DOI: 10.1111/cgf.70164
Eurographics Symposium on Rendering 2025
B. Wang and A. Wilkie
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 44 (2025), Number 4

A Data-Driven Approach to Analytical Dwivedi Guiding

Darryl Gouder[1,2], Jiří Vorba[2], Marc Droske[2], Alexander Wilkie[1]

[1]Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic
[2]Wētā FX, New Zealand

Figure 1: *We improve analytical Dwivedi sampling by informing the slab normal selection with the boundary incident illumination information. The left side of the image uses the canonical Dwivedi Point of Entry Slab Normal, where paths are guided to exit the volume from the side of the point of entry. The right shows our data-driven approach, which adapts the guiding direction depending on the energy transport that arrives at the boundary. In the denser parts of the medium on the face, the paths are mostly guided towards the point of entry, whilst on the ears, which are backlit, the paths are guided to the back.*

Abstract

Path tracing remains the gold standard for high-fidelity subsurface scattering despite requiring numerous paths for noise-free estimates. We introduce a novel variance-reduction method based on two complementary zero-variance-theory-based approaches. The first one, analytical Dwivedi sampling, is lightweight but struggles with complex lighting. The second one, surface path guiding, learns incident illumination at boundaries to guide sampled paths, but it does not reduce variance from subsurface scattering. In our novel method, we enhance Dwivedi sampling by incorporating the radiance field learned only at the volume boundary. We use the average normal of points on an illuminated boundary region or directions sampled from distributions of incident light at the boundary as our analytical Dwivedi slab normals. Unlike previous methods based on Dwivedi sampling, our method is efficient even in scenes with complex light rigs typical for movie production and under indirect illumination. We achieve comparable noise reduction and even slightly improved estimates in some scenes compared to volume path guiding, and our method can be easily added on top of any existing surface path guiding system. Our method is particularly effective for homogeneous, isotropic media, bypassing the extensive training and caching inside the 3D volume that volume path guiding requires.

CCS Concepts

• ***Computing methodologies*** → ***Ray tracing;***

1. Introduction

Realistic light transport simulation is now prevalent in visual effects, architectural visualization, and computer games. The simulation involves sampling random light paths in the virtual scene while integrating the transport along them using Monte Carlo (MC) integration. For the realistic appearance of materials like milk, marble, and human skin, accurate subsurface scattering simulation is needed when light enters the volume through its boundary and scatters. If not absorbed, light exits at a different boundary point.

Precise Monte Carlo simulation of subsurface scattering is expensive. The simulation requires tracing many long paths through the volume to average out the MC estimator noise. Our primary goal is to reduce subsurface scattering simulation noise via importance sampling. To that end, we follow the works based on zero-variance theory [Hoo08], which can be classified into two lines of work. One line of work on so-called *Dwivedi sampling* [Dwi82b; Dwi82a; Kd14; MHD16] employs fully analytical sampling based on a set of simplifying assumptions about the integrand. The assumptions —a half-infinite homogeneous volume slab with isotropic scattering and the equal importance of each point at the boundary —are often violated in practice, which leads to poor sample contributions.

Works in the other line, known as *path guiding*, are data-driven and follow the rules of zero-variance walks while using learned approximations of the light field in the scene. Surface path guiding usually involves learning the light field distribution on surfaces in the scene and has become mainstream in production for surface light transport [VHH*19]. However, to our knowledge, surface path guiding caches have not been used to guide volumetric paths. Herholz et al. propose an adjoint-driven pre-process to learn the volumetric radiance field represented by spatially cached 2D directional distributions [HZE*19]. This involves building an internal spatial structure and fitting 2D directional distributions within volumes. We argue that implementing volumetric path guiding is challenging and requires guiding caches inside volumes. We have found that volumetric guiding caches can become dense when approximating optically thick media, even if the media is homogeneous.

Our method uses analytical Dwivedi kernels for path sampling and surface path guiding caches for Dwivedi slab normal sampling. The goal is to retain Dwivedi sampling's simplicity while increasing robustness under complex lighting typical in movie production and other applications. It can be implemented on top of regular surface path guiding, available in many production renderers [VHH*19]. If surface path guiding is the baseline, our approach has negligible memory overhead and may require no additional training. In theory, based on the application and specifics of the rendering system, the incident illumination could be represented differently (e.g., neural networks, light probes, or textures), but this is left for future work.

Our contributions can be summarized as follows:

- We propose exploiting the learned light field at the surface boundary to sample the slab normal for Dwivedi kernels.
- We use the Burley diffusion profile to identify regions on the boundary that contribute energy to some reference point on the boundary.

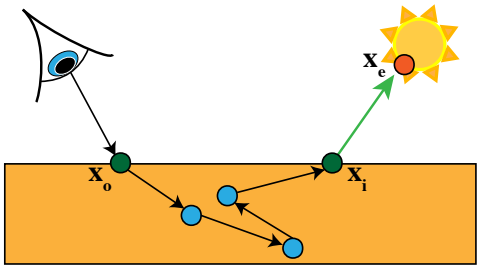

Figure 2: *The random walk is constructed via volume distance and directional sampling. We continue path sampling beyond the exit point* $\mathbf{x_i}$ *to find emitter source vertices* $\mathbf{x_e}$*. Our importance sampling guides the subsurface path towards exit points transmitting high energy into the volume.*

- The *Boundary Normal (BONO)* technique, which uses the average normal of some region on the boundary, that contributes energy to the volume entry point.
- The *Boundary Incident Illumination (BILL)* technique, which samples surface path guiding directional distributions for a direction and applies it as a slab normal for analytical Dwivedi sampling.

Our methods stabilize and improve Dwivedi in complex lighting scenarios whilst outperforming VPG in some situations.

2. Background and related work

Computing outgoing radiance due to subsurface scattering

$$L_o^{\mathrm{sss}}(\mathbf{x_o}, \omega_\mathbf{o}) = \int_{\Omega^-} \rho^{\perp}(\mathbf{x_o}, \omega_\mathbf{o}, \omega) L_i(\mathbf{x_o}, \omega) \mathrm{d}\omega \tag{1}$$

from the *point of entry* $\mathbf{x_o}$ at the volume boundary towards the direction $\omega_\mathbf{o}$, corresponds to integrating incoming radiance L_i from the volume multiplied by *bidirectional scattering distribution function* (BSDF) $\rho^{\perp}$ describing material properties and cosine foreshortening. The integration domain is the hemisphere of directions Ω^- at the boundary side oriented towards the volume.

The incoming radiance from the volume at the point $\mathbf{x}$ from the direction ω (pointing away from $\mathbf{x}$) is described by *volume rendering equation*

$$L_i(\mathbf{x}, \omega) = \int_t T(\mathbf{x}, \mathbf{x}) \sigma_s(\mathbf{x}) L_s(\mathbf{x}, -\omega) \, \mathrm{d}t + T(\mathbf{x}, \mathbf{x_i}) L_o(\mathbf{x_i}, -\omega), \tag{2}$$

which is derived from the *radiative transfer equation* [Cha60]. Here, the in-scattered radiance L_s, multiplied by the scattering coefficient σ_s and the transmittance

$$T(\mathbf{x}, \mathbf{y}) = e^{-\tau(|\mathbf{x}-\mathbf{y}|)}, \tag{3}$$

is integrated over the points $\mathbf{x} = \mathbf{x} + \omega t$. The optical thickness $\tau(l) = \int_{s=0}^{l} \sigma_t \mathrm{d}s$ between the two points in a homogeneous volume depends on the extinction coefficient $\sigma_t = \sigma_a + \sigma_s$. The inverse of the extinction coefficient is the *mean free path* (mfp). The in-scattered radiance $L_s = \int_{\mathcal{S}} \rho_{ph}(\mathbf{x}, -\omega, \omega_\mathbf{i}) L_i(\mathbf{x}, \omega_\mathbf{i}) \, \mathrm{d}\omega_\mathbf{i}$ integrates the product of the phase function ρ_{ph} and the incoming radiance

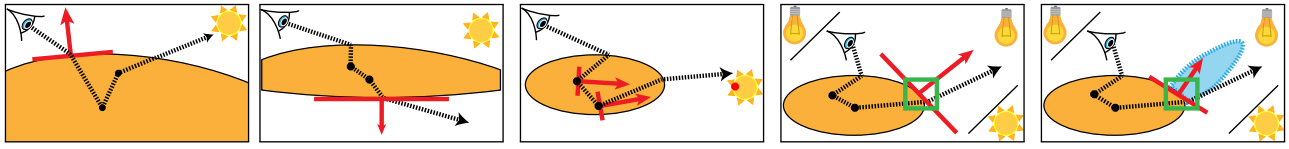

(a) *PoE* **(b)** *CP* **(c)** *II* **(d)** *BONO* **(e)** *BILL*

Figure 3: *The slab is oriented based on the selected normal, shown in red. Previous methods have used:* ***(a)*** *Point of Entry PoE normal,* ***(b)*** *Closest Point CP normal, or* ***(c)*** *a direction towards the emitter, called Incident Illumination II. Our methods first require a training phase to spatially divide the boundary and learn average normals and incident illumination directional distributions. Post-training, when a path enters the volume, an illuminated region (green box) is sampled and a slab normal from this region is selected. In* ***(d)*** *our first method, Boundary Normal BONO, the average of the training data point normals is used as the slab normal. Our second method,* ***(e)*** *BILL, uses a direction from the surface path guiding incident illumination directional distribution as the slab normal.*

L_i over the sphere of directions at each point $\mathbf{x}$. We also add outgoing radiance L_o from the *point of exit* $\mathbf{x_i}$, which describes both the radiance transmitted into the volume and the inner reflections. We name points of entry and exit with respect to the direction of paths traced from the camera when we estimate Eq. 1. Also note we assume all light sources are outside of the volume.

