

Neural Importance Sampling of Many Lights

PEDRO FIGUEIREDO, Texas A&M University, USA

QIHAO HE, Texas A&M University, USA

STEVE BAKO, Aurora Innovation, USA

NIMA KHADEMI KALANTARI, Texas A&M University, USA

We propose a neural approach for estimating spatially varying light selection distributions to improve importance sampling in Monte Carlo rendering, particularly for complex scenes with many light sources. Our method uses a neural network to predict the light selection distribution at each shading point based on local information, trained by minimizing the KL-divergence between the learned and target distributions in an online manner. To efficiently manage hundreds or thousands of lights, we integrate our neural approach with light hierarchy techniques, where the network predicts cluster-level distributions and existing methods sample lights within clusters. Additionally, we introduce a residual learning strategy that leverages initial distributions from existing techniques, accelerating convergence during training. Our method achieves superior performance across diverse and challenging scenes.

CCS Concepts: • **Computing methodologies** → **Ray tracing; Neural networks**; • **Mathematics of computing** → *Sequential Monte Carlo methods*.

Additional Key Words and Phrases: many lights rendering, ray tracing, importance sampling, neural networks, next event estimation

ACM Reference Format:

Pedro Figueiredo, Qihao He, Steve Bako, and Nima Khademi Kalantari. 2025. Neural Importance Sampling of Many Lights. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25)*, August 10–14, 2025, Vancouver, BC, Canada. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3721238.3730754>

1 Introduction

Monte Carlo (MC) rendering is the gold standard for offline rendering and is widely adopted in the film industry, where photorealistic results are critical. With this widespread adoption, the demand for rendering increasingly complex scenes has grown, particularly those with many light sources. Even computing direct illumination in such scenarios presents significant challenges, as efficient rendering requires identifying the light sources that contribute the most at each shading point by accounting for the light intensity, geometry, visibility, and material properties (see Fig. 1).

A group of techniques [Conty Estevez and Kulla 2018; Fernandez et al. 2002; Liu et al. 2019; Moreau et al. 2022; Pantaleoni 2019; Paquette et al. 1998; Tokuyoshi et al. 2024; Vévoda et al. 2018; Walter et al. 2005; Wang et al. 2021; Yuksel 2019] address this problem by

constructing a light hierarchy (e.g., using a binary tree) where lights are grouped based on criteria such as spatial position, intensity, or other characteristics. These methods then adaptively determine a cut through the hierarchy and define the importance of the clusters along the cut. Clusters are sampled with probabilities proportional to their importance, with either a representative light from the selected cluster being evaluated [Walter et al. 2005] or subtrees being traversed to evaluate individual lights [Conty Estevez and Kulla 2018; Yuksel 2019]. A key limitation of these approaches is that importance is computed without considering visibility, which can result in frequently sampling occluded lights (see Fig. 1).

To address this problem, several methods [Donikian et al. 2006; Fernandez et al. 2002; Pantaleoni 2019; Vévoda et al. 2018; Wang et al. 2021] propose progressively improving the sampling distributions during rendering by incorporating visibility. All these techniques utilize complex spatial data structures to track cluster importance during the online learning process and they mainly differ in their optimization strategies. Despite their ability to account for visibility, these methods optimize a single distribution for groups of shading points in each subspace, ignoring significant variations within the subareas (see Fig. 1).

In this paper, we address these limitations by estimating a spatially varying light selection distribution using a neural network. Unlike prior approaches that rely on complex spatial data structures to track and optimize cluster importance, our network directly outputs the distribution for each shading point based on its local information (e.g., position and outgoing direction). This eliminates the need for maintaining and querying such data structures, while enabling the estimation of spatially varying distributions across all shading points. To train the network, we minimize the KL-divergence between the learned and target distributions, where the target distribution is approximated using Monte Carlo (MC) sampling.

To efficiently handle scenes with a large number of light sources (hundreds or thousands), we propose combining our neural approach with existing light hierarchy techniques. Specifically, we construct a light hierarchy and define our cut as all the nodes at a certain height of the tree. The neural network predicts the distribution over these clusters, while existing methods are used to sample lights within each cluster. Furthermore, we enhance convergence by utilizing the cluster distribution from existing techniques as an initialization and training the network to estimate the residual.

Although inspired by neural path guiding methods [Dong et al. 2023; Huang et al. 2024; Müller et al. 2019] that focus on continuous selection of the ray directions, our formulation is specifically designed for discrete light selection. This distinction is critical for handling scenes with a large number of light sources, where finding many important, but sparsely distributed lights, in the directional domain may be challenging.

