

Neural Path Guiding with Distribution Factorization

Pedro Figueiredo^{id}, Qihao He^{id} and Nima Khademi Kalantari^{id}

Texas A&M University

Abstract

In this paper, we present a neural path guiding method to aid with Monte Carlo (MC) integration in rendering. Existing neural methods utilize distribution representations that are either fast or expressive, but not both. We propose a simple, but effective, representation that is sufficiently expressive and reasonably fast. Specifically, we break down the 2D distribution over the directional domain into two 1D probability distribution functions (PDF). We propose to model each 1D PDF using a neural network that estimates the distribution at a set of discrete coordinates. The PDF at an arbitrary location can then be evaluated and sampled through interpolation. To train the network, we maximize the similarity of the learned and target distributions. To reduce the variance of the gradient during optimizations and estimate the normalization factor, we propose to cache the incoming radiance using an additional network. Through extensive experiments, we demonstrate that our approach is better than the existing methods, particularly in challenging scenes with complex light transport.

1. Introduction

Accurate simulation of light transport necessitates computing a complex multidimensional integral. Monte Carlo (MC) rendering, a common approach for physically-based light transport simulation, estimates this integral by averaging samples of the integrand. While this approach can produce compelling effects, the variance of the estimator decreases linearly with the number of samples. Therefore, accurate estimation of the integral requires evaluating a large number of expensive samples (paths), making the approach slow.

An effective way of reducing the variance of MC rendering is through importance sampling. Specifically, the goal of importance sampling is to sample the paths with higher contribution to the integral more frequently. A large number of methods [GBBE18, VKv*14, MGN17] focus on devising strategies to approximate sampling distributions, typically in an online fashion, that are proportional to the integrand. Particularly, the *path guiding* approaches adopt a local view and perform the guided sampling at every intersection point along a path. These methods use efficient data structures, such as k-d tree [DPÖM22], octree [BB17], and quadtree [MGN17], as well as parametric models like Gaussian mixture model (GMM) [HEV*16, VKv*14], and von Mises-Fisher (vMF) mixture models [RHL20], to estimate the distributions in a data-driven manner. Most of these methods learn a marginalized distribution within each subdivision, thus failing to capture the spatial variations in each subarea.

More recently, a few techniques [MMR*19, DWL23, HIT*24] overcome the limitations in explicit representations by utilizing neural networks. The network in these methods takes a continuous coordinate (spatial as well as directional in case of product sampling) as the input, and estimates the distribution at that par-

ticular location. The key difference between these methods is in how they model the probability distribution function (PDF). Specifically, Müller et al. [MMR*19] use normalizing flows [DSB17], an expressive model with the ability to represent detailed distributions at the cost of being slower and more difficult to optimize. Dong et al. [DWL23] and Huang et al. [HIT*24], on the other hand, use mixture models that are relatively easy to optimize. In practice, these models are used with only a few components to ensure efficiency. Unfortunately, such models are not able to handle challenging light transport scenarios.

In this paper, we propose a simple representation that is sufficiently expressive and reasonably fast. Specifically, we represent the joint PDF over the 2D directional domain in terms of the product of two 1D distributions: a marginal and a conditional PDF. We model each distribution using a separate neural network that estimates the PDF at uniformly spaced discrete locations. The PDF at an arbitrary coordinate is obtained by interpolating the estimated values. We explore two simple interpolation strategies and present approaches to evaluate and sample from our learned representation.

We train our networks by minimizing the loss between the predicted and target distributions. Since the target distribution is unknown, existing methods [DWL23, MMR*19] use its MC samples without normalization during optimization. Because of this, the gradients are highly noisy, making the training unstable. To address this issue, we identify the incoming radiance as the source of the variance, and propose to cache it using an additional network which is in turn used to estimate a better target distribution. We demonstrate that the combination of radiance caching and our simple PDF representation is highly effective and outperforms the existing techniques, particularly in challenging scenes.

2. Related Work

Global Path Guiding These approaches attempt to sample a full path (from camera to light sources) at once according to the distribution of the integrand. This is typically done in the primary sample space (PSS), which provides an elegant way of parameterizing a path, while ensuring each vertex lies on the surfaces. Various approaches propose to do this using k-d trees [GBBE18] or neural networks [ZZ19, MMR*19]. However, these approaches suffer from the curse of dimensionality and their performance degrade with longer paths [MMR*19]. Specialized guiding methods have been developed to efficiently deal with glossy interactions, offering improved performance in specific scenarios [LWT*22, FHG*23], which is not the primary focus of our framework.

Local Path Guiding An alternative method is to perform guiding at every intersection point along the path. Early approaches represent the guiding distribution using a 2D histogram and store it through photon mapping [Jen95] or a 5D tree [LW95]. Follow up work refine this strategy through a combination of BRDF and photon map before sampling [SL06] and improvements in caustics sampling [BAJ08]. We also use a discretized representation, but factorize it to two 1D distributions and estimate them using neural networks. Moreover, we present evaluation and sampling using linear interpolation, in addition to the commonly-used nearest neighbor strategy.

A variety of other representations including hemispherical particle footprint [HP02], cosine lobes [BRDC12], B-splines [PWP08], binary and quad tree [MGN17, VHH*19, RGH*20, ZXS*21a], double trees [BB17], and discretized hemispheres using reinforcement learning [DK17] have also been proposed. Moreover, several methods use parametric models, such as Gaussian mixture [VKv*14, HEV*16, DPÖM22, VK16] and von Mises-Fisher (vMF) mixture models [HZE*19, RHL20]. While most of these approaches fail to capture the spatial variations due to their limited spatial partitioning, a couple of methods address this issue [RHL20, DPÖM22]. Ruppert et al. [RHL20] introduce a parallax compensation strategy that readjusts vMF lobes to point to a virtual light source. Unfortunately, they do not account for occlusion in such adjustments. Dodik et al. [DPÖM22] propose a higher-dimensional representation that can perform product sampling, but this comes at the cost of increased computational cost.

More recently, path guiding methods utilizing neural networks have emerged. Instead of using complex data structures, a few techniques [MMR*19, LHL*24, DWL23, HIT*24] use neural networks to estimate the distributions in an online manner. Particularly, a couple of methods [MMR*19, LHL*24] use normalizing flows [DSB17] to map a known distribution (e.g., uniform) to the target distribution. Unfortunately, evaluating and sampling the PDF from normalizing flows is computationally expensive. A couple of methods attempt to address this problem by using vMF [DWL23] and normalized anisotropic spherical Gaussian (NASG) [HIT*24] mixtures. The main problem with these two approaches is that the mixture model representations have limited expressiveness, and thus struggle in scenes with complex light transport. In contrast to all the discussed online methods, a few approaches propose to perform path guiding using convolutional neural networks (CNN) [BMDS19, HWZ*20, ZXS*21a, ZXS*21b]



Figure 1: Method Overview. During path tracing, we generate samples using our guiding distribution p_Θ to increase the number of paths that reach light sources. These paths are then used to train our radiance caching f_Φ . We leverage the cached radiance as a smoother objective to improve p_Θ , which in turn is used on the next sample generation.

in an offline manner. These methods train a CNN on a large number of scenes and use the trained network to adaptively refine the distribution [HWZ*20] or denoise/improve the guiding maps [BMDS19, ZXS*21a, ZXS*21b] during rendering. In addition to having a large computational cost associated with the CNNs, some of these techniques [BMDS19, HWZ*20] only allow guiding at the first intersection.

Radiance Caching Since the seminal work by Ward et al. [WRC88], a plethora of methods aiming for both offline [GKB09, DBN17, MJJG18, ZBN19] and real-time [BFK18, Pan20, SNRS12, RZD14, SL17, VPG14] rendering have been proposed. The key challenge in radiance caching is placement of the cache records and interpolation, particularly in presence of glossy surfaces. The recent neural radiance caching (NRC) [MRNK21] tackles these challenges by caching the reflected radiance along an outgoing direction at each surface point using a highly efficient implementation of a multilayer perceptron (MLP). We use this approach to cache the incoming radiance, but in the context of path guiding, and for reducing the variance of the gradients and estimating the normalization factor.