Estimating subsurface scattering. To estimate L_o^{sss}, we incrementally sample volume random walks starting from point of entry $\mathbf{x_0}$ towards the point of exit $\mathbf{x_i}$ (see Figure 2). We sample from the camera direction, thus, we continue sampling towards the lights behind $\mathbf{x_i}$. The *Classical* approach [PJH16] repeatedly samples first the path segment distance proportional to transmittance T, and then a new ray direction proportionally to the phase function ρ_{ph}. Optionally, Russian roulette can be played with survival probability equal to single scattering albedo or a user-chosen value for an unbiased estimator. The Classical approach's disadvantage is the local-term importance sampling without respecting the incoming radiance. This can result in high-variance estimates.

Dwivedi sampling. Dwivedi et al. proposed distance and directional importance sampling using analytical functions (kernels) derived under simplifying assumptions [Dwi82a]. The volume is considered a uniformly-lit, homogeneous, isotropic, half-infinite slab. Heterogeneous volumes composed of finite homogeneous slabs were also discussed. Křivánek et al. applied these variance-reducing kernels to subsurface scattering [Kd14] in computer graphics. They parameterize the kernels by the slab normal $\vec{n}_D$ at the *Point of Entry (PoE)* $\mathbf{x_0}$. This strategy guides subsurface chains back towards the boundary, which is considered a uniform source of importance. However, such sampling can be highly sub-optimal when the boundary is not illuminated uniformly.

Sampling a direction corresponds to sampling a cosine of the angle θ between the slab normal $\vec{n}_D$ and the sampled direction ω in the shading space, where the z-axis is the up vector, and therefore $\cos\theta$ corresponds to the z component of the sampled direction. Given an i.i.d. number ξ_0, and the diffusion length ν_0, we compute

$$\cos\theta = \nu_0 - (\nu_0 + 1) \cdot \left(\frac{\nu_0 - 1}{\nu_0 + 1} \right)^{\xi_0}. \tag{4}$$

We sample ω_z with the density

$$p_D(\omega_z | \vec{n}_D) = \frac{1}{\log \frac{\nu_0 + 1}{\nu_0 - 1}} \cdot \frac{1}{\nu_0 - \cos(\theta)}. \tag{5}$$

To find x and y components of the sampled direction, we sample the azimuth uniformly. Thus, the total probability of sampling the direction ω in this case is $p(\omega) = p_D(\omega_z | \vec{n}_D) \cdot \frac{1}{2\pi}$. Note, that this kernel samples the directions around $\vec{n}_D$ with high density.

The diffusion length ν_0 can be computed analytically, but it is a non-trivial task. Nevertheless, D'Eon provides the most accurate approximation to our knowledge [KGV*20]

$$\nu_0 \approx \ell \frac{1}{\sqrt{1 - \alpha^{2.44294 - 0.0215813\alpha + \frac{0.578637}{\alpha}}}}, \tag{6}$$

which is computed from the single scattering albedo $\alpha = \sigma_s / \sigma_t$ and mean free path $\ell = 1/\sigma_t$. Obviously, the Dwivedi sampling kernels are dependent on the volume properties. If the medium has high absorption and thus low albedo, the sampling is biased more strongly towards the slab normal $\vec{n}_D$. On the other hand, no absorption results in the Dwivedi kernel taking the shape of isotropic scattering.

The distance sampling kernel stretches distances along the rays sampled in direction ω which form small angle θ between the slab normal $\vec{n}_D$ and ω. On the other hand, it shortens the sampled distance when ω points away from $\vec{n}_D$. Distance t is sampled proportionally to transmittance T within the Classical approach. However, the distance sampling is parameterized by the modified extinction coefficient $\sigma_t' = \sigma_t \cdot (1 - \cos(\theta)/\nu_0)$. The actual transmittance is still calculated with the original σ_t.

Dwivedi sampling with incident illumination. Meng et al. propose alternative options for selecting the slab normal $\vec{n}_D$ [MHD16]. They choose a point $\mathbf{y}$ either by finding the closest point on the boundary with respect to the first sampled interaction $\mathbf{x_1}$ within the volume or they sample a point on an emitter. They define the slab normal $\vec{n}_D$ for the Dwivedi kernels as a direction from $\mathbf{x_1}$ to $\mathbf{y}$. These methods are called *Closest Point (CP)* and *Incident Illumination (II)*, respectively. This choice works well in scenes with optically thin volumes that are back lit with one light source which is a difficult scenario for the Point of Entry sampling. Unfortunately, the method can struggle in scenes with complex lighting, multiple emitters, and occlusion because it does not recognize the actual illumination at the volume boundary.

Volume path guiding. This problem is addressed in a fundamentally different, fully data-driven approach. Herholz et al. learn an estimate of the volumetric radiance field in a pre-

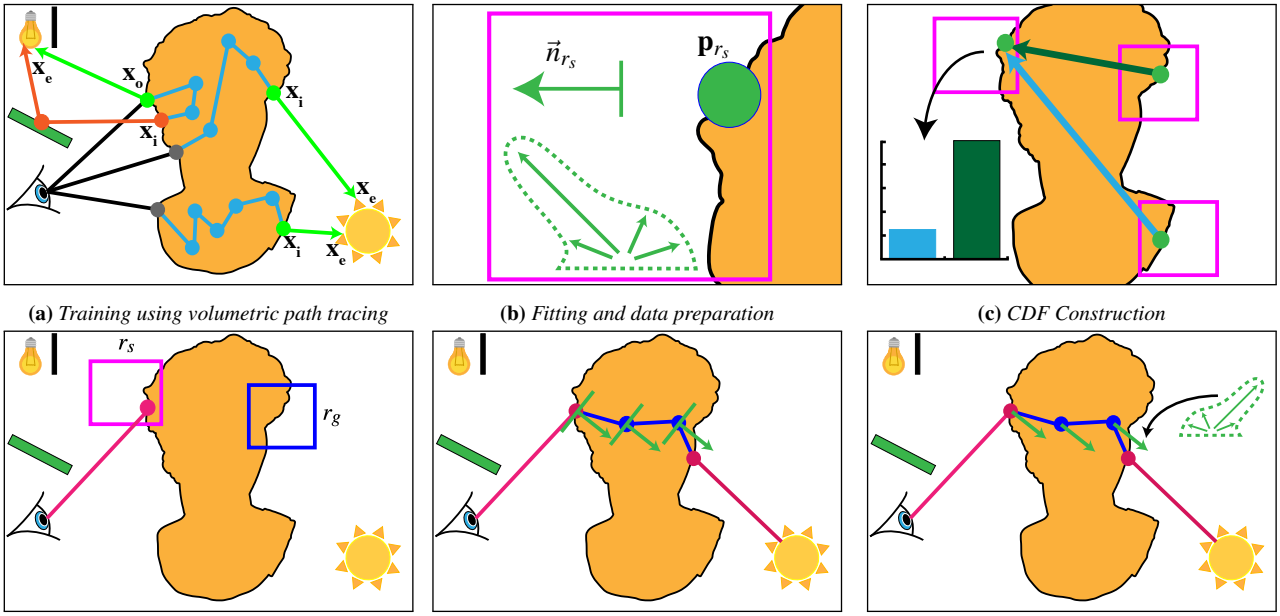

(a) *Training using volumetric path tracing* **(b)** *Fitting and data preparation* **(c)** *CDF Construction*

(d) *Sampling an illuminated region* r_g **(e)** *BONO subsurface path guiding* **(f)** *BILL subsurface path guiding*

