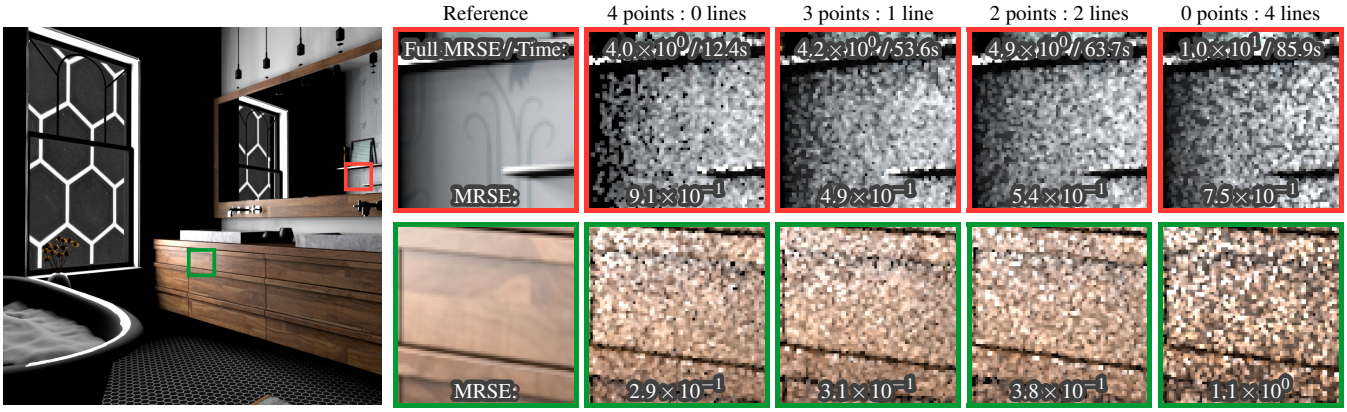

Figure 11: *We compare MISing between texture-based point samples and horizontal solid-angle line samples using a random sampler with four light samples per scene intersection per pixel. MRSE and render time for the full image is given at the top of the top row and for each highlighted region at the bottom of the corresponding row. We see that while textured-based points are clearly more effective in general, adding just one line sample quashes some variance on the left side of the wall (top row) while keeping the overall MRSE comparable to points alone.*

[BP93] Bao, H. and Peng, Q. "Shading models for linear and area light sources". *Computers & Graphics* 17.2 (1993), 137–145. ISSN: 0097-8493. DOI: 10/cf4f9p 2.

[BXH*18] Belcour, L., Xie, G., Hery, C., Meyer, M., Jarosz, W., and Nowrouzezahrai, D. "Integrating clipped spherical harmonics expansions". *ACM Transactions on Graphics* 37.2 (Mar. 2018), 19:1–19:12. ISSN: 0730-0301. DOI: 10/gd52pf 2.

[CA00] Chen, M. and Arvo, J. "A closed-form solution for the irradiance due to linearly-varying luminaires". *Rendering Techniques (Proceedings of the Eurographics Workshop on Rendering)*. Vienna: Springer-Verlag, June 2000, 137–148. ISBN: 978-3-7091-6303-0. DOI: 10/gfz9gv 2.

[CA01] Chen, M. and Arvo, J. R. "Simulating non-Lambertian phenomena involving linearly-varying luminaires". *Rendering Techniques (Proceedings of the Eurographics Workshop on Rendering)*. Springer-Verlag, June 2001, 25–38. ISBN: 978-3-7091-6242-2. DOI: 10/chb4qg 2.

[Coo86] Cook, R. L. "Stochastic sampling in computer graphics". *ACM Transactions on Graphics* 5.1 (Jan. 1986), 51–72. ISSN: 0730-0301. DOI: 10/cqwhcc 2.

[CPC84] Cook, R. L., Porter, T., and Carpenter, L. "Distributed ray tracing". *Computer Graphics (Proceedings of SIGGRAPH)* 18.3 (July 1, 1984), 137–145. ISSN: 0097-8930. DOI: 10/c9thc3 2.

[CSW94] Chiu, K., Shirley, P., and Wang, C. "Multi-jittered sampling". *Graphics Gems IV*. Ed. by Heckbert, P. S. San Diego, CA, USA: Academic Press, 1994, 370–374. ISBN: 0-12-336155-9 2, 6, 7.

[Dur11] Durand, F. *A Frequency Analysis of Monte-Carlo and Other Numerical Integration Schemes*. TR-2011-052. MIT CSAIL, Feb. 2011 2.