Histogram Prediction A couple of methods use discretized histograms to model PDFs in closely related problems. For example, Zhu et al. [ZBX*21] learn two-dimensional discrete importance weights to sample complex luminaires from different perspectives. Xu et al. [XWH*23] extend this idea by decomposing the histogram into weighted combinations of shared-basis histograms, allowing a more accurate representation of diverse PDF lobes in the context of neural materials. In both cases, optimization is performed as a preprocessing step, allowing ample time and data samples to effectively learn complex target PDFs. However, for online path guiding, it is critical to balance expressiveness with training and inference costs. This motivates our efficient formulation, which represents the joint 2D PDF as a product of two 1D distributions.

3. Method

We begin by formulating the light transport in a local manner using the rendering equation [Kaj86]. Based on this formulation, the outgoing radiance L_o at a surface point \mathbf{x} in direction ω_o can be written in terms the emitted L_e and incident L_i radiance as follows:

$$L_o(\mathbf{x}, \omega_o) = L_e(\mathbf{x}, \omega_o) + \int_{\Omega} \rho(\mathbf{x}, \omega_o, \omega) L_i(\mathbf{x}, \omega) |\cos(\theta)| d\omega, \quad (1)$$

where ρ is the bidirectional scattering distribution function (BSDF), Ω is the unit sphere on \mathbf{x} containing all the incoming directions

ω_i , and θ is the angle between the incoming direction and surface normal at \mathbf{x} . The key challenge here is estimating reflected radiance by computing the integral term which is typically done through MC integration as follows:

$$\langle L_r(\mathbf{x}, \omega_o) \rangle = \frac{1}{N} \sum_{i=1}^N \frac{\rho(\mathbf{x}, \omega_o, \omega_i) L_i(\mathbf{x}, \omega_i) |\cos(\theta_i)|}{\hat{p}(\omega_i | \mathbf{x}, \omega_o)}, \quad (2)$$

where $\hat{p}(\omega_i | \mathbf{x}, \omega_o)$ is the distribution from which the sampled incoming directions ω_i are drawn from. To lower the variance, this distribution should be proportional to the integrand, i.e., $\hat{p}(\omega | \mathbf{x}, \omega_o) \propto \rho(\mathbf{x}, \omega_o, \omega) L_i(\mathbf{x}, \omega) |\cos(\theta)|$.

Our goal is to model the guiding distribution using a neural network, $\hat{p}_{\Theta}(\omega | \mathbf{x}, \omega_o)$ where Θ represents the network parameters, and train it in a data-driven manner. In the following sections, we explain our distribution model and optimization strategies. The overview of our entire pipeline is illustrated in Fig. 1.

3.1. Distribution Factorization

The success of path guiding methods depends in part on the choice of PDF representation. For such a representation to be effective, it should possess two main properties. 1) It should be sufficiently expressive to match the shape of the integrand (target distribution), even in complex settings. 2) Generating samples ω_i from the distribution and evaluating their PDF $\hat{p}_{\Theta}(\omega_i | \mathbf{x}, \omega_o)$, which are both necessary for MC integration (Eq. 2), should be fast.

To this end, we propose to model the distribution explicitly, by estimating the PDF at a set of discrete locations and obtaining the continuous PDF through interpolation. To ensure the discretization is done uniformly over the directional domain, we define the PDF in uniform square space with coordinates ϵ_1 and ϵ_2 ranging between 0 to 1. These coordinates can be easily converted to spherical domain through $\phi = 2\pi\epsilon_1$ and $\theta = \cos^{-1}(1 - 2\epsilon_2)$. The PDF defined over the uniform square space $\hat{p}(\epsilon_1, \epsilon_2 | \mathbf{x}, \omega_o)$ can be converted to distribution over ω , $\hat{p}(\omega | \mathbf{x}, \omega_o)$, by taking the Jacobian of the transformation into account.

Estimating the PDF over a grid of size $M_1 \times M_2$ (discretization resolution) using a neural network could potentially be computationally expensive. To avoid this problem, we use the product rule to represent the joint PDF over the uniform square domain as the multiplication of a marginal and a conditional PDF as follows:

$$\hat{p}_{\Theta}(\epsilon_1, \epsilon_2 | \mathbf{x}, \omega_o) = \hat{p}_{\mathbf{w}_1}(\epsilon_1 | \mathbf{x}, \omega_o) \hat{p}_{\mathbf{w}_2}(\epsilon_2 | \epsilon_1, \mathbf{x}, \omega_o). \quad (3)$$

Through this representation, estimating the joint PDF boils down to predicting two 1D distributions. We discretize the two domains, ϵ_1 and ϵ_2 , into M_1 and M_2 discrete locations, respectively, and use two independent networks to estimate the PDF at those locations. Specifically, one network $f_{\mathbf{w}_1}$ models $\hat{p}_{\mathbf{w}_1}(\epsilon_1 | \mathbf{x}, \omega_o)$ by taking \mathbf{x} and ω_o as the input and estimating an M_1 dimensional vector \mathbf{v}_1 , while the other network $f_{\mathbf{w}_2}$ approximates $\hat{p}_{\mathbf{w}_2}(\epsilon_2 | \epsilon_1, \mathbf{x}, \omega_o)$ by taking ϵ_1 , in addition to \mathbf{x} and ω_o , and predicting an M_2 dimensional vector \mathbf{v}_2 . The PDF at an arbitrary location can then be obtained by interpolating the estimated PDFs at discrete locations.

Care should be taken when estimating the PDF vectors to ensure

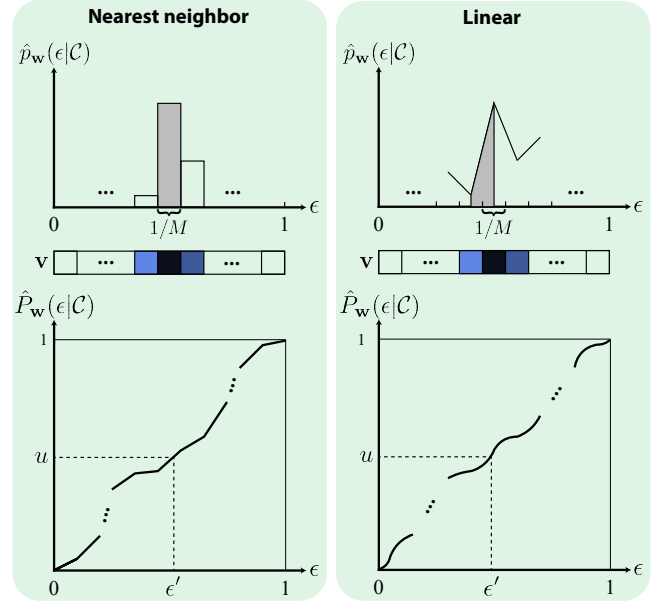


Figure 2: We demonstrate the PDF evaluation and sampling process for nearest neighbor and linear interpolation. To model the 1D distributions (marginal and conditional in Eq. 3) our network first predicts a vector \mathbf{v} containing estimates of the PDF at discrete locations. To obtain the PDF at an arbitrary location, we either use the PDF estimate at the closest sample (top-left), or linearly interpolate between the two closest samples (top-right). Sampling is done by evaluating the inverse CDF at a randomly generated value u with a uniform distribution. Note that the CDF for nearest neighbor interpolation is piecewise linear, while it is piecewise quadratic for linear interpolation.

the interpolated continuous functions are valid distributions and integrate to one. Below, we discuss our network design to estimate valid PDFs for two interpolation strategies: nearest neighbor and linear. We explain each case in a general setting since the two networks are similar and only differ in their input conditions. In Sec. 5, we compare the performance of these two methods.