Figure 4: ***(a)*** *Random walks using Classical sampling provide training data for the caches. NEE at* $\mathbf{x_o}$ *and* $\mathbf{x_i}$ *(green paths) and indirect illumination at* $\mathbf{x_i}$ *(orange paths) populate the cache.* ***(b)*** *In each region, we use the training data to fit vMF proportionally to incident illumination, compute the averaged normal* $\vec{n}_{r_s}$*, and the averaged point* $\mathbf{p}_{r_s}$*.* ***(c)*** *The CDFs are constructed to allow importance sampling of illuminated regions.* ***(d)*** *When a path hits the region* r_s *at* $\mathbf{x_o}$*, we sample an illuminated region* r_g *using the CDF stored in* r_s*.* ***(e)*** *To guide subsurface sampling of the path, BONO uses the average normal* $\vec{n}_{r_g}$ *from the selected region* r_g *(green arrows) as the Dwivedi slab normal.* ***(f)*** *Alternatively, in our BILL, we sample the slab normal (green arrows) from the vMF obtained from* r_g*.*

training pass[HZE*19]. The direction and distance sampling decisions are guided at every vertex by exploiting the learned radiance field. Furthermore, low-contributing paths are terminated using adjoint-driven Russian Roulette, while high-contributing paths split [AK90; VK16; RGH*22]. Herholz's proposal targets more general cases than our approach including unbounded volumes, anisotropic phase function. It is also robust under various lighting conditions. However, implementing it is non-trivial and requires training and caching directional distributions inside the whole volume. Deng et al. propose a similar approach focused only on guiding the scattering direction sampling [DWWH20]. Wu et al. propose using differentiable regularization to reduce variance, however, this introduces bias to the estimator [WWH*24]. Leonard et al. proposed a data-driven approach using Conditional Variational Auto-encoders (CVAEs) to model photon path distributions within spherical regions, approximating multiple scattering events [LHW21]. Their sphere-tracing algorithm efficiently condenses long scattering chains into single representative steps, but introduces statistical bias that deviates from converged solutions in thin, translucent media and requires substantial memory. In contrast, we propose a method that requires caches only at the volume boundary and can work even without learning directional distributions. On the other hand, our method is limited to bounded homogeneous volumes with (nearly) isotropic scattering.

Surface path guiding. Similar to volume path guiding, paths can also be guided outside volumes when we learn the approximation of incident illumination at scene surfaces. This approach is widespread in industry [VHH*19], with many methods using different radiance field representations. Possible choices include photon maps [Jen95], cosine lobes [BDC12], Gaussian Mixture Models (GMMs) [VKŠ*14; HEV*16],vMF mixtures [RHL20] or quadtrees [MGN17]. Simon et al. [SHJD18] represent only the noisiest light field part as full light paths, but this approach is inefficient for subsurface scattering with varying path directions and distances. Other methods are based on neural networks [MMR*19; HIT*24; DWL23] or formulated as reinforcement learning [DK17]. Our method learns an approximation of the boundary incident illumination to sample the Dwivedi guiding slab normal $\vec{n}_D$, which guides subsurface paths towards illuminated boundary regions. We also show that it is possible to avoid learning directional distributions, which might be useful for some applications.

BSSRDFs. Subsurface scattering can also be described by the BSSRDF, modelling complete subsurface light transport integration between two points at the surface boundary. Assuming a homogeneous, half-infinite slab with isotropic scattering, the BSSRDF can be precisely represented by dipole diffusion [JMLH01; FHK15] or the quantized-diffusion model [DI11]. Using surface boundary irradiance estimates to accelerate BSSRDF evaluation has also been explored [JB02]. Alternative BSSRDF approaches using neural networks have also been explored [VKJ19; TTJ*24]. We use the Burley diffusion profile [HMB*15; Chr15] to approximate the distribution of energy transported between two volume

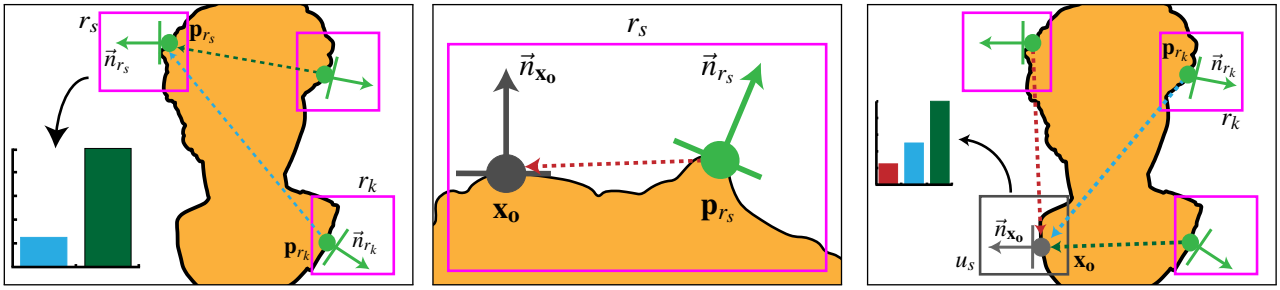

(a) *Illuminated region CDF construction* **(b)** *Same illuminated region weight calculation* **(c)** *Unilluminated region CDF construction*

Figure 5: ***(a)*** *The CDF weight of r_s with respect to r_k is parameterized by the length of the segment between the mean boundary point $\mathbf{p}_{r_s}$ and $\mathbf{p}_{r_k}$. These CDF weights are cached and reused for all paths at the region sampling the CDF.* ***(b)****The CDF weight of some r_s with respect to r_s is parametrized by the distance between the first path vertex $\mathbf{x_o}$ to encounter the region after training and the mean boundary point $\mathbf{p}_{r_s}$. This weight is computed at each path entering the medium at some x_o at some region r_s.* ***(c)*** *A region without the average statistics is considered to be an unilluminated region, denoted by u_s. We compute its CDF weights by considering the first $\mathbf{x_o}$ at u_s and the mean boundary points of the other illuminated regions, $\mathbf{p}_{r_k}$. These CDF weights are cached for all subsequent paths at u_s.*

boundary points. Since we do not directly use the Burley diffusion profile to evaluate the integrand, our method remains unbiased.

3. A Data-Driven Approach to Dwivedi Guiding

When a path enters the medium at $\mathbf{x_o}$, as shown in Figure 2, we need to sample a subsurface path that transports a significant amount of energy between an exit $\mathbf{x_i}$ point and $\mathbf{x_o}$. To that end, we importance sample an illuminated region to select a slab normal for Dwivedi sampling, thereby guiding the volume path sampling towards this region. We propose two alternative techniques for slab normal sampling based on the selected region:

1. The boundary average normal of the exit region, which we call Boundary Normal (BONO). This method samples a region on the boundary and uses the mean of all the normals associated with boundary points at that region that receive radiance.
2. The direction from some boundary region's incident illumination directional distribution. We call this method Boundary Incident Illumination (BILL) and can leverage surface path guiding directional distributions.

In Section 3.1, we describe learning and caching of the incident illumination at the boundary. In Section 3.2, we use this information to construct cumulative distribution functions (CDFs) for sampling boundary regions proportionally to their energy contributions to a given entry point. In Section 3.3, we describe the situation after sampling of the boundary region and how we use it to guide the scattering and free-flight sampling roughly towards the selected region. For clarity of exposition, we will describe our training and sampling in the context of BONO, followed by the necessary tweaks required for BILL.

3.1. Training the Boundary Volume Guiding Cache

We spatially discretise the boundary using a sparse grid structure that partitions the boundary into non-overlapping regions where we store information about the incident illumination at the boundary. Thus, our method is independent of the surface representation and does not require geometric storage.

Figure 4 illustrates the training process and the grid structure. We adopt a forward-training approach [MGN17], that is, we trace paths from the camera as in regular path tracing, and we evaluate subpaths to learn illumination at specific path vertices. Specifically, at each $\mathbf{x_o}$ and $\mathbf{x_i}$ on the volume boundary, we learn direct illumination by using next-event estimation to connect them to some point $\mathbf{x_e}$ on an emitter. When a path exits the medium at $\mathbf{x_i}$, we continue sampling the path to estimate the indirect illumination at $\mathbf{x_i}$. In our implementation, if a region receives at least 32 samples, then we consider it to be an *illuminated region* and is denoted as r_s, where s is the index of the region. Conversely, we refer to regions u_s that do not meet this threshold as *unilluminated* regions. We use the training data to construct an approximation of the radiance field. Specifically, we estimate the average irradiance $\tilde{E}_{r_s}$ for each illuminated region r_s from all direct and indirect illumination samples collected within r_s. Furthermore, we compute the average surface normal $\vec{n}_{r_s}$ and average point position $\mathbf{p}_{r_s}$ from the training boundary points (the $\mathbf{x_o}$ and $\mathbf{x_i}$ locations) in that region.