Authors' Contact Information: Pedro Figueiredo, Texas A&M University, USA, pedrofigueiredo@tamu.edu; Qihao He, Texas A&M University, USA, phyqh@tamu.edu; Steve Bako, Aurora Innovation, USA, sbako@aurora.tech; Nima Khademi Kalantari, Texas A&M University, USA, nimak@tamu.edu.

SIGGRAPH Conference Papers '25, Vancouver, BC, Canada

© 2025 Copyright held by the owner/author(s).

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25)*, August 10–14, 2025, Vancouver, BC, Canada, <https://doi.org/10.1145/3721238.3730754>.




Fig. 1. We illustrate a scene with three light sources of increasing brightness. Accurate light sampling requires considering factors such as light intensity, the BRDF, and visibility. The light with the highest contribution at each shading point, accounting for all factors, is indicated by red arrows. Sampling based solely on illumination magnitude leads to over-sampling light 2 at shading point x_1 , even though it does not contribute due to the directional constraints of the BRDF. Similarly, methods that consider both light intensity and the BRDF [Conty Estevez and Kulla 2018; Yuksel 2019] heavily sample light 1, which is occluded and thus does not contribute to the reflected light. Properly incorporating all three factors reveals that light 0 should be sampled more frequently at x_1 . While learning-based techniques [Pantaleoni 2019; Vévoda et al. 2018; Wang et al. 2021] incorporate visibility into their framework, they compute a single distribution for clusters of shading points (indicated by the gray dashed box), failing to capture critical local variations. For instance, light 0 dominates at x_1 , whereas light 2 contributes most at the neighboring point x_2 , underscoring the importance of spatially varying distributions.

We demonstrate that our method outperforms existing approaches on diverse and challenging scenes.

In summary, our contributions are as follows:

- A neural approach for learning light selection distributions for effective importance sampling.
- A hybrid method that integrates the neural approach with light hierarchy techniques to efficiently manage scenes with large numbers of light sources.
- A residual learning strategy to reduce the startup cost and improve convergence during training.

2 Related Work

Many Light Rendering. Early approaches [Kok and Jansen 1994; Shirley et al. 1996; Ward 1994] address scenes with many lights by distinguishing between important and unimportant light sources, thereby reducing the number of shadow rays traced. Several methods propose clustering lights by sampling the entries of the light transport matrix [Hašan et al. 2007; Hašan et al. 2008; Huo et al. 2015; Ou and Pellacini 2011], assuming that it is of low-rank.

A more widely adopted approach to tackle this problem involves building a light hierarchy to form clusters [Fernandez et al. 2002; Paquette et al. 1998]. Specifically, Walter et al. [2005] define clusters by finding a cut through the light tree and sampling representative lights at each cluster. Subsequent works [Conty Estevez and Kulla 2018; Liu et al. 2019; Moreau et al. 2022; Tokuyoshi et al. 2024; Yuksel 2019] make the sampling process unbiased by traversing the tree

stochastically and sampling individual light sources. However, a significant limitation of these techniques is their failure to account for visibility during sampling.

Learning-Based Light Sampling. A large number of recent methods incorporate visibility into light sampling in a data-driven manner, either through preprocessing [Georgiev et al. 2012; Wu and Chuang 2013] or online updates [Donikian et al. 2006; Fernandez et al. 2002; Pantaleoni 2019; Vévoda et al. 2018; Wang et al. 2021]. These approaches typically rely on spatial data structures to track distributions, learning a single distribution for a group of shading points within a subregion. However, this approach fails to capture fine-grained variations in the distribution. We address this limitation by learning spatially varying distributions using a neural network.

Resampled Importance Sampling. Resampled importance sampling (RIS) [Talbot et al. 2005] provides a framework for rendering methods to importance sample from complex distributions. It achieves this by first drawing samples from a simple proposal distribution, then selecting a subset of these samples based on weights derived from a complex, unnormalized distribution that accounts for factors such as material, geometry, and visibility at each shading point. Since its introduction, resampling techniques have been extended to real-time rendering, most notably in spatiotemporal reservoir resampling (ReSTIR) [Bitterli et al. 2020] and its variants for global illumination, including ReSTIR-GI [Ouyang et al. 2021] and ReSTIR-PT [Lin et al. 2022]. These methods demonstrate exceptional real-time performance in complex, many-light scenarios. Many-light sampling approaches, including ours, are complementary to these techniques, as they can be used to improve the candidate distributions for resampled importance sampling frameworks like RIS and ReSTIR, ultimately enhancing their results.