Nearest Neighbor In this case, as shown in Fig. 2 (top-left), the domain is divided into a set of M bins and the PDF inside each bin is obtained from the corresponding element of the estimated vector \mathbf{v} as follows:

$$\hat{p}_{\mathbf{w}}(\epsilon | \mathcal{C}) = \mathbf{v}[\lfloor \epsilon M \rfloor]. \quad (4)$$

Here, we drop the subscripts since the process for the two PDFs are similar. Moreover, \mathcal{C} refers to the condition which is different for the two PDF terms in Eq. 3. To have a valid PDF, the cumulative distribution function (CDF) at one ($\hat{P}_{\mathbf{w}}(1)$) should be equal to one, i.e., the PDF should integrate to one. We can formally write this as:

$$\hat{P}_{\mathbf{w}}(1) = \int_0^1 \hat{p}_{\mathbf{w}}(\epsilon | \mathcal{C}) d\epsilon = \sum_{i=0}^{M-1} \mathbf{v}[i] \frac{1}{M} = 1, \quad (5)$$

where the summation calculates the integral by computing the area underneath each bar, one of which shaded in gray in Fig. 2 (top-left). For this equation to hold, the elements of the output vector




Figure 3: Our multilayer perceptron (MLP) takes the condition C as the input and estimates an M dimensional vector. We then apply softmax function to this vector and multiply each element by M to obtain a vector containing the PDF estimate at discrete locations. We use this network to model the marginal and conditional distributions in Eq. 3. Note that we use \mathbf{x} and ω_o as the condition when modeling the marginal distribution, but for the conditional one, we additionally pass the first dimension ϵ_1 to the network. We use two separate networks to model the two distributions, but we follow the process illustrated in this figure in both cases.

should sum up to M , i.e., $\sum_{i=0}^{M-1} \mathbf{v}[i] = M$. To achieve this, we obtain \mathbf{v} by applying softmax to the output of the network and multiplying each element by M as follows (see Fig. 3):

$$\mathbf{v} = M \sigma(f_w(C)), \quad (6)$$

where σ is the standard softmax operator. We can thus evaluate the PDF at an arbitrary location through combination of Eqs. 4 and 6.

To sample from the learned distribution $\hat{p}(\epsilon|C)$, we use the inverse transform sampling technique. As shown in Fig. 2 (bottom-left), we evaluate the inverse CDF at a randomly generated number with uniform distribution u , i.e. $\epsilon' = \hat{P}_w^{-1}(u|C)$. The exact formulation of the inverse CDF is provided in the supplementary material.

Linear Here, the PDF at an arbitrary location is obtained by linearly interpolating the estimated PDF at the two nearest discrete coordinates as follows (see Fig. 2 top-right):

$$\hat{p}_w(\epsilon|C) = (1 - \alpha) \mathbf{v}[\lfloor m \rfloor] + \alpha \mathbf{v}[\lceil m \rceil], \quad (7)$$

where

$$m = \epsilon M - 0.5, \quad \text{and} \quad \alpha = m - \lfloor m \rfloor. \quad (8)$$

Note that the subtraction of 0.5 is made to ensure the elements of the vector \mathbf{v} provide PDF values at the center of each bin. Again, to ensure \hat{p}_w is a valid PDF, it should integrate to one, i.e.,

$$\hat{P}_w(1) = \int_0^1 \hat{p}_w(\epsilon|C) d\epsilon = \sum_{i=0}^{M-1} \frac{\mathbf{v}[i] + \mathbf{v}[i+1]}{2} \frac{1}{M} = 1. \quad (9)$$

In this case, the integral is calculated by summation each trapezoid's area, one of which shaded in gray in Fig. 2 (top-right).

Note that the boundaries ($\epsilon < \frac{1}{2M}$ and $\epsilon > 1 - \frac{1}{2M}$) should be handled with care. We treat these regions differently when dealing

with ϵ_1 and ϵ_2 . For ϵ_1 (corresponding to spherical coordinate ϕ) we handle the boundaries by linearly interpolating $\mathbf{v}[0]$ and $\mathbf{v}[M-1]$, as $\epsilon_1 = 0$ and 1 point to identical directions. For ϵ_2 , on the other hand, we use nearest neighbor to handle the boundaries. In both cases, the formulation in Eq. 9 remains valid and thus we use the same architecture, as shown in Fig. 3.

Sampling from the learned distribution is again performed using inverse CDF sampling, as shown in Fig. 2 (bottom-right). The main difference is that the CDF is piecewise quadratic (instead of linear). Refer to the supplementary materials for the full derivation.

Discussion Our representation has similarities to action space discretization in reinforcement learning, as proposed by Tang and Agrawal [TA20]. However, there are two key differences. First, this technique discretizes the PDF, so evaluation and sampling are done at discrete locations. In contrast, we estimate PDF at discrete locations, but evaluation and sampling remain continuous. The second, more subtle difference is that they assume different dimensions of the action space are independent, representing the joint PDF as a product of marginal distributions. In our case, the dimensions are correlated, so our second PDF is conditional (see Eq. 3).

3.1.1. Connection to NIS

Our method addresses the problem of learning a PDF and sampling from it in a fundamentally different way than NIS [MMR*19]. We focus on directly modeling a multidimensional PDF by estimating a series of marginal and conditional distributions; for example, in our approach, we estimate four 1D PDFs to model a four-dimensional PDF. Sampling is a byproduct of our system, performed through inverse CDF sampling. In contrast, NIS uses normalizing flows, where the primary process maps samples from a known (e.g., uniform) distribution to the target distribution through a set of coupling layers, with PDF evaluation being a byproduct of this sampling process. Specifically, NIS splits the input dimensions into two equally sized partitions, using one as input to the network to estimate the mapping function that warps the other dimensions. This process is repeated in the remaining coupling layers in an alternating pattern.

Our factorization is motivated by the product relationship between joint and 1D PDFs, whereas NIS splits the input dimensions due to its use of normalizing flows with coupling layers. This fundamental difference in approach leads to a significant performance gap, even in the more constrained path guiding scenario with only two input dimensions. Specifically, the key difference in two dimensions is that *evaluating* the PDF of a given sample in NIS requires evaluating the two networks sequentially, as the warped sample serves as the condition for the network in the next coupling layer. In contrast, we evaluate the PDF by executing the two networks in parallel. This formulation results in simpler gradients (avoiding backpropagation through two networks), lower training times, and faster inference, without sacrificing expressiveness (see supplementary material).

3.2. Optimization with Radiance Caching

Since our goal is to approximate the target distribution using our learned PDF, we optimize our networks by minimizing the KL divergence, defined as:

$$D(p(\omega) \parallel \hat{p}_\Theta(\omega)) = \int_{\Omega} p(\omega) \log \frac{p(\omega)}{\hat{p}_\Theta(\omega)} d\omega \quad (10)$$

where p and p_Θ are the target and guiding distributions, respectively. Note that the guiding distribution is obtained through the product of the two 1D PDFs in Eq. 3, evaluated using the two networks. Moreover, all the distributions are conditioned on both \mathbf{x} and ω_o , but we remove the condition for simplicity of notation.

Minimizing this loss necessitates taking its derivative with respect to the parameters of the guiding distribution (the two networks parameters $\Theta = \{\mathbf{w}_1, \mathbf{w}_2\}$) as follows:

$$\nabla_\Theta D(p(\omega) \parallel \hat{p}_\Theta(\omega)) = - \int_{\Omega} p(\omega) \nabla_\Theta \log \hat{p}_\Theta(\omega) d\omega, \quad (11)$$

where we omit the gradient with respect to $\log p(\omega)$, since it is independent of Θ . Calculating this derivative requires evaluating the integral which does not have a closed form solution. Therefore, similar to the existing methods [DWL23, MMR*19], we approximate it through MC integration as follows:

$$\langle \nabla_\Theta D(p(\omega) \parallel \hat{p}_\Theta(\omega)) \rangle = - \frac{1}{N} \sum_{i=1}^N \frac{p(\omega_i)}{q(\omega_i)} \nabla_\Theta \log \hat{p}_\Theta(\omega_i) \quad (12)$$

where the samples ω_i are drawn from $q(\omega_i)$, which could be BSDF importance sampling, guiding distribution, or a combination of the two. Moreover, the ideal target distribution is:

$$p(\omega_i) = \underbrace{\rho(\mathbf{x}, \omega_o, \omega_i) L_i(\mathbf{x}, \omega_i) |\cos(\theta_i)|}_{\text{integrand}} \underbrace{L_r(\mathbf{x}, \omega_o)^{-1}}_{\text{normalization factor}}, \quad (13)$$

where the left term is the integrand in Eq. 1 and the normalization factor is to ensure the PDF integrates to one.