[GBA11] Gribel, C. J., Barringer, R., and Akenine-Möller, T. "High-quality spatio-temporal rendering using semi-analytical visibility". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 30.4 (July 1, 2011), 54:1–54:11. ISSN: 07300301. DOI: 10/fgq7b9 2.

[GDA10] Gribel, C. J., Doggett, M., and Akenine-Möller, T. "Analytical motion blur rasterization with compression". *Proceedings of High Performance Graphics*. Eurographics Association, 2010, 163–172. ISBN: 978-3-905674-26-2. DOI: 10/f2z5ds 2.

[GKDS12] Georgiev, I., Křivánek, J., Davidovič, T., and Slusallek, P. "Light transport simulation with vertex connection and merging". *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 31.6 (Nov. 2012), 192:1–192:10. ISSN: 0730-0301. DOI: 10/gbb6q7 2.

[GKH*13] Georgiev, I., Křivánek, J., Hachisuka, T., Nowrouzezahrai, D., and Jarosz, W. "Joint importance sampling of low-order volumetric scattering". *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 32.6 (Nov. 2013), 1–14. ISSN: 0730-0301. DOI: 10/gbd5qs 2.

[HCJ13] Habel, R., Christensen, P. H., and Jarosz, W. "Photon beam diffusion: a hybrid Monte Carlo method for subsurface scattering". *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 32.4 (June 2013), 27–37. ISSN: 1467-8659. DOI: 10/f445m4 2.

[Hei18] Heitz, E. *Generating Random Segments from Non-Uniform Distributions*. Research Report hal-01745725. Unity Technologies, Mar. 2018 2.

[HGJ*17] Hachisuka, T., Georgiev, I., Jarosz, W., Křivánek, J., and Nowrouzezahrai, D. "Extended path integral formulation for volumetric transport". *Proceedings of EGSR (Experimental Ideas & Implementations)*. Eurographics Association, June 2017. DOI: 10/gfznb3 2.

[HHM18] Heitz, E., Hill, S., and McGuire, M. "Combining analytic direct illumination and stochastic shadows". *Proceedings of the Symposium on Interactive 3D Graphics and Games* (Montreal, Quebec, Canada). New York, NY, USA: ACM Press, 2018, 2:1–2:11. ISBN: 978-1-4503-5705-0. DOI: 10/gfznb7 2.

[HPJ12] Hachisuka, T., Pantaleoni, J., and Jensen, H. W. "A path space extension for robust light transport simulation". *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 31.6 (Jan. 11, 2012), 191:1–191:10. ISSN: 0730-0301. DOI: 10/gbb6n3 2.

[ICG86] Immel, D. S., Cohen, M. F., and Greenberg, D. P. "A radiosity method for non-diffuse environments". *Computer Graphics (Proceedings of SIGGRAPH)* 20.4 (Aug. 1986), 133–142. ISSN: 0097-8930. DOI: 10/dmjm9t 2.

[Jen01] Jensen, H. W. *Realistic Image Synthesis Using Photon Mapping*. Natick, MA, USA: AK Peters, Ltd., 2001. ISBN: 1-56881-147-0 2.

[Jen96] Jensen, H. W. "Global illumination using photon maps". *Rendering Techniques (Proceedings of the Eurographics Workshop on Rendering)*. Vienna: Springer-Verlag, June 1996, 21–30. ISBN: 978-3-211-82883-0 978-3-7091-7484-5. DOI: 10/fzc6t9 2.

[JMM*14] Jarabo, A., Marco, J., Munoz, A., Buisan, R., Jarosz, W., and Gutierrez, D. "A framework for transient rendering". *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 33.6 (Nov. 2014), 177:1–177:10. ISSN: 15577368. DOI: 10/gfznb8 2.

[JNSJ11] Jarosz, W., Nowrouzezahrai, D., Sadeghi, I., and Jensen, H. W. "A comprehensive theory of volumetric radiance estimation using photon points and beams". *ACM Transactions on Graphics* 30.1 (Jan. 1, 2011), 5:1–5:19. ISSN: 0730-0301. DOI: 10/fcdh2f 2.

[JNT*11] Jarosz, W., Nowrouzezahrai, D., Thomas, R., Sloan, P.-P., and Zwicker, M. "Progressive photon beams". *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 30.6 (Dec. 2011), 181:1–181:12. ISSN: 07300301. DOI: 10/fn5xzj 2.