3.2. CDF Construction

Given a point $\mathbf{x_o}$, we construct a discrete cumulative distribution function (CDF) over the estimated contribution of each illuminated region to $\mathbf{x_o}$. However, building a CDF at each $\mathbf{x_o}$ is expensive, particularly when the number of illuminated regions is large. To mitigate this cost, we cache a common CDF for the whole region containing $\mathbf{x_o}$. Therefore, the next sample that intersects the region at $\mathbf{x'_o}$, we reuse the previously constructed CDF. We observed that caching did not increase variance in equal-sample renders.

To compute the CDF weights, we approximate the volume transport between two regions using a diffusion profile, denoted by $R_d(t)$ and transmittance term, $Tr(t)$, where t is the distance between the two regions. The distance between two illuminated regions, r_s and r_k is approximated by computing the distance between their respective averaged points $\mathbf{p}_{r_s}$ and $\mathbf{p}_{r_k}$.

The diffusion profile captures the multiple and single scattering terms, whilst the transmittance captures the direct transmission between the two points. The weighting function requires the reference point $\mathbf{y_0}$ that receives energy and $\mathbf{y_1}$ that is considered the source. The weighting function is defined as:

$$w(\mathbf{y}_0,\vec{n}_0,\mathbf{y}_1,\vec{n}_1,\tilde{E}_r) := \Big(R_d(||\mathbf{y}_0-\mathbf{y}_1||)+ Tr(\|\mathbf{y}_0-\mathbf{y}_1\|)\cdot f(\vec{n}_0,\vec{n}_1)\Big)\cdot\tilde{E}_r. \quad (7)$$

Here, $\tilde{E}_r$ denotes the irradiance of the illuminated region that is considered the illuminant for $\mathbf{y}_0$. We apply $f(\vec{n}_0,\vec{n}_1)$ to control the transmittance between the 2 points by cutting it off depending on the orientation of the normals, $\vec{n}_{\mathbf{y_0}}$ and $\vec{n}_{\mathbf{y_1}}$. The cutoff function is:

$$f(\vec{n}_0,\vec{n}_1) = \begin{cases} 0 & \text{if } \vec{n}_0\cdot\vec{n}_1 \geq \cos(\frac{\pi}{4}) \\ 1, & \text{otherwise} \end{cases} \quad (8)$$

As mentioned, in Equation 7 $\mathbf{y}_0$ is the reference point that receives energy, and $\mathbf{y}_1$ is considered the source. More specifically for its application in this work:

- $\mathbf{y}_0,\vec{n}_0$ represent the position and normal characterizing the receiving point. This can be a specific point of entry $\mathbf{x_o}$ with its normal $\vec{n}_{\mathbf{x_o}}$ (as used when $\mathbf{x_o}$ is at an illuminated region r_s or an unilluminated region u_s), or the average point $\mathbf{p}_{r_s}$ of a receiving region r_s with its average normal $\vec{n}_{r_s}$ (used for building the CDF for r_s based on other regions r_k).
- $\mathbf{y}_1,\vec{n}_1$ characterize the source of illumination. These are typically the average point $\mathbf{p}_{r_k}$ (or $\mathbf{p}_{r_s}$ when $\mathbf{x_o}$ is in r_s and r_s itself is considered the source) and average normal $\vec{n}_{r_k}$ (or $\vec{n}_{r_s}$) of an illuminated region that acts as the illuminant.
- $\tilde{E}_r$ is the average irradiance of this illuminated region that acts as the source for $\mathbf{y}_0$.

This interpretation allows Equation 7 to be applied consistently in the following scenarios for CDF construction.

The CDF weight for the energy arriving at r_s from r_k is computed using $w(\mathbf{p}_{r_s},\vec{n}_{r_s},\mathbf{p}_{r_k},\vec{n}_{r_k},\tilde{E}_{r_k})$. Computing this distribution for all illuminated regions r_k (where k is not the index for r_s) provides the mechanism to stochastically select another region r_k that most likely contributes energy to r_s. However, when $\mathbf{x_o}$ is at an illuminated region r_s, we must decide whether to seed the Dwivedi sampling by the slab normal from r_s or whether we acquire the slab normal from one of the other illuminated regions r_k. To make this decision, we use the sum of the unnormalized CDF weights from other regions, $t_{r_s} = \sum_{k\neq s} w(\mathbf{p}_{r_s},\vec{n}_{r_s},\mathbf{p}_{r_k},\vec{n}_{r_k},\tilde{E}_{r_k})$, and a weight $b = w(\mathbf{x_o},\vec{n}_{\mathbf{x_o}},\mathbf{p}_{r_s},\vec{n}_{r_s},\tilde{E}_{r_s})$. The probability of acquiring the slab normal from r_s is $\frac{b}{b+t_{r_s}}$. Otherwise, a region r_k is sampled from the CDF constructed for r_s, based on weights $w(\mathbf{p}_{r_s},\vec{n}_{r_s},\mathbf{p}_{r_k},\vec{n}_{r_k},\tilde{E}_{r_k})$.

The first time $\mathbf{x_o}$ lands on an unilluminated region, denoted by u_s, the CDF is constructed. The weight of the energy distribution for $\mathbf{x_o}$ receiving from each illuminated region r_k is $w(\mathbf{x_o},\vec{n}_{\mathbf{x_o}},\mathbf{p}_{r_k},\vec{n}_{r_k},\tilde{E}_{r_k})$. The CDF is cached and used for all subsequent $\mathbf{x'_o}$ that land in the region u_s. Unlike illuminated regions where there is a decision step involving b and t_{r_s}, when $\mathbf{x_o}$ is at an unilluminated region, an illuminated region is sampled using only this CDF.

Figure 6: *Yellow regions indicate where the point of entry slab normal matches the slab normal in use. The Point Of Entry (PoE) image is mostly yellow because the normals are equal. In Incident Illumination (II), the distribution of slab normals is more uniform, as this technique simply selects one of the illuminants as the slab normal. In contrast, our method shows that paths interacting with the ear are guided toward the backlight, while those in the front-lit region align more closely with the point of entry normal.*

We chose the analytical Burley diffusion profile to approximate volume transport, for its simplicity and efficiency. Although it does not accurately approximate the volume transport between two non-coplanar points, we show in Section 4.2 that using a more accurate diffusion profile in such scenarios did not improve the CDF quality.

3.3. Our Dwivedi Guiding

We now describe how to use the structure to guide our volumetric paths. When a path hits the medium at $\mathbf{x_o}$, we sample a region r_g. We then use r_g to acquire the slab normal we will use in the analytical Dwivedi sampling kernels to guide the path. The slab normal used depends on the guiding technique used. We discuss BONO, guiding using the mean of the boundary normals.

BONO (Boundary Normal) Given a sampled region, r_g, the slab normal $\vec{n}_D$ is set to the region's average boundary normal n_{r_g}. The slab normal is fixed for all scattering and free-flight sampling decisions until the path exits the volume.

BILL (Boundary Incident Illumination) Inspired by the work of Meng et al., we propose learning the directional distribution of the incident illumination at the boundary and use that distribution as the sample space for the slab normal. Unlike II, our directional distribution captures both the direct and indirect illumination incident on the boundary. The training for this method requires incoming directions of incident illumination that can be obtained by the training method mentioned in Sec. 3.1. When we sample the region, r_g, we subsequently sample the slab normal $\vec{n}_D$ from the incident illumination directional distribution stored in r_g. Once $\vec{n}_D$ has been selected, the scattering directions and free-flight distance are sampled using the Dwivedi kernels. $\vec{n}_D$ is fixed for all the segments of the random walk inside the medium until the path exits the volume. In our implementation, we represented the incident illumination directional distributions using vMF mixtures.