Since the true incoming radiance $L_i(\mathbf{x}, \omega_i)$ is not available, existing methods [DWL23, MMR*19] estimate it using MC samples. Unfortunately, this increases the variance of the gradients, and consequently makes the optimization difficult, particularly in complex scenes. Additionally, the normalization term is the reflected radiance, a quantity that we are aiming to estimate. Current techniques ignore this factor with the hope that the scale is canceled when using optimization approaches, such as Adam [KB15], that work with the ratio of the current and historic gradients. This is true if the optimization is performed only at a fixed \mathbf{x} and ω_o , as the normalization factor will be the same for all the gradients. However, in practice, samples from a large number of surface locations with various outgoing directions are used to compute the gradients. Therefore, ignoring the normalization factor could potentially hurt the performance of the optimization.

To address the first issue and reduce the variance of the gradient, we propose to utilize radiance caching, a well-explored topic in rendering. In particular, we follow the technique by Müller et al. [MRNK21] (NRC) and cache the radiance using a small MLP network that takes the surface location and ray direction as the input and estimates the corresponding radiance. Similar to NRC, instead of caching the incoming radiance at the current location $L_i(\mathbf{x}, \omega_i)$, we learn the reflected radiance at the next intersection point $L_r(\mathbf{x}', \omega'_o)$. Despite the equivalence of these two quantities




Figure 4: Computing the target distribution requires obtaining $L_r(\mathbf{x}, \omega_o)$ and $L_i(\mathbf{x}, \omega_i)$ (see Eq. 13). We use a neural network that takes location and direction as the input and estimates the cached reflected radiance along that particular ray. By evaluating the network at the current (\mathbf{x}', ω'_o) and next (\mathbf{x}, ω_i) intersection points, we obtain an estimate of the reflected L_r and incoming L_i radiance. Note that here the incoming radiance at \mathbf{x}, ω_i is equal to the reflected radiance at \mathbf{x}', ω'_o .

(refer to Fig. 4), we observe superior performance in learning L_r over L_i due to differences in input coordinates. In our system, the network f_Φ takes the surface location at the next intersection point along with the corresponding outgoing direction and estimates the reflected radiance, i.e., $\hat{L}_r(\mathbf{x}', \omega'_o) = f_\Phi(\mathbf{x}', \omega'_o)$. Since $L_i(\mathbf{x}, \omega_i) = L_r(\mathbf{x}', \omega'_o)$, we can leverage this network to estimate the incoming radiance in Eq. 13.

Since the network estimates the reflected radiance, it can be used to estimate the normalization factor in Eq. 13, thereby addressing the second problem. To summarize, using our radiance caching network, the target distribution at ω_i is estimated as follows:

$$\hat{p}(\omega_i) = \rho(\mathbf{x}, \omega_o, \omega_i) f_\Phi(\mathbf{x}', \omega'_o) |\cos(\theta_i)| f_\Phi(\mathbf{x}, \omega_o)^{-1}. \quad (14)$$

We use this estimated target distribution $\hat{p}(\omega_i)$ in place of $p(\omega_i)$ in Eq. 12 to estimate the gradient of the KL divergence in training.

Note that theoretically, caching radiance and using it as L_i should not help with variance reduction; the variance will just be passed from here to radiance caching optimization. However, as shown in Fig. 9, we find this to be helpful in practice. This could be attributed to NRC's [MRNK21] highly effective approach to deal with the variance through the relative \mathcal{L}^2 loss, as well as the strategy to learn L_r as opposed to L_i .

Discussion We note an interesting connection between our method and actor-critic techniques in reinforcement learning [SB18, BSA83]. In reinforcement learning, the goal is to minimize the expected reward over all the trajectories. During training, this is often approximated using sampled trajectories making the gradients highly noisy. The idea behind actor-critic methods is to cache the reward (critic) to reduce the variance of gradient for optimizing the policy (actor). In this context, our radiance caching network f_Φ is the critic, while the distribution learning network f_Θ can be thought of as the actor.

A similar caching strategy has been applied to path guiding by Dahm and Keller [DK17] using Q -learning, which is a closely related, but different concept. Pantaleoni [Pan20] further explores caching by comparing different basis functions and data structures on both unbiased and biased alternatives. The main discerning characteristics of our approach are: 1) modeling the cache implicitly, which avoids spatial discretization artifacts, and 2) caching the out-

Table 1: Equal-time comparisons measured in relMSE (lower is better) for seven representative scenes. We compare our approach with nearest neighbor (DF-N) and linear (DF-L) interpolation against unidirectional path tracing (PT), PPG [MGN17], Variance [RGH*20], NIS [MMR*19], and NPM [DWL23]. We color code the first, second, and third lowest numbers.

	PT		PPG		Variance		NIS		NPM		DF-N		DF-L	
BATHROOM	1.1140	3774spp	0.5152	1598spp	0.5729	1445spp	0.2738	789spp	0.3761	1507spp	0.2601	1166spp	0.2106	1131spp
BEDROOM	0.0270	4759spp	0.0389	1634spp	0.0318	1591spp	0.0201	928spp	0.0182	1830spp	0.0167	1365spp	0.0138	1328spp
BREAKFAST	0.1668	6675spp	0.1578	1978spp	0.1080	1981spp	0.0263	1414spp	0.0286	3267spp	0.0199	2161spp	0.0171	2046spp
KITCHEN	0.0821	4697spp	0.0386	2644spp	0.0323	2240spp	0.0663	960spp	0.0196	2091spp	0.0186	1427spp	0.0146	1382spp
SALLE DE BAIN	0.0654	4969spp	0.1067	1462spp	0.0625	1469spp	0.0251	937spp	0.0233	2249spp	0.0193	1384spp	0.0155	1355spp
STAIRCASE	2.5989	5403spp	0.6820	2584spp	1.1779	2510spp	1.3275	1202spp	1.1155	1981spp	0.3494	1763spp	0.2861	1729spp
VEACH DOOR	14.182	6054spp	3.0940	1755spp	9.3482	1738spp	1.9415	1116spp	1.8092	2828spp	0.7132	1524spp	0.4129	1454spp

going radiance L_r with f_Φ to more effectively learn the product of incoming radiance L_i and BSDF with f_Θ .

4. Implementation

Online Optimization with Radiance Caching We optimize our approach in an online fashion using the prototype GPU-accelerated renderer of [DWL23]. For all experiments, we simultaneously train and guide paths for the first 30% of the allocated budget of time or spp; after which optimization is halted for the remainder of the task. To increase efficiency, we reuse rendering rays for training our radiance caching network at the end of each rendered frame. The training data is split into at least four mini batches of equal size to ensure an aggressive optimization of f_Φ before it is used to train f_{w_1} and f_{w_2} . We use Adam [KB15] with default parameters and a fixed learning rate of 10^{-2} for f_Φ . For the PDF estimation networks, we found that increasing the learning rate further to 3×10^{-2} is beneficial and stable thanks to our smoother cached objectives. We leverage the learned distributions to guide 70% of paths using our method, while the remainder of the paths is generated with unbiased BSDF sampling to allow exploration. We provide a memory and runtime analysis of our approach in the supplementary.

Networks Our networks, implemented in tiny-cuda-nn [M21], share the same underlying architecture of 3 hidden layers of 64 neurons and ReLU intermediary activation functions. We encode the inputs using [DWL23]’s learnable dense grid encoding for the intersection position, and spherical harmonics with degree 4 for normalized outgoing ray direction ω_o . Similar to existing methods, we additionally provide normal and roughness as the input to the networks and encode them through one-blob encoding [MMR*19] using 4 bins. Our conditional PDF estimation network f_{w_2} requires the extra input of ϵ_1 for which we use triangle wave encoding [MRNK21] with 12 frequencies. We set the discretization resolution M_1 to 32 and M_2 to 16. The reduction in resolution in the second dimension is due to the smaller angular range of ϵ_2 corresponding to θ (0 to π).