[JP00] Jones, T. R. and Perry, R. N. "Antialiasing with line samples". *Rendering Techniques (Proceedings of the Eurographics Workshop on Rendering)*. Springer-Verlag, 2000, 197–205. ISBN: 978-3-7091-6303-0. DOI: 10/gfznb9 2.

[JZJ08] Jarosz, W., Zwicker, M., and Jensen, H. W. "The beam radiance estimate for volumetric photon mapping". *Computer Graphics Forum (Proceedings of Eurographics)* 27.2 (Apr. 2008), 557–566. ISSN: 1467-8659. DOI: 10/bjsfsx 2.

[Kaj86] Kajiya, J. T. "The rendering equation". *Computer Graphics (Proceedings of SIGGRAPH)* 20.4 (Aug. 1986), 143–150. ISSN: 0097-8930. DOI: 10/cvf53j 2.

[Ken13] Kensler, A. *Correlated Multi-Jittered Sampling*. 13-01. Pixar Animation Studios, Mar. 2013 2.

[KGH*14] Křivánek, J., Georgiev, I., Hachisuka, T., Vévoda, P., Šik, M., Nowrouzezahrai, D., and Jarosz, W. "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. ISSN: 0730-0301. DOI: 10/f6cz72 2.

[LW93] Lafortune, E. P. and Willems, Y. D. "Bi-directional path tracing". *Proceedings of the International Conference on Computational Graphics and Visualization Techniques (Compugraphics)* (Alvor, Portugal). Vol. 93. Alvor, Portugal, Dec. 1993, 145–153 2.

[Max86] Max, N. L. "Atmospheric illumination and shadows". *Computer Graphics (Proceedings of SIGGRAPH)* 20.4 (Aug. 31, 1986), 117–124. ISSN: 0097-8930. DOI: 10/djr6wm 2.

[Max90] Max, N. L. "Antialiasing scan-line data". *IEEE Computer Graphics & Applications* 10.1 (Jan. 1990), 18–30. ISSN: 0272-1716. DOI: 10/dswxdv 2.

[MGJ*19] Marco, J., Guillén, I., Jarosz, W., Gutierrez, D., and Jarabo, A. "Progressive transient photon beams". *Computer Graphics Forum* 38.1 (Mar. 2019). ISSN: 1467-8659. DOI: 10/gfvr9w 2.

[MJGJ17] Marco, J., Jarosz, W., Gutierrez, D., and Jarabo, A. "Transient photon beams". *Congreso Espanol de Informatica Grafica*. Eurographics Association, June 2017. ISBN: 978-3-03868-046-8. DOI: 10/gfzncc 2.

[NBMJ14] Nowrouzezahrai, D., Baran, I., Mitchell, K., and Jarosz, W. "Visibility silhouettes for semi-analytic spherical integration". *Computer Graphics Forum* 33.1 (Feb. 2014), 105–117. ISSN: 1467-8659. DOI: 10/f5t6tf 2.

[NNDJ12a] Novák, J., Nowrouzezahrai, D., Dachsbacher, C., and Jarosz, W. "Progressive virtual beam lights". *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 31.4 (June 2012), 1407–1413. ISSN: 01677055. DOI: 10/gfzndw 2.

[NNDJ12b] Novák, J., Nowrouzezahrai, D., Dachsbacher, C., and Jarosz, W. "Virtual ray lights for rendering scenes with participating media". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 31.4 (July 2012), 60:1–60:11. ISSN: 0730-0301. DOI: 10/gbbwk2 2.

[NON85] Nishita, T., Okamura, I., and Nakamae, E. "Shading models for point and linear sources". *ACM Transactions on Graphics* 4.2 (Apr. 1985), 124–146. ISSN: 0730-0301. DOI: 10/cvrns9 2.

[Pho75] Phong, B. T. "Illumination for computer generated pictures". *Communications of the ACM* 18.6 (Jan. 6, 1975), 311–317. ISSN: 0001-0782. DOI: 10/bkfrm9 4.

[Pic92] Picott, K. P. "Extensions of the linear and area lighting models". *IEEE Computer Graphics & Applications* 12.2 (Mar. 1992), 31–38. DOI: 10/drm4j5 2.

[PJH16] Pharr, M., Jakob, W., and Humphreys, G. *Physically Based Rendering: From Theory to Implementation*. 3rd. Cambridge, MA: Morgan Kaufmann, 2016. ISBN: 978-0-12-800645-0 2, 6.