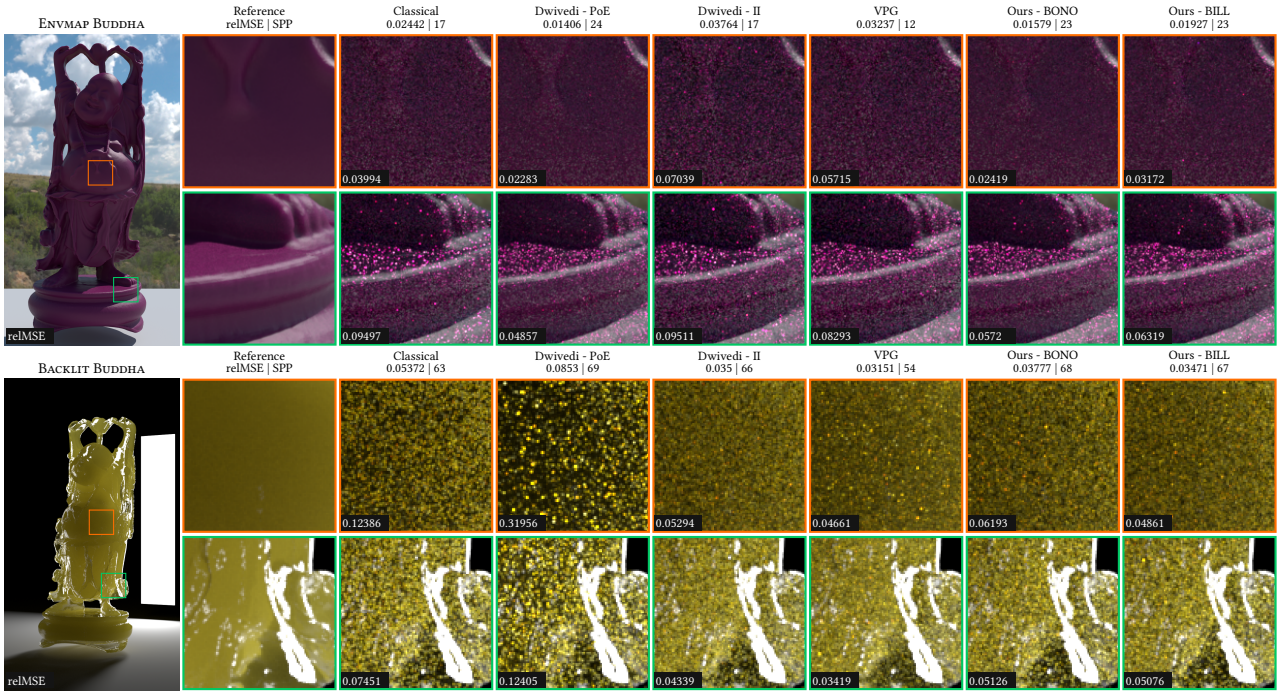

Figure 7: *Equal-time comparisons of approximately 8 thread minutes each. The crops show the various parts of the medium, requiring multiple scattering events to exit and find the emitter. The top row shows a dense medium, and uniform lighting which provides the ideal scenario for PoE. BONO adapts to this quite well, while BILL is also able to reduce variance compared to Classical, but not to the level of BONO. The bottow row shows a translucent medium with a backlight providing a slight parallax effect, which BONO does not handle as well as BILL.*

3.4. PDF Evaluation for Sampled Slab Normals

A key consideration for BONO and BILL, where the Dwivedi slab normal $\vec{n}_D$ is sampled rather than deterministic, is the evaluation of the correct probability density function (PDF) for importance sampling weights. The true underlying PDF involves integrating the Dwivedi kernel over the distribution of potential slab normals, which is generally intractable.

To address this efficiently, we leverage a one-sample estimator based on Stochastic Multiple Importance Sampling (SMIS) [WGGH20] to approximate the required densities. As derived in the supplemental material, this approach yields a practical simplification: the effective density for the polar angle ω_z (or distance) given a chosen slab normal $\vec{n}_D$ reduces to the standard Dwivedi PDF conditioned *only* on that single sampled normal, i.e., $p_D(\omega_z|\vec{n}_D)$. This result applies analogously to free-flight distance sampling and holds for BONO, BILL, and II. The full derivation can be found in the supplemental.

4. Results and Discussion

We first evaluate the various components of our method. Section 4.2 evaluates the Burley diffusion profile as a weight-generating function for the CDFs. We follow that with a comparison of the performance given various resolutions of the guiding cache in Section 4.3. We perform the algorithmic performance evaluation comparisons in Section 4.4. We compare against *Classical* sampling, Dwivedi using the *PoE* or *II* slab normal, *VPG*, and our proposed *BONO* and *BILL* slab normals. When evaluating different slab normal variants, we assess each technique independently without combining them with Classical sampling through MIS. This allows us to isolate and directly compare their individual performance. VPG is combined with Classical as per the authors' recommendation in their work. VPG's scattering direction sampling requires combining with Classical using MIS because poorly fitted directional distributions can introduce sampling bias. The final part of the evaluation examines how our technique combines with *Classical* using MIS in Section 4.5.

4.1. Implementation

All our results were rendered on a machine with an AMD Ryzen 3950X 16-Core Processor running at 3.5GHz. We ran our renders using 32 threads. We implemented our technique in the *Mitsuba* renderer [Jak10]. The implementation includes our description of the volume boundary irradiance approximation. The vMF mixture fitting implementation is shared between our method and that used in VPG. For further details on parameter estimation, we refer the reader to Herholz et al.'s paper and supplemental [HZE*19]. Note that while our method constructs directional distributions at the boundary surface, these are not used for surface light transport and are only used as a sample space for the slab normals during volume light transport.

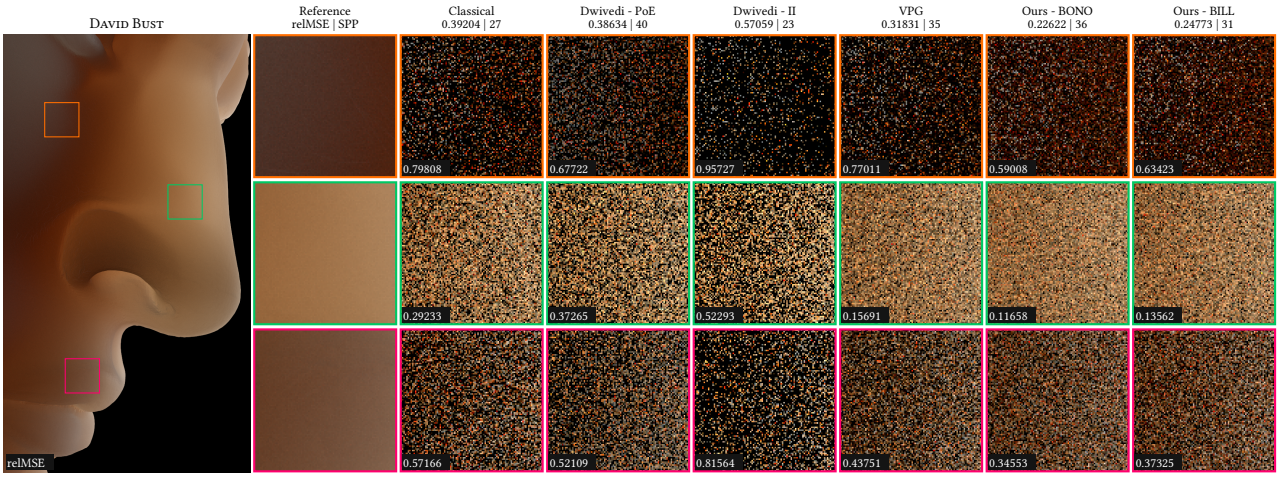

Figure 8: *Equal-time comparisons of about 16 thread minutes each. The low-frequency indirect illumination poses a problem for II, in which paths are guided to low-contribution regions and directions in volume. Our boundary cache guides the path to exit at a region where the boundary can connect to the indirect source of illumination.*

4.2. Evaluating the Burley Diffusion Profile

We use the Burley diffusion profile to compute the CDF weights to determine the distribution of energy on the boundary. The Burley diffusion profile is fitted using surface integration only, neglecting transport between surface and volume points or points on opposite sides of a volume boundary. It can inadequately represent the energy transport between points that deviate from the original geometric configuration. We conducted two experiments to validate whether the Burley profile remains viable for volume transport distribution estimation despite its theoretical limitations. The full details will be provided in the supplemental. To conduct the experiments, we designed a diffusion profile that extends the Burley diffusion profile by providing the diffusion value between a point on the surface and a point in the volume. We call this a *depth-aware tabulated diffusion profile*. Our first experiment measured fluence distribution at varying depths within a half-infinite slab and a finite slab, with a homogeneous volume. The depth-aware profile consistently matched reference Monte Carlo solutions, while the standard Burley profile underestimated fluence. However, all methods produced similar distribution shapes, differing primarily in magnitude. The second experiment showed the CDFs for some pixels using both the tabulated and Burley diffusion profiles. The CDFs were very similar because the diffusion profiles differ primarily in magnitude rather than shape, and CDF normalization eliminates these magnitude differences. Hence, the Burley diffusion profile is an adequate heuristic for volume transport distribution estimation.

4.3. Guiding Cache Spatial Resolution

This evaluation examines the performance of our algorithm with respect to the granularity of the spatial structure over the boundary. While fine caches accurately represent boundary distributions, dense ones can hurt performance due to CDF construction overhead; however, performance remains robust across granularities.