5. Results

In this section, we first describe the evaluation scenes, then show extensive visual and numerical comparisons to demonstrate the effectiveness of our approach. For all the results, we use a unidi-

rectional path tracer with BSDF sampling and disable next-event estimation (NEE), as done in prior methods [MMR*19, ZXS*21a, DWL23], to better illustrate the effectiveness of path guiding. Additionally, we disable Russian roulette and limit the path length to 6. We perform all comparisons and ablations on an AMD Ryzen 5800X CPU and an NVIDIA RTX3080 GPU. Here, we report quantitative results using the trimmed relative mean squared error [RKZ11] (relMSE) with a threshold of 0.1% averaged over 10 independent runs. Evaluation using additional metrics can be found in the supplementary.

5.1. Scenes

To better showcase complex lighting scenarios that are commonplace in real-world applications, seven representative scenes are used throughout this section. Here, we briefly describe each scene’s lighting setup and challenges. The BATHROOM and SALLE DE BAIN scenes show examples of bathrooms featuring cone-shaped and cube-shaped light fixtures, respectively. In order to reach a light source, paths require finer guidance from the learnable methods given the spacing between lights and their directional component. The semi-transparent curtains of the BEDROOM scene create a barrier to the light sources, resulting in indirect illumination for most of the paths, leaving only a small opening at the sides for direct lighting. The BREAKFAST scene features directional lighting through the blinds in addition to two high-intensity light fixtures that are focused on the table. When combined with the metallic BSDFs of the teapot and cups, it results in a small quantity of viable paths to light, increasing complexity. Similarly, the KITCHEN scene also features a directional light coming from the window in addition to a second light source on top of the microwave. The STAIRCASE is illuminated by three small high-intensity light fixtures placed at an angle that result in overlapping conic sections on the main atrium. Lastly, the VEACH DOOR scene showcases an example of indirect lighting, as the scene is lit through the reflections on the right wall. Three separate light sources sit behind the partially-open door.

5.2. Comparisons

Throughout this section, we compare our approach against several state-of-the-art path guiding methods. Specifically, we show comparisons against Müller et al. [MGN17] (PPG), Rath et

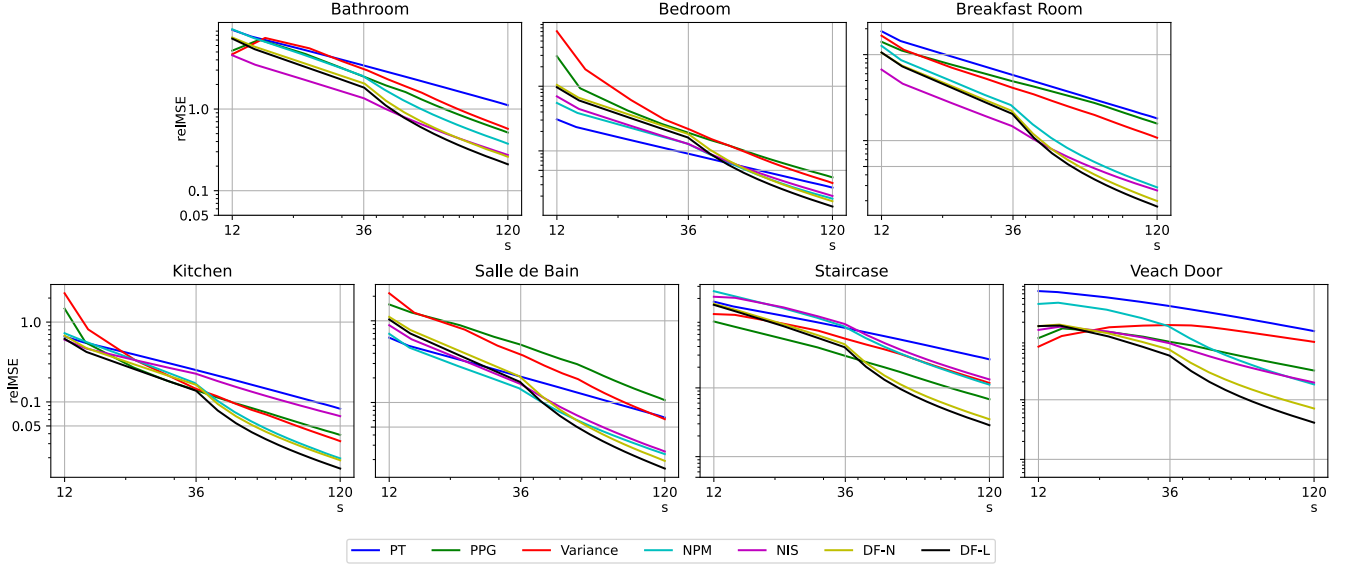


Figure 5: We show convergence plot of all the approaches on the seven scenes from 12 to 120 seconds (s). The 36 s line is where all the approaches (except the unidirectional path tracer) stop learning and use the learned distributions to sample the remaining paths.

al. [RGH*20] (Variance), Müller et al. [MMR*19] (NIS), and Dong et al. [DWL23] (NPM). We implement all methods on the fast prototype GPU path tracer provided by Dong et al. [DWL23]. We show the results of two variants of our distribution factorization approach with nearest (DF-N) and linear (DF-L) interpolation. We implement NIS with two fully-fused networks ($L = 2$) using piecewise quadratic coupling layers and fixed bin size and resolutions of 32×16 to match the capacity of our approach. Note that NIS, NPM and our method are product-based (learn the product of radiance and BSDF), but the other two methods are radiance-based.

Equal Time We begin by showing equal-time numerical and visual comparisons against the other approaches in Table 1 and Fig. 6 (see supplementary material for equal sample comparisons). To obtain these results, we provide a budget of 120 seconds to all the approaches, including the standard path tracer (PT). As shown, our method with linear interpolation produces the best results in all the scenes, with the nearest-neighbor variant being second-place. Faster methods such as PPG, Variance and NPM can render significantly more samples than our approach in some scenes, but fail to model PDFs accurately. In contrast, NIS traces fewer samples than our method in the same time.

We further demonstrate the encoding quality of our approach in an equal-time setting by comparing the learned PDFs of our method against NIS [MMR*19] and NPM [DWL23] on the VEACH DOOR scene in Fig. 7. NIS trains on fewer noisy MC samples, which results in a less detailed representation. NPM produces blurry estimations, leading to less samples reaching the light. Our approach with nearest-neighbor interpolation (DF-N) roughly approximates the target distribution, but fails to encode the light sources with detail due to a resolution limitation. In contrast, DF-L produces sharper approximations that correlate the most with the reference. Note that we do not show visualizations for Müller et al. [MGN17]

or Rath et al. [RGH*20] since these methods encode the radiance while our approach is designed for the product with BSDF.

Convergence Analysis To understand the convergence behavior of our approach, we plot the relMSE of the rendered images in equal-time (Fig. 5) and equal-sample (supplementary) settings. Overall, our approach produces better results than all the other methods in all cases. We note that, in some scenes (e.g. SALLE DE BAIN), other methods are able to learn a better map initially, but our results improve significantly once the training phase ends and we can trace more samples. Moreover, our method with linear interpolation is consistently better than nearest neighbor. Finally, the improvement becomes larger over time, as we learn better guiding distributions than the other approaches.

5.3. Ablations

Discretization Resolution Here, we evaluate the effect of resolution (M_1 and M_2) on the quality of results. We show our rendered results with 750 spp using three different resolutions in Fig. 8. Note that M_2 is always half of M_1 as it corresponds to θ which is between 0 and π . As seen, by increasing the resolution, our networks can model the target distribution more accurately, leading to better results. Particularly, increasing the resolution from 16×8 to 32×16 significantly improves the results, but the improvement beyond this resolution is smaller. Given the computational cost associated with higher resolutions, we choose 32×16 to generate all the results.

Optimization with Radiance Caching We further evaluate the impact of caching radiance during optimization in Fig. 9 for the VEACH DOOR scene (and for all scenes in the supplementary). The impact of cache varies by scene, with higher gains in scenes with complex light transport. For the VEACH DOOR, using the cached radiance to estimate both L_r and L_i in the target distribution (Eq. 14) significantly improves the results.