Path Guiding. Path guiding techniques perform importance sampling over the directional domain. Some methods rely on data structures, such as binary and quad trees [Müller et al. 2017; Rath et al. 2020; Vorba et al. 2019; Zhu et al. 2021b], to store fitted distributions. More recent approaches utilize neural networks to fit target distributions, either offline [Bako et al. 2019; Huo et al. 2020; Zhu et al. 2021b,c] or online [Dong et al. 2023; Huang et al. 2024; Lu et al. 2024; Müller et al. 2019]. However, fitting accurate distributions for scenes with many sparsely distributed lights remains challenging, limiting the effectiveness of these techniques in such scenarios. Nonetheless, path guiding methods are complementary to light sampling approaches, as their combination with multiple importance sampling (MIS) is often required to handle diverse scenes effectively. While our approach draws inspiration from neural path guiding, we diverge by focusing on a novel formulation for the discrete domain, rather than learning a continuous distribution over the 2D directional domain.

3 Neural Light Selection Importance Sampling

We aim to estimate the reflected light, $L_o(\mathbf{x}, \omega_o)$, in direction ω_o at point \mathbf{x} , due to direct lighting. This is defined as the sum of contributions from all the light sources:

$$L_o(\mathbf{x}, \omega_o) = \sum_{y=1}^M L_y(\mathbf{x}, \omega_o) = \sum_{y=1}^M \int_{A_y} F(\mathbf{x}, \mathbf{z}, \omega_o) dA_y(\mathbf{z}) \quad (1)$$

where M is the total number of light sources. Moreover, the contribution of y^{th} light source is obtained by integrating F over its area A_y . The integrand, F , models the interaction of light emitted from \mathbf{z} , a point on the light source, with the surface and is defined as:

$$F(\mathbf{x}, \mathbf{z}, \omega_o) = L_i(\mathbf{x}, \mathbf{z}) f_s(\mathbf{x}, \omega_o, \mathbf{z}) G(\mathbf{x}, \mathbf{z}), \quad (2)$$

where $L_i(\mathbf{x}, \mathbf{z})$ represents the radiance emitted from \mathbf{z} toward \mathbf{x} , $f_s(\mathbf{x}, \omega_o, \mathbf{z})$ refers to the bidirectional reflectance distribution function (BRDF), and $G(\mathbf{x}, \mathbf{z})$ accounts for geometry and visibility terms.

To facilitate efficient computation of the summation in Eq. 1, we express it as an expectation over $Y \sim p(y)$ as follows:

$$L_o(\mathbf{x}, \omega_o) = \mathbb{E} \left[\frac{L_Y(\mathbf{x}, \omega_o)}{p(Y)} \right] \quad (3)$$

where $p(y)$ is the probability mass function (PMF) and defines the probability of selecting the y^{th} light at point \mathbf{x} and for outgoing direction ω_o . For brevity, we omit the explicit conditioning on \mathbf{x} and ω_o in the notation of the distribution, though these dependencies remain implied. Note that this equality is valid as long as $p(y)$ is nonzero wherever the light contribution $L_y(\mathbf{x}, \omega_o)$ is nonzero.

This expectation can be approximated using N independent samples of light indices, $Y_j \in \{1, \dots, M\}$ and $Y_j \sim p(y)$, with the following Monte Carlo (MC) estimator:

$$\langle L_o(\mathbf{x}, \omega_o) \rangle = \frac{1}{N} \sum_{j=1}^N \frac{L_{Y_j}(\mathbf{x}, \omega_o)}{p(Y_j)}. \quad (4)$$

The choice of $p(y)$ has a substantial impact on the variance of this estimator. To reduce the variance, the light selection PMF, $p(y)$, should closely match the distribution of the numerator. To achieve this, we propose to model $p(y)$ using a neural network, $p_\theta(y)$, parameterized by θ , and train it in an online manner.

Note that computing $L_y(\mathbf{x}, \omega_o)$ requires evaluating the integral in Eq. 1 and is also performed through MC integration by randomly sampling a point \mathbf{z} on the light source. As in prior work on light sampling [Conty Estevez and Kulla 2018; Wang et al. 2021; Yuksel 2019], we use standard techniques [Hart et al. 2020; Shirley et al. 1996] to perform this.