At finer resolutions, the CDF size can grow very large, for example, in the Envmap Buddha scene, a $40 \times 40 \times 40$ cache can have CDFs with 9000 records. Such a scenario is ideal for reservoir sampling, however, suboptimal uniform candidate generation and expensive evaluation (as the sampled normal is used for the whole path) reduce reservoir sampling's benefits. We were unable to find candidate generation techniques that would consider the volume transport. Our method needs high-quality individual samples due to costly evaluation.

Finally, caching CDFs is necessary as per-random walk construction is prohibitively complex. The illuminated CDFs in total require $O(n^2)$ time, and the unilluminated regions require $O(mn)$ time, where m is the number of unilluminated regions. Paying that cost per random walk is not feasible, and thus, the caching significantly improves the performance. Note that we do not fully occupy the 3D grid. A $40 \times 40 \times 40$ grid has a theoretical capacity of 64,000 records. However, our actual allocation requirements are significantly reduced since we restrict our coverage to only the boundary and use sparse-voxel hashing. In the Orc scene, for instance, a grid of this size ultimately requires only 1,262 illuminated regions, representing roughly just 2% of the maximum capacity.

4.4. Algorithmic Comparison

We present the comparison of our method with *Classical*, Dwivedi *Point of Entry*, Dwivedi *Incident Illumination* sampling and Volume Path Guiding. We present scenes that exemplify the strength of each method and compare their performance with ours. We note that we do not combine the Dwivedi strategies with *Classical* to ensure an equivalent evaluation of the performance of the sampling strategy compared with ours. Details on the material properties for the boundary and volume can be found in the supplemental.

Envmap Buddha Scene Figure 7 (top) shows the Buddha scene. The scene is lit with an environment map, and the strongest emissive part of the environment is the sun shining from the right, as observed in the highlight on the belly. The statue has a rough

dielectric boundary and an optically thick medium. This setup suits PoE, as paths should exit without going deep into the volume. II struggles as many paths are guided towards the sun. Our method performs comparably to *Classical*. By learning the distribution of energy on the boundary, we are not making uninformed guiding towards the strongest illumination. Instead, our method infers that we should steer paths back out from the same side of entry. The optically thick medium results in a very dense volume cache, which adds additional performance overhead to VPG.

BACKLIT BUDDHA Scene Figure 7 (bottom) shows the BACKLIT BUDDHA scene. The statue is backlit with a single source of illumination, similar to the setup presented in Meng et al [MHD16]. A smooth dielectric is used for the boundary, and the lighting setup requires paths to exit from the backside. Classical and PoE both perform poorly. Classical's uniform sampling is uninformed, and PoE steers the paths to the front, which means most paths will not exit to find an emission contribution. II correctly steers the path to the back-lit boundary area by sampling an emitter point. The optically thin volume in the BUDDHA scene combined with the smooth dielectric plays to VPG's strengths. The connection with the emitter on exit requires a refracted event, which VPG learns and can accurately represent in its internal volume caches. Additionally, VPG's guided distance sampling handles no-scatter events directly, thus reducing variance by directly transmitting through the volume. Our method correctly learns that the important boundary source is at the back and guides paths to exit there to connect with the emitter. BILL slightly outperforms BONO

DAVID BUST Scene In the DAVID BUST scene, the statue is backlit using only indirect lighting, as shown in Figure 8. The medium is optically thick, requiring multiple scattering events to traverse through the nose. The boundary is diffuse transmissive. Classical's unguided paths either exit at incorrect locations or bounce deeper into the volume. PoE's paths, uninformed about boundary energy, rarely exit on the side allowing an indirect illumination connection. This is particularly evident in thinner regions of the medium like the nose. II performs worst because the emission source is on the opposite side of the incident illumination source, above the statue, and does not directly light the statue. As a result, II paths are primarily guided to the top of the head, leading to low path throughput, and the paths exit on sides that receive no contribution from indirect lighting or direct lighting. VPG, learning the illumination's directional distribution, outperforms Classical and other Dwivedi methods, but not ours. Notably, it still struggles with thicker medium parts, highlighted in Figure 8's orange and pink crops. This medium is optically thinner than the ENVMAP BUDDHA scene, hence the issue is less exacerbated but still present. Our method correctly infers that the primary illumination source is on the opposite side of the face. Steering paths to the other side increases the likelihood of exiting BSDF-sampled paths landing on surfaces contributing to the indirect illumination reflector.

ORC Scene The ORC scene features a typical light rig used for dramatic character lighting in production shots. Two emitters positioned at the front illuminate the face from the sides, complemented by strong backlighting. Additionally, a targeted light is directed at the eye on the left side of the image to create a highlight on the iris and pupil, a common technique in headshots.

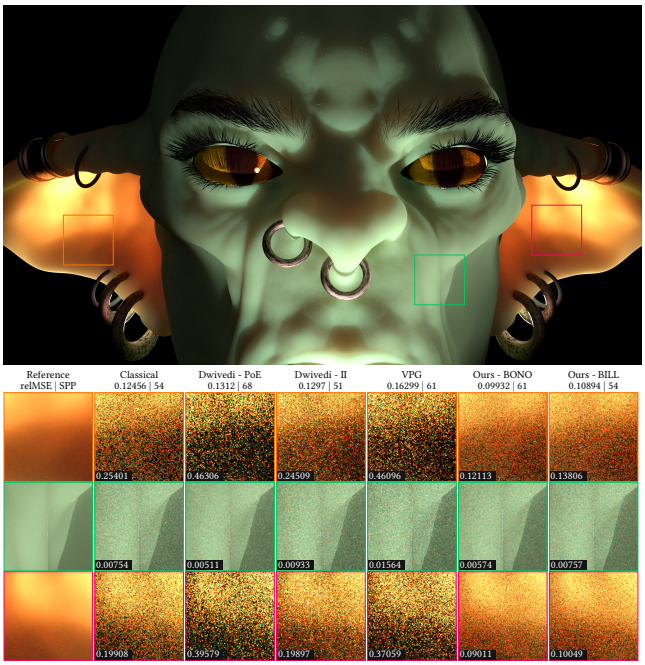

Figure 9: *Equal-time comparisons of about 2.5 thread hours each. The Dwivedi methods guide only the volume transport only. The crops show the various regions of the medium that require bespoke slab normal selection. Note how our method reduces variance on the ears and does not deteriorate on the flatter regions where PoE performs best.*

The volume is very optically thick, requiring multiple scattering events even in the ear to exit to the other side. The ears are lit only by the backlighting and any indirect light that bounces off the face itself. Classical's unguided paths do not favor any side of the ear, and paths in flatter, front-lit regions can penetrate deeper. PoE performs well on flatter face surfaces but poorly on ears as paths are steered opposite the incident illumination. II does not significantly improve ear path distribution, as many guided paths go to front emitters not directly illuminating ears.

The optically thick volume results in VPG building very dense internal caches. In our various experiments, we found that this is a drawback of VPG, especially the extended training time and memory requirements required to build the volume guiding cache. VPG needs to learn both the directional distribution and the spatial distribution means that it requires large amounts of training data, and these additional degrees of freedom offer more opportunities for error. The density of the medium adversely affects the performance of VPG, even though the medium is homogeneous. By limiting our structure to cover only the boundary, we overcome this limitation. Our methods correctly guide the paths to the regions of the boundary that receive energy. The ears' paths are directed backward, while front-lit face paths behave more like the *PoE* normal. Our method significantly improves analytical Dwivedi guiding by providing a slab normal that more accurately describes the most significant illumination source direction. Figure 6 shows a com-

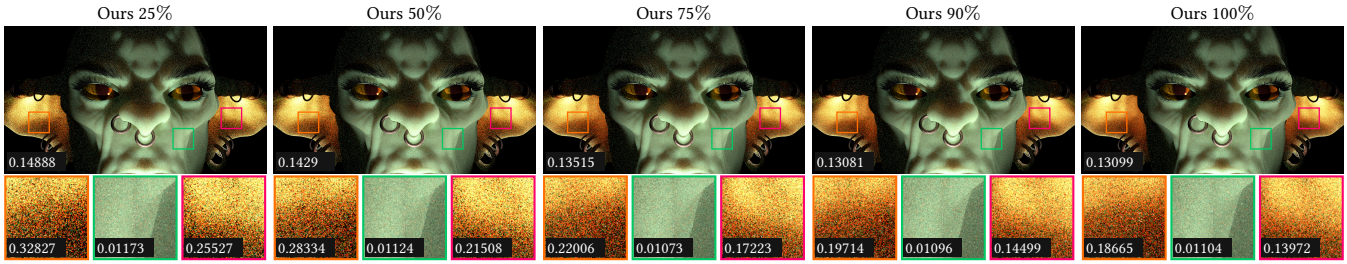

Figure 10: *Orc MIS comparisons show our method is best not combined with Classical sampling.*

parison of the distribution of the slab normals depending on the technique used.