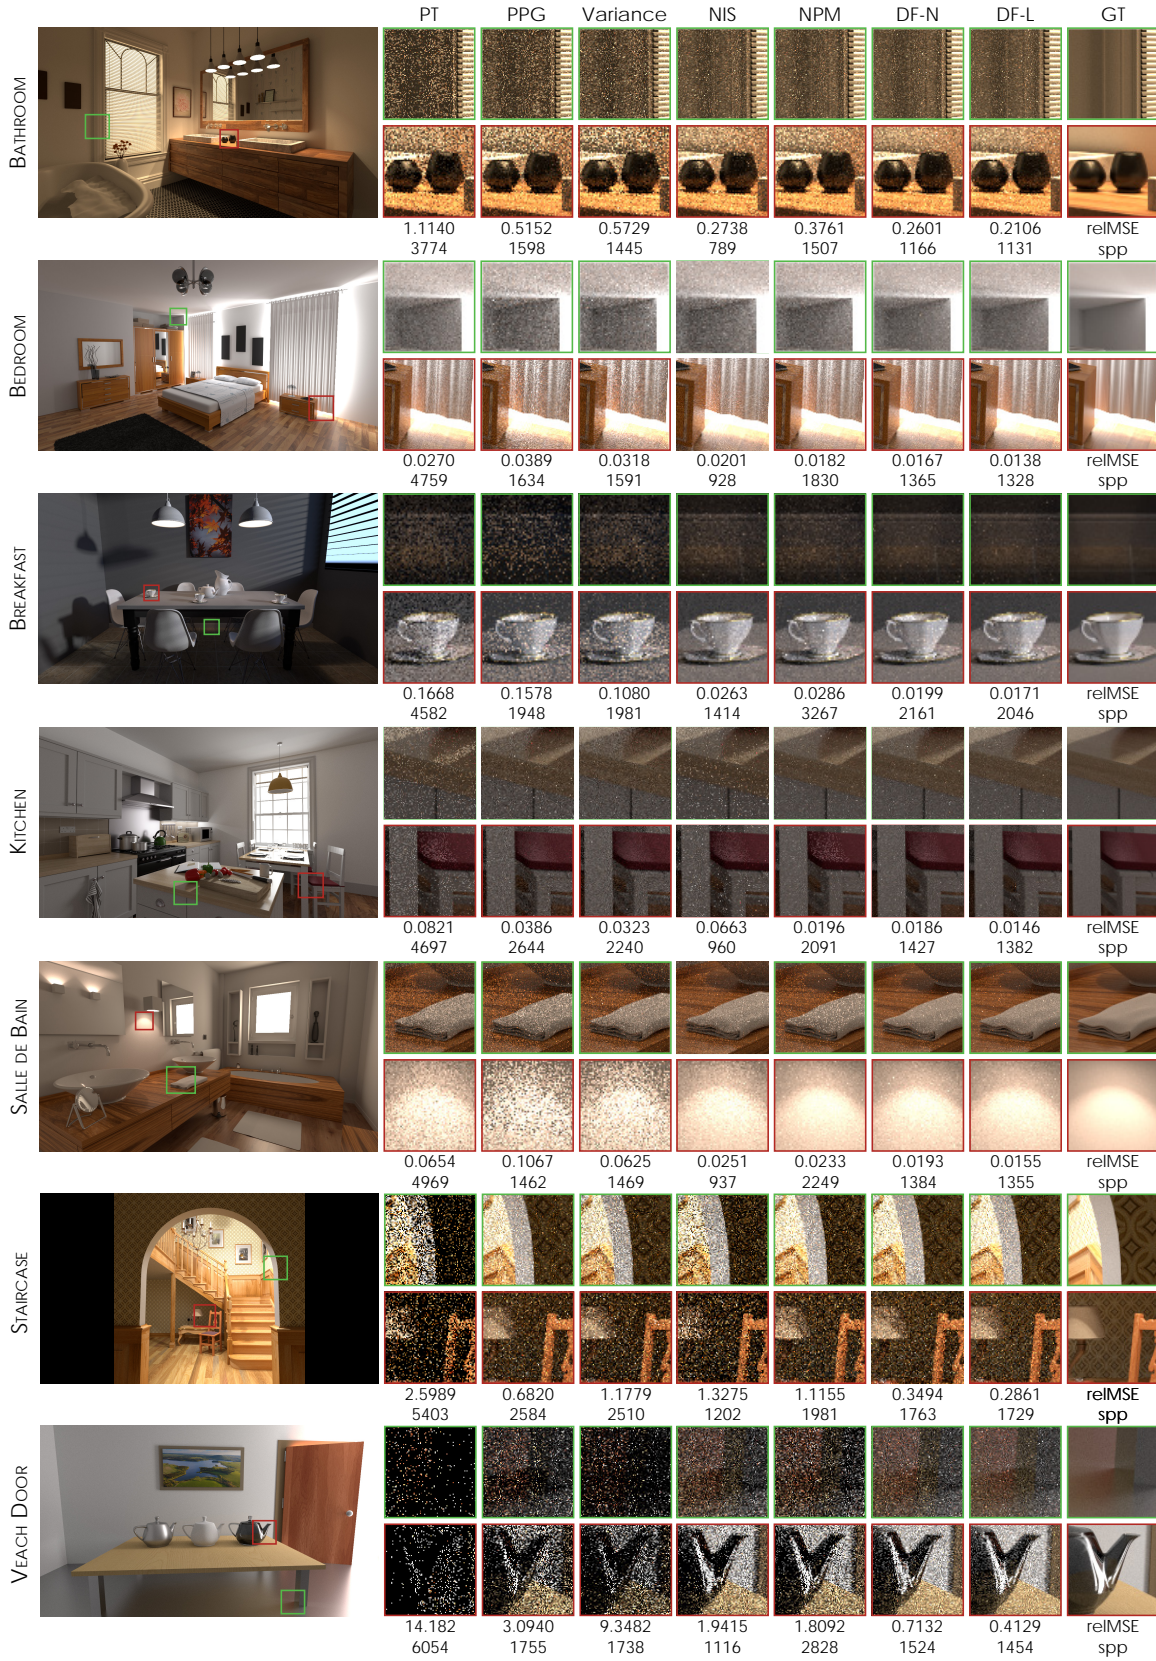


Figure 6: Equal-time comparison against Müller et al. [MGN17] (PPG), Rath et al. [RGH*20] (Variance), Müller et al. [MMR*19] (NIS), Dong et al. [DWL23] (NPM), and our method with nearest neighbor (DF-N) and linear (DF-L) interpolation. The time budget for all the methods is 120 seconds, with a training budget of 30%.




Figure 7: Equal-time comparison of learned distributions for the VEACH DOOR scene. The time budget for all methods is 120s.

6. Conclusion, Limitations, and Future Work

In this paper, we have presented a novel neural path guiding method. We first factorize the distribution over the 2D directional domain into two 1D PDFs and estimate them using two separate neural networks at discrete locations. The PDF at an arbitrary coordinate is then obtained through interpolation. We explore two interpolation strategies and discuss evaluation and sampling process for each. Moreover, we propose a strategy to reduce the variance and normalize the samples used as the target distribution during optimization. Specifically, we demonstrate that we can cache the radiance using an additional network and use the estimates for incoming radiance and normalization. We demonstrate the effectiveness of our approach with extensive experiments.

While our approach works well in most tested scenarios, it uses a fixed resolution, and thus struggles to encode features that are significantly smaller than each bin. We visualize this limitation on an equal-time comparison of the outdoor SWIMMING POOL scene in Fig. 10. As shown, our method does not properly model the Sun, producing results with excessive noise. This problem is directly related to the resolution since increasing it from 32×16 to 64×32 improves the results. Similarly, NIS [MMR*19] (with fixed bin size) suffers from the same limitation, but produces slightly worse results because of higher computational cost. In contrast, PPG [MGN17] and Variance [RGH*20] quickly adapt their data structures to focus on the directional light of the Sun, resulting in the least noise. Although NPM [DWL23] uses a continuous representation, surprisingly, it struggles the most here, as its optimization becomes unstable with the concentrated high intensity illumination. In the future, it would be interesting to combine our approach with explicit spatial data structures to vary the resolutions based on the complexity. Alternatively, we could adopt the variable bin size strategy, as




Figure 8: Effect of discretization resolution $M_1 \times M_2$ on the quality of the results. As the resolution increases, the quality improves, but this comes at the cost of increased computational complexity.