[PP09] Pegoraro, V. and Parker, S. G. "An analytical solution to single scattering in homogeneous participating media". *Computer Graphics Forum (Proceedings of Eurographics)* 28.2 (2009), 329–335. DOI: 10/c9zhxn 2.

[PSC*15] Pilleboue, A., Singh, G., Coeurjolly, D., Kazhdan, M., and Ostromoukhov, V. "Variance analysis for Monte Carlo integration". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 34.4 (July 2015), 124:1–124:14. ISSN: 0730-0301. DOI: 10/f7m28c 2, 5.

[SJ17] Singh, G. and Jarosz, W. "Convergence analysis for anisotropic Monte Carlo sampling spectra". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 36.4 (July 2017), 137:1–137:14. ISSN: 0730-0301. DOI: 10/gbxfhj 2.

[SK13] Subr, K. and Kautz, J. "Fourier analysis of stochastic sampling strategies for assessing bias and variance in integration". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 32.4 (July 2013), 128:1–128:12. ISSN: 0730-0301. DOI: 10/gbdg7c 2.

[SMJ17] Singh, G., Miller, B., and Jarosz, W. "Variance and convergence analysis of Monte Carlo line and segment sampling". *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 36.4 (July 2017), 79–89. ISSN: 01677055. DOI: 10/gfzncj 2, 4.

[SNJ*14] Subr, K., Nowrouzezahrai, D., Jarosz, W., Kautz, J., and Mitchell, K. "Error analysis of estimators that use combinations of stochastic sampling strategies for direct illumination". *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 33.4 (June 2014), 93–102. ISSN: 1467-8659. DOI: 10/f6fgw4 2.

[SRNN05] Sun, B., Ramamoorthi, R., Narasimhan, S. G., and Nayar, S. K. "A practical analytic single scattering model for real time rendering". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 24.3 (July 2005), 1040–1049. ISSN: 0730-0301. DOI: 10/fgnbqt 2.

[SSC*19] Singh, G., Subr, K., Coeurjolly, D., Ostromoukhov, V., and Jarosz, W. "Fourier analysis of correlated Monte Carlo importance sampling". *Computer Graphics Forum* 38.1 (Apr. 5, 2019). ISSN: 0167-7055, 1467-8659. DOI: 10/gfznck 2, 5.

[SSJ16] Subr, K., Singh, G., and Jarosz, W. "Fourier analysis of numerical integration in Monte Carlo rendering: theory and practice". *ACM SIGGRAPH Course Notes*. ACM Press, July 2016. ISBN: 978-1-4503-4289-6. DOI: 10/gfzncn 6.

[SZLG10] Sun, X., Zhou, K., Lin, S., and Guo, B. "Line space gathering for single scattering in large scenes". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 29.4 (July 26, 2010), 54:1–54:8. ISSN: 0730-0301. DOI: 10/dzxvvr 2.

[TPD*12] Tzeng, S., Patney, A., Davidson, A., Ebeida, M. S., Mitchell, S. A., and Owens, J. D. "High-quality parallel depth-of-field using line samples". *Proceedings of High Performance Graphics*. Eurographics Association, 2012, 23–31. ISBN: 978-3-905674-41-5. DOI: 10/gfzncq 2.

[UFK13] Ureña, C., Fajardo, M., and King, A. "An area-preserving parametrization for spherical rectangles". *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 32.4 (2013), 59–66. ISSN: 1467-8659. DOI: 10/gfzncr 1, 2, 4–6, 8.

[VG95a] Veach, E. and Guibas, L. J. "Bidirectional estimators for light transport". *Photorealistic Rendering Techniques (Proceedings of the Eurographics Workshop on Rendering)*. Springer-Verlag, 1995, 145–167. ISBN: 978-3-642-87825-1. DOI: 10/gfznbh 2.

[VG95b] Veach, E. and Guibas, L. J. "Optimally combining sampling techniques for Monte Carlo rendering". *Annual Conference Series (Proceedings of SIGGRAPH)*. Vol. 29. ACM Press, Aug. 1995, 419–428. ISBN: 978-0-89791-701-8. DOI: 10/d7b6n4 2, 3.

[WR18] Wang, J. and Ramamoorthi, R. "Analytic spherical harmonic coefficients for polygonal area lights". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 37.4 (July 30, 2018), 1–11. ISSN: 07300301. DOI: 10/gd52s2 2.