4.5. Multiple Importance Sampling

Previous approaches in Dwivedi sampling [Kd14; MHD16; KGV*20] typically combine analytical techniques with Classical sampling via MIS to mitigate fireflies from low-probability path events. However, our analysis shows that since our method generalizes other Dwivedi slab normal approaches and outperforms Classical in most scenarios, combining it with Classical using MIS deteriorates our results (Figure 10 and supplemental). We ran equal render samples, increasing the probability of selecting our technique.

Meng et al. enforce Classical sampling for initial free-flight distances [MHD16]. However, we observe limitations in this approach, particularly for optically and geometrically thin, backlit media. Stretching the initial free-flight distance via the Dwivedi distribution increases the sampling probability for no-scatter events, a dominant energy transport component in such configurations.

4.6. Which method to use?

Whilst both methods offer improvements over existing Dwivedi sampling techniques, BONO demonstrates superior performance relative to BILL in the majority of evaluated scenarios, with the exception of the isolated case presented in Figure 7 (bottom row, top crop). Despite this singular underperformance, BONO maintains better robustness across all other evaluated contexts. Based on these results, BONO may be more suitable for a wider range of applications. The simpler training methodology and elimination of directional distribution requirements (along with the associated fitting procedure) result in easier integration into existing rendering pipelines. On the other hand, renderers with existing surface path guiding schemes that have directional distributions on the surface may find BILL to be an easier drop-in method.

5. Limitations and Future Work

Guided Russian Roulette and Splitting. VPG's Guided RR enables more samples to be evaluated for the same render time. Our method does not have an adjoint estimate as accurate as that of VPG; however, we believe some form of Guided RR and Splitting using the radiance field information at the surface can still lead to variance reduction. This follows from the ZVT principle, which states that an approximation of the solution can be enough of a catalyst to reduce variance.

Heterogeneous Media. The Dwivedi sampling kernels are derived from an adjoint solution for homogeneous media, making them unpredictable when handling complex heterogeneous media interactions. Solutions for heterogeneous media are only available under the assumption of homogenization of the medium [Dwi82a].

Anisotropic Phase Functions. Meng et al. demonstrate the use of Dwivedi with anisotropic phase functions; however, as the mean cosine gets closer to the endpoints, the results deteriorate [MHD16]. D'Eon hints at possible analytical solutions [KGV*20] that can be derived from previous linear transport work [MMN*80].

Refractive Boundaries. VPG effectively resolves the energy concentration in the volume caused by refraction due to the internal volumetric adjoint approximation. Our method does not have this fine-grained information, which leads to less informed sampling decisions. Learning the refracted boundary exit directions can prove useful to guide volume paths when close to the boundary. There are also existing techniques such as Subdivision Next Event Estimation [KNK*16] and Manifold Next Event Estimation [HDF15], that can construct paths that connect with emitters through refractive boundaries.

6. Conclusions

Our core idea is to inform analytical Dwivedi sampling with information about boundary illumination to reduce variance under complex lighting and in the context of homogeneous, isotropic media. We approximate the incident illumination by a sparse grid structure that stores an averaged normal and directional distribution, per region. During rendering, we first importance-sample a voxel at the boundary based on volume transport properties, using an analytical diffusion profile as an approximation of volume transport. Then we either use the voxel's averaged normal from all illuminated training data points or the direction proportional to the incident illumination at the surface, and interpret it as the slab normal for the Dwivedi kernels. Under complex lighting scenarios and indirect illumination, our new *BONO* and *BILL* methods outperform existing Dwivedi guiding techniques, Classical Volume Path Tracing, and, in some cases, even Volume Path Guiding.

7. Acknowledgements

We want to thank Sebastian Herholz for sharing his Mitsuba Volume Path Guiding implementation with us. We also want to thank Vincent Pegoraro and Farhan Wali for numerous discussions. We acknowledge the following persons and institutions for providing some of the models and scenes used in this work: Stanford 3D scanning repository (BUDDHA), Statens Museum for Kunst (DAVID) and Adam McAnderson (ORC). We acknowledge funding by Wētā FX, the Charles University Grant Number SVV 112423 and the Grantová Agentura České Republiky with grant number GAČR-22-22875S. Open access publishing facilitated by Univerzita Karlova, as part of the Wiley - CzechELib agreement.

References

[AK90] ARVO, JAMES and KIRK, DAVID. "Particle transport and image synthesis". *Proc. SIGGRAPH '90*. Dallas, TX, USA: ACM, 1990, 63–66. ISBN: 0-89791-344-2. DOI: 10.1145/97879.97886 4.

[BDC12] BASHFORD-ROGERS, THOMAS, DEBATTISTA, KURT, and CHALMERS, ALAN. "A Significance Cache for Accelerating Global Illumination". *Computer Graphics Forum* 31.6 (2012), 1837–51 4.

[Cha60] CHANDRASEKHAR, SUBRAHMANYAN. *Radiative Transfer*. Dover Publications, 1960 2.

[Chr15] CHRISTENSEN, PER H. "An approximate reflectance profile for efficient subsurface scattering". *ACM SIGGRAPH 2015 Talks*. SIGGRAPH '15. Los Angeles, California: Association for Computing Machinery, 2015. ISBN: 9781450336369. DOI: 10.1145/2775280.2792555 4.

[DI11] D'EON, EUGENE and IRVING, GEOFFREY. "A quantized-diffusion model for rendering translucent materials". *ACM Transactions on Graphics* 30.4 (July 2011), 56:1–56:14. ISSN: 0730-0301. DOI: 10.1145/2010324.1964951. (Visited on 01/10/2022) 4.

[DK17] DAHM, KEN and KELLER, ALEXANDER. "Learning light transport the reinforced way". *ACM SIGGRAPH 2017 Talks*. SIGGRAPH '17. Los Angeles, California: Association for Computing Machinery, July 2017, 1–2. ISBN: 978-1-4503-5008-2. DOI: 10.1145/3084363.3085032. (Visited on 05/15/2020) 4.

[Dwi82a] DWIVEDI, S. R. "A new importance biasing scheme for deep-penetration Monte Carlo". *Annals of Nuclear Energy* 9.7 (1982). Publisher: Elsevier, 359–368 2, 3, 10.

[Dwi82b] DWIVEDI, S. R. "Zero Variance Biasing Schemes for Monte Carlo Calculations of Neutron and Radiation Transport Problems". *Nuclear Science and Engineering* 80.1 (1982), 172–178. DOI: 10.13182/NSE82-A21413 2.

[DWL23] DONG, HONGHAO, WANG, GUOPING, and LI, SHENG. "Neural Parametric Mixtures for Path Guiding". *ACM SIGGRAPH 2023 Conference Proceedings*. SIGGRAPH '23. Los Angeles, CA, USA: Association for Computing Machinery, 2023. ISBN: 9798400701597. URL: https://doi.org/10.1145/3588432.3591533 4.

[DWWH20] DENG, HONG, WANG, BEIBEI, WANG, RUI, and HOLZSCHUCH, NICOLAS. "A practical path guiding method for participating media". *Computational Visual Media* 6.1 (2020), 37–51. DOI: 10.1007/s41095-020-0160-1 4.

[FHK15] FRISVAD, JEPPE REVALL, HACHISUKA, TOSHIYA, and KJELDSEN, THOMAS KIM. "Directional Dipole Model for Subsurface Scattering". *ACM Trans. Graph.* 34.1 (Dec. 2015). ISSN: 0730-0301. DOI: 10.1145/2682629. URL: https://doi.org/10.1145/2682629 4.

[HDF15] HANIKA, JOHANNES, DROSKE, MARC, and FASCIONE, LUCA. "Manifold Next Event Estimation". *Computer Graphics Forum* (2015). DOI: 10.1111/cgf.12681 10.

[HEV*16] HERHOLZ, SEBASTIAN, ELEK, OSKAR, VORBA, JIŘÍ, et al. "Product Importance Sampling for Light Transport Path Guiding". en. *Computer Graphics Forum* 35.4 (2016), 67–77. ISSN: 1467-8659. DOI: 10.1111/cgf.12950 4.

[HIT*24] HUANG, JIAWEI, IIZUKA, AKITO, TANAKA, HAJIME, et al. "Online Neural Path Guiding with Normalized Anisotropic Spherical Gaussians". *ACM Trans. Graph.* 43.3 (Apr. 2024). ISSN: 0730-0301. URL: https://doi.org/10.1145/3649310 4.