Figure 9: Effect of radiance caching during equal-time comparison of 120s on the VEACH DOOR. We compare the results of our approach ($L_r + L_i$) against not using the cache (similar to existing methods) as well as using the cache only for L_i in Eq. 14. Using the cache for both L_i and L_r improves the results.

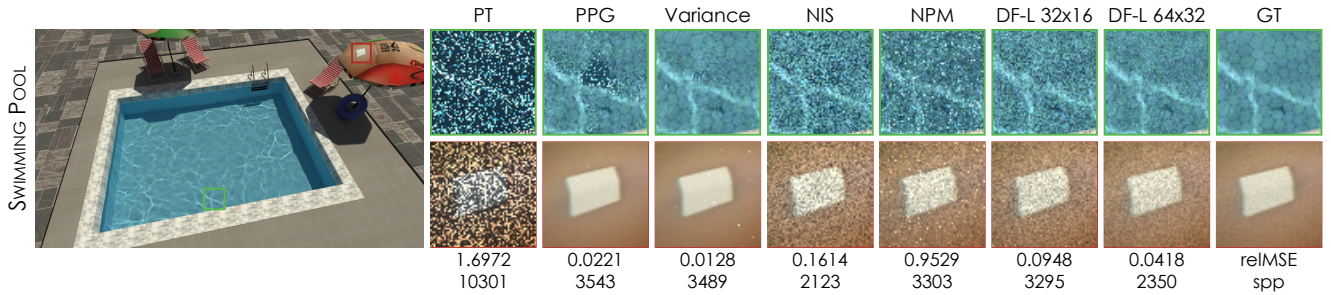


Figure 10: Equal-time comparison (120s) on the SWIMMING POOL scene. The pool shows direct lighting from an environment map that models the Sun forming caustics in an outdoor swimming pool. Our approach struggles in this case as it does not sharply encode the directional light from the Sun. This is because of our fixed discretization resolution since, as shown, results improve as we increase the resolution from 32×16 to 64×32 .




Figure 11: Equal-time comparison (30 seconds budget) on the CORNELL BOX with flipped light. Although DF-N and DF-L encode higher quality PDFs, Dong et al. [DWL23] (NPM) generate the result with least noise because they trace more samples.

proposed by Müller et al. [MMR*19], to increase the flexibility of fixed resolution.

In addition, we designed our approach to be highly effective specifically when handling complex scenes. On simple scenes with easy to model light transport, the advantage of our method is not significant. For such cases, when using a fast GPU path tracer, simpler and faster methods might produce better results because of their ability to trace more samples. Fig. 11 demonstrates this phenomenon through an equal-time experiment of 30 seconds using the CORNELL BOX with flipped light. Although the learned PDF of the two variants of our approach better correlate with the target distribution compared to Dong et al. [DWL23] (NPM), NPM's ability to throw more samples results in less noise overall.

Acknowledgments

This project was funded in part by the NSF CAREER Award #2238193. We are grateful to Dong et al. [DWL23] for releasing the source code of their work. We would like to thank the following artists for sharing their scenes and models that appear in our figures: Mareck (BATHROOM), Slyk-Drako (BEDROOM), Wig42 (BREAKFAST, STAIRCASE), Jay-Artist (KITCHEN), nacimus (SALLE DE BAIN), Benedikt Bitterli (VEACH DOOR, CORNELL BOX) and CN Entertainment, LordSamueliSolo (SWIMMING POOL).

References

- [BAJ08] BUDGE B. C., ANDERSON J. C., JOY K. I.: Caustic forecasting: Unbiased estimation of caustic lighting for global illumination. *Computer Graphics Forum* 27, 7 (2008), 1963–1970. 2
- [BB17] BUS N., BOUBEKEUR T.: Double hierarchies for directional importance sampling in monte carlo rendering. *Journal of Computer Graphics Techniques (JCGT)* 6, 3 (August 2017), 25–37. 1, 2
- [BFK18] BINDER N., FRICKE S., KELLER A.: Fast path space filtering by jittered spatial hashing. In *ACM SIGGRAPH 2018 Talks* (New York, NY, USA, 2018), SIGGRAPH '18, Association for Computing Machinery. 2
- [BMDS19] BAKO S., MEYER M., DEROSE T., SEN P.: Offline deep importance sampling for monte carlo path tracing. *Computer Graphics Forum* 38, 7 (2019), 527–542. 2
- [BRDC12] BASHFORD-ROGERS T., DEBATTISTA K., CHALMERS A.: A significance cache for accelerating global illumination. *Computer Graphics Forum* 31, 6 (2012), 1837–1851. 2
- [BSA83] BARTO A. G., SUTTON R. S., ANDERSON C. W.: Neuron-like adaptive elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, and Cybernetics SMC-13*, 5 (1983), 834–846. 5
- [DBN17] DUBOUCHET R. A., BELCOUR L., NOWROUZEZAHRAI D.: Frequency based radiance cache for rendering animations. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations* (Goslar, DEU, 2017), EGSR '17, Eurographics Association, p. 41–53. 2
- [DK17] DAHM K., KELLER A.: Learning light transport the reinforced way. In *ACM SIGGRAPH 2017 Talks* (New York, NY, USA, 2017), SIGGRAPH '17, Association for Computing Machinery. 2, 5
- [DPÖM22] DODIK A., PAPAS M., ÖZTIRELI C., MÜLLER T.: Path guiding using spatio-directional mixture models. *Computer Graphics Forum* 41, 1 (2022), 172–189. 1, 2
- [DSB17] DINH L., SOHL-DICKSTEIN J., BENGIO S.: Density estimation using real NVP. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings* (2017), OpenReview.net. 1, 2
- [DWL23] DONG H., WANG G., LI S.: Neural parametric mixtures for path guiding. In *ACM SIGGRAPH 2023 Conference Proceedings* (New York, NY, USA, 2023), SIGGRAPH '23, Association for Computing Machinery. 1, 2, 5, 6, 7, 8, 9, 10
- [FHG*23] FAN Z., HONG P., GUO J., ZOU C., GUO Y., YAN L.-Q.: Manifold path guiding for importance sampling specular chains. *ACM Trans. Graph.* 42, 6 (dec 2023). 2
- [GBBE18] GUO J. J., BAUSZAT P., BIKKER J., EISEMANN E.: Primary sample space path guiding. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations* (Goslar, DEU, 2018), SR '18, Eurographics Association, p. 73–82. 1, 2

- [GKB09] GASSENBAUER V., KŘIVÁNEK J., BOUATOUCH K.: Spatial directional radiance caching. *Computer Graphics Forum* 28, 4 (2009), 1189–1198. 2
- [HEV*16] HERHOLZ S., ELEK O., VORBA J., LENSCH H., KŘIVÁNEK J.: Product importance sampling for light transport path guiding. *Computer Graphics Forum* 35, 4 (2016), 67–77. 1, 2
- [HIT*24] HUANG J., IIZUKA A., TANAKA H., KOMURA T., KITAMURA Y.: Online neural path guiding with normalized anisotropic spherical gaussians. *ACM Trans. Graph.* 43, 3 (Apr. 2024). 1, 2
- [HP02] HEY H., PURGATHOFER W.: Importance sampling with hemispherical particle footprints. In *Proceedings of the 18th Spring Conference on Computer Graphics* (New York, NY, USA, 2002), SCCG '02, Association for Computing Machinery, p. 107–114. 2
- [HWZ*20] HUO Y., WANG R., ZHENG R., XU H., BAO H., YOON S.-E.: Adaptive incident radiance field sampling and reconstruction using deep reinforcement learning. *ACM Trans. Graph.* 39, 1 (jan 2020). 2
- [HZE*19] HERHOLZ S., ZHAO Y., ELEK O., NOWROUZEZAHRAI D., LENSCH H. P. A., KŘIVÁNEK J.: Volume path guiding based on zero-variance random walk theory. *ACM Trans. Graph.* 38, 3 (jun 2019). 2
- [Jen95] JENSEN H. W.: Importance driven path tracing using the photon map. In *Rendering Techniques '95* (Vienna, 1995), Hanrahan P. M., Purgathofer W., (Eds.), Springer Vienna, pp. 326–335. 2
- [Kaj86] KAJIYA J. T.: The rendering equation. 143–150. 2
- [KB15] KINGMA D., BA J.: Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)* (2015). 5, 6
- [LHL*24] LITALIEN J., HAŠAN M., LUAN F., MULLIA K., GEORGIEV I.: Neural product importance sampling via warp composition. In *SIGGRAPH Asia 2024 Conference Papers* (New York, NY, USA, 2024), SA '24, Association for Computing Machinery. 2
- [LW95] LAFORTUNE E. P., WILLEMS Y. D.: A 5d tree to reduce the variance of monte carlo ray tracing. In *Rendering Techniques '95* (Vienna, 1995), Hanrahan P. M., Purgathofer W., (Eds.), Springer Vienna, pp. 11–20. 2
- [LWT*22] LI H., WANG B., TU C., XU K., HOLZSCHUCH N., YAN L.-Q.: Unbiased caustics rendering guided by representative specular paths. In *SIGGRAPH Asia 2022 Conference Papers* (New York, NY, USA, 2022), SA '22, Association for Computing Machinery. 2
- [M21] MÜLLER T.: tiny-cuda-nn, 4 2021. URL: <https://github.com/NVlabs/tiny-cuda-nn>. 6
- [MGN17] MÜLLER T., GROSS M., NOVÁK J.: Practical path guiding for efficient light-transport simulation. *Computer Graphics Forum* 36, 4 (2017), 91–100. 1, 2, 6, 7, 8, 9
- [MJG18] MARCO J., JARABO A., JAROSZ W., GUTIERREZ D.: Second-order occlusion-aware volumetric radiance caching. *ACM Trans. Graph.* 37, 2 (jul 2018). 2
- [MMR*19] MÜLLER T., MCWILLIAMS B., ROUSSELLE F., GROSS M., NOVÁK J.: Neural importance sampling. *ACM Trans. Graph.* 38, 5 (oct 2019). 1, 2, 4, 5, 6, 7, 8, 9, 10
- [MRNK21] MÜLLER T., ROUSSELLE F., NOVÁK J., KELLER A.: Real-time neural radiance caching for path tracing. *ACM Trans. Graph.* 40, 4 (jul 2021). 2, 5, 6
- [Pan20] PANTALEONI J.: Online path sampling control with progressive spatio-temporal filtering, 2020. [arXiv:2005.07547](https://arxiv.org/abs/2005.07547). 2, 5
- [PWP08] PEGORARO V., WALD I., PARKER S. G.: Sequential monte carlo adaptation in low-anisotropy participating media. *Computer Graphics Forum* 27, 4 (2008), 1097–1104. 2
- [RGH*20] RATH A., GRITTMANN P., HERHOLZ S., VÉVODA P., SLUSALLEK P., KŘIVÁNEK J.: Variance-aware path guiding. *ACM Trans. Graph.* 39, 4 (aug 2020). 2, 6, 7, 8, 9
- [RHL20] RUPPERT L., HERHOLZ S., LENSCH H. P. A.: Robust fitting of parallax-aware mixtures for path guiding. *ACM Trans. Graph.* 39, 4 (aug 2020). 1, 2
- [RKZ11] ROUSSELLE F., KNAUS C., ZWICKER M.: Adaptive sampling and reconstruction using greedy error minimization. *ACM Transactions on Graphics (TOG)* 30, 6 (2011), 1–12. 6
- [RZD14] REHFELD H., ZIRR T., DACHSBACHER C.: Clustered pre-convolved radiance caching. In *Proceedings of the 14th Eurographics Symposium on Parallel Graphics and Visualization* (Goslar, DEU, 2014), PGV '14, Eurographics Association, p. 25–32. 2
- [SB18] SUTTON R. S., BARTO A. G.: *Reinforcement Learning: An Introduction*, second ed. The MIT Press, 2018. 5
- [SL06] STEINHURST J., LASTRA A.: Global importance sampling of glossy surfaces using the photon map. *2006 IEEE Symposium on Interactive Ray Tracing* (2006), 133–138. 2
- [SL17] SILVENNOINEN A., LEHTINEN J.: Real-time global illumination by precomputed local reconstruction from sparse radiance probes. *ACM Trans. Graph.* 36, 6 (nov 2017). 2
- [SNRS12] SCHERZER D., NGUYEN C. H., RITSCHER T., SEIDEL H.-P.: Pre-convolved radiance caching. *Comput. Graph. Forum* 31, 4 (jun 2012), 1391–1397. 2
- [TA20] TANG Y., AGRAWAL S.: Discretizing continuous action space for on-policy optimization. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 04 (Apr. 2020), 5981–5988. 4
- [VHH*19] VORBA J., HANIKA J., HERHOLZ S., MÜLLER T., KŘIVÁNEK J., KELLER A.: Path guiding in production. In *ACM SIGGRAPH 2019 Courses* (New York, NY, USA, 2019), SIGGRAPH '19, ACM, pp. 18:1–18:77. 2
- [VK16] VORBA J., KŘIVÁNEK J.: Adjoint-driven russian roulette and splitting in light transport simulation. *ACM Trans. Graph.* 35, 4 (jul 2016). 2
- [VKv*14] VORBA J., KARLÍK O., ŠIK M., RITSCHER T., KŘIVÁNEK J.: On-line learning of parametric mixture models for light transport simulation. *ACM Trans. Graph.* 33, 4 (jul 2014). 1, 2
- [VPG14] VARDIS K., PAPAIOANNOU G., GKARAVELIS A.: Real-time radiance caching using chrominance compression. *Journal of Computer Graphics Techniques (JCGT)* 3, 4 (December 2014), 111–131. 2
- [WRC88] WARD G. J., RUBINSTEIN F. M., CLEAR R. D.: A ray tracing solution for diffuse interreflection. In *Proceedings of the 15th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1988), SIGGRAPH '88, Association for Computing Machinery, p. 85–92. 2
- [XWH*23] XU B., WU L., HASAN M., LUAN F., GEORGIEV I., XU Z., RAMAMOORTHY R.: Neusample: Importance sampling for neural materials. In *ACM SIGGRAPH 2023 Conference Proceedings* (2023), pp. 1–10. 2
- [ZBN19] ZHAO Y., BELCOUR L., NOWROUZEZAHRAI D.: View-dependent radiance caching. In *Proceedings of the 45th Graphics Interface Conference on Proceedings of Graphics Interface 2019* (Waterloo, CAN, 2019), GI'19, Canadian Human-Computer Communications Society. 2
- [ZBX*21] ZHU J., BAI Y., XU Z., BAKO S., VELÁZQUEZ-ARMENDÁRIZ E., WANG L., SEN P., HASAN M., YAN L.-Q.: Neural complex luminaires: representation and rendering. *ACM Trans. Graph.* 40, 4 (2021), 57–1. 2
- [ZXS*21a] ZHU S., XU Z., SUN T., KUZNETSOV A., MEYER M., JENSEN H. W., SU H., RAMAMOORTHY R.: Hierarchical neural reconstruction for path guiding using hybrid path and photon samples. *ACM Trans. Graph.* 40, 4 (jul 2021). 2, 6
- [ZXS*21b] ZHU S., XU Z., SUN T., KUZNETSOV A., MEYER M., JENSEN H. W., SU H., RAMAMOORTHY R.: Photon-driven neural reconstruction for path guiding. *ACM Trans. Graph.* 41, 1 (Nov. 2021). 2
- [ZZ19] ZHENG Q., ZWICKER M.: Learning to importance sample in primary sample space. *Computer Graphics Forum* 38, 2 (2019), 169–179. 2