[HMB*15] HILL, STEPHEN, MCAULEY, STEPHEN, BURLEY, BRENT, et al. "Physically Based Shading in Theory and Practice". *ACM SIGGRAPH 2015 Courses*. SIGGRAPH '15. Los Angeles, California: Association for Computing Machinery, 2015. ISBN: 9781450336345. DOI: 10.1145/2776880.2787670 4.

[Hoo08] HOOGENBOOM, J. EDUARD. "Zero-Variance Monte Carlo Schemes Revisited". *Nuclear Science and Engineering* 160.1 (Sept. 2008). Publisher: Taylor & Francis _eprint: https://doi.org/10.13182/NSE160-01, 1–22. ISSN: 0029-5639. DOI: 10.13182/NSE160-01 2.

[HZE*19] HERHOLZ, SEBASTIAN, ZHAO, YANGYANG, ELEK, OSKAR, et al. "Volume Path Guiding Based on Zero-Variance Random Walk Theory". *ACM Trans. Graph.* 38.3 (June 2019), 25:1–25:19. ISSN: 0730-0301. DOI: 10.1145/3230635 2, 4, 7.

[Jak10] JAKOB, WENZEL. *Mitsuba renderer*. http://www.mitsuba-renderer.org. 2010 7.

[JB02] JENSEN, HENRIK WANN and BUHLER, JUAN. "A rapid hierarchical rendering technique for translucent materials". *ACM Trans. Graph.* 21.3 (July 2002), 576–581. ISSN: 0730-0301. DOI: 10.1145/566654.566619 4.

[Jen95] JENSEN, HENRIK WANN. "Importance Driven Path Tracing using the Photon Map". *Rendering Techniques '95*. Ed. by HANRAHAN, PATRICK M. and PURGATHOFER, WERNER. Vienna: Springer Vienna, 1995, 326–335. ISBN: 978-3-7091-9430-0 4.

[JMLH01] JENSEN, HENRIK WANN, MARSCHNER, STEPHEN R., LEVOY, MARC, and HANRAHAN, PAT. "A practical model for subsurface light transport". en. *Proceedings of the 28th annual conference on Computer graphics and interactive techniques - SIGGRAPH '01*. Not Known: ACM Press, 2001, 511–518. ISBN: 978-1-58113-374-5. DOI: 10.1145/383259.383319 4.

[Kd14] KŘIVÁNEK, JAROSLAV and D'EON, EUGENE. "A zero-variance-based sampling scheme for Monte Carlo subsurface scattering". *ACM SIGGRAPH 2014 Talks*. SIGGRAPH '14. New York, NY, USA: Association for Computing Machinery, July 2014, 1. ISBN: 978-1-4503-2960-6. DOI: 10.1145/2614106.2614138 2, 3, 10.

[KGV*20] KELLER, ALEXANDER, GRITTMANN, PASCAL, VORBA, JIŘÍ, et al. "Advances in Monte Carlo rendering: the legacy of Jaroslav Křivánek". *ACM SIGGRAPH 2020 Courses*. SIGGRAPH '20. New York, NY, USA: Association for Computing Machinery, Aug. 2020, 1–366. ISBN: 978-1-4503-7972-4. DOI: 10.1145/3388769.3407458. (Visited on 10/28/2022) 3, 10.

[KNK*16] KOERNER, DAVID, NOVAK, JAN, KUTZ, PETER, et al. *Subdivision Next-Event Estimation for Path-Traced Subsurface Scattering*. en. Accepted: 2016-06-17T14:12:43Z ISSN: 1727-3463. The Eurographics Association, 2016. ISBN: 978-3-03868-019-2. DOI: 10.2312/sre.20161214 10.

[LHW21] LEONARD, L., HÖHLEIN, K., and WESTERMANN, R. "Learning Multiple-Scattering Solutions for Sphere-Tracing of Volumetric Subsurface Effects". *Computer Graphics Forum* 40.2 (2021), 165–178. DOI: https://doi.org/10.1111/cgf.142623 4.

[MGN17] MÜLLER, THOMAS, GROSS, MARKUS, and NOVÁK, JAN. "Practical Path Guiding for Efficient Light-Transport Simulation". *Computer Graphics Forum* 36.4 (June 2017), 91–100. ISSN: 1467-8659. DOI: 10.1111/cgf.13227 4, 5.

[MHD16] MENG, JOHANNES, HANIKA, JOHANNES, and DACHSBACHER, CARSTEN. "Improving the Dwivedi Sampling Scheme". *Computer Graphics Forum* 35.4 (July 2016), 37–44. ISSN: 0167-7055 2, 3, 9, 10.

[MMN*80] MARCHUK, GURI I., MIKHAILOV, GENNADI A., NAZARALIEV, MAGAMEDSHAFI A., et al. *The Monte Carlo Methods in Atmospheric Optics*. Ed. by MACADAM, DAVID L. Vol. 12. Springer Series in Optical Sciences. Berlin, Heidelberg: Springer Berlin Heidelberg, 1980. DOI: 10.1007/978-3-540-35237-2. (Visited on 11/04/2022) 10.

[MMR*19] MÜLLER, THOMAS, MCWILLIAMS, BRIAN, ROUSSELLE, FABRICE, et al. "Neural Importance Sampling". *ACM Transactions on Graphics* 38.5 (Oct. 2019), 145:1–145:19. ISSN: 0730-0301. DOI: 10.1145/3341156. (Visited on 04/29/2020) 4.

[PJH16] PHARR, MATT, JAKOB, WENZEL, and HUMPHREYS, GREG. *Physically Based Rendering: From Theory to Implementation*. 3rd. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2016. ISBN: 0128006455 3.

[RGH*22] RATH, ALEXANDER, GRITTMANN, PASCAL, HERHOLZ, SEBASTIAN, et al. "EARS: Efficiency-Aware Russian Roulette and Splitting". *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2022)* 41.4 (July 2022). DOI: 10.1145/3528223.3530168 4.

[RHL20] RUPPERT, LUKAS, HERHOLZ, SEBASTIAN, and LENSCH, HENDRIK P. A. "Robust Fitting of Parallax-Aware Mixtures for Path Guiding". *ACM Trans. Graph.* 39.4 (Aug. 2020). ISSN: 0730-0301. DOI: 10.1145/3386569.3392421 4.

[SHJD18] SIMON, FLORIAN, HANIKA, JOHANNES, JUNG, ALISA, and DACHSBACHER, CARSTEN. "Selective guided sampling with complete light transport paths". *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 37.6 (Dec. 2018). DOI: 10.1145/3272127.3275030 4.

[TTJ*24] TG, THOMSON, TRAN, DUC MINH, JENSEN, HENRIK W., et al. "Neural SSS: Lightweight Object Appearance Representation". *Computer Graphics Forum* (2024). ISSN: 1467-8659. DOI: 10.1111/cgf.15158 4.

[VHH*19] VORBA, JIŘÍ, HANIKA, JOHANNES, HERHOLZ, SEBASTIAN, et al. "Path Guiding in Production". *ACM SIGGRAPH 2019 Courses*. SIGGRAPH '19. event-place: Los Angeles, California. New York, NY, USA: ACM, 2019, 18:1–18:77. ISBN: 978-1-4503-6307-5. DOI: 10.1145/3305366.3328091 2, 4.

[VK16] VORBA, JIŘÍ and KŘIVÁNEK, JAROSLAV. "Adjoint-Driven Russian Roulette and Splitting in Light Transport Simulation". *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2016)* 35.4 (July 2016) 4.

[VKJ19] VICINI, DELIO, KOLTUN, VLADLEN, and JAKOB, WENZEL. "A Learned Shape-Adaptive Subsurface Scattering Model". *ACM Transactions on Graphics* 38.4 (July 2019), 127:1–127:15. ISSN: 0730-0301. DOI: 10.1145/3306346.3322974 4.

[VKŠ*14] VORBA, JIŘÍ, KARLÍK, ONDŘEJ, ŠIK, MARTIN, et al. "On-line Learning of Parametric Mixture Models for Light Transport Simulation". *ACM Trans. Graph.* 33.4 (July 2014), 101:1–101:11. ISSN: 0730-0301. DOI: 10.1145/2601097.2601203 4.

[WGGH20] WEST, REX, GEORGIEV, ILIYAN, GRUSON, ADRIEN, and HACHISUKA, TOSHIYA. "Continuous Multiple Importance Sampling". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 39.4 (July 2020). DOI: 10.1145/3386569.3392436 7.

[WWH*24] WU, WENSHI, WANG, BEIBEI, HAŠAN, MILOŠ, et al. "Efficient participating media rendering with differentiable regularization". *Computational Visual Media* 10.5 (2024), 937–948. DOI: 10.1007/s41095-023-0372-2 4.