In the subsequent sections, we detail our optimization strategy for training $p_\theta(y)$ and describe our efficient network architecture.

3.1 Optimizing the Light Selection Network

Our goal is to train our light selection network so that our learned PMF $p_\theta(y)$ closely matches the target distribution $q(y)$, the details of which will be discussed later. To this end, we propose to minimize the KL-divergence between the target and learned PMFs as follows:

$$D_{\text{KL}}(q, p_\theta) = \sum_{y=1}^M q(y) \log \frac{q(y)}{p_\theta(y)}. \quad (5)$$

To minimize this loss, we need to take its derivative with respect to the network parameters θ as follows:

$$\nabla_\theta D_{\text{KL}}(q, p_\theta) = - \sum_{y=1}^M q(y) \nabla_\theta \log p_\theta(y), \quad (6)$$

where the gradient of $q(y)$ is omitted as the target PMF is independent of θ . Expressing this gradient as an expectation over $Y \sim p_\theta(y)$:

$$\nabla_\theta D_{\text{KL}}(q, p_\theta) = - \mathbb{E} \left[\frac{q(Y)}{p_\theta(Y)} \nabla_\theta \log p_\theta(Y) \right], \quad (7)$$

allows us to effectively approximate it using an MC estimator:

$$\langle \nabla_\theta D_{\text{KL}}(q, p_\theta) \rangle = - \frac{1}{N} \sum_{j=1}^N \left[\frac{q(Y_j)}{p_\theta(Y_j)} \nabla_\theta \log p_\theta(Y_j) \right], \quad (8)$$

where the samples Y_j are drawn from the learned PMF $p_\theta(y)$. We now explain our process of obtaining the target distribution. As discussed, to reduce the variance of the MC estimator in Eq. 4, we would like our learned distribution to match the light contributions. Therefore, we define the target distribution as follows:

$$q(y|\mathbf{x}, \omega_o) = L_y(\mathbf{x}, \omega_o) L_o(\mathbf{x}, \omega_o)^{-1}. \quad (9)$$

Here, $L_y(\mathbf{x}, \omega_o)$ represents the reflected radiance at point \mathbf{x} in direction ω_o caused by a single light source. Similarly, $L_o(\mathbf{x}, \omega_o)$ denotes the radiance due to all light sources combined (see Eq. 1). Note that the second term in the above equation ensures that q is a valid PMF (summing to one).

Obtaining $L_y(\mathbf{x}, \omega_o)$ requires evaluating the integral in Eq. 1, which we approximate using Monte Carlo (MC) integration with point sampling. Specifically, we sample the integrand, $F(\mathbf{x}, \mathbf{z}, \omega_o)$, at randomly selected points \mathbf{z} on the y^{th} light. The second term, $L_o(\mathbf{x}, \omega_o)$, is unknown because it represents the quantity we aim to estimate. However, since our optimization process relies on Adam [Kingma and Ba 2015], which updates parameters based on the ratio of the current gradient to its historical average, this term can safely be omitted [Dong et al. 2023; Müller et al. 2019].

Combining all the components, the final MC estimate of the gradient can be written as:

$$\langle \nabla_\theta D_{\text{KL}}(q, p_\theta) \rangle = - \frac{1}{N} \sum_{j=1}^N \left[\frac{F(\mathbf{x}, \mathbf{Z}_j, \omega_o)}{p(\mathbf{Z}_j|Y_j)p_\theta(Y_j)} \nabla_\theta \log p_\theta(Y_j) \right]. \quad (10)$$

Here, $p(\mathbf{z}|y)$ is the probability of selecting point \mathbf{z} on the y^{th} light source, which is evaluated using existing methods [Hart et al. 2020; Shirley et al. 1996] and is not learned in our method.

This formulation shows that optimizing the loss is equivalent to minimizing the negative log likelihood weighted by the MC estimates of the integrand F . Therefore, the loss can be efficiently optimized during path tracing in an online manner.

3.2 Network Architecture

Our network, implemented in tiny-cuda-nn [Müller 2021], is composed of 3 hidden layers of 64 neurons and ReLU intermediary activation. We encode the inputs using [Dong et al. 2023]’s learnable dense grid encoding for the intersection position, and spherical harmonics with degree 4 for normalized outgoing ray direction ω_o .




Fig. 2. We present an overview of our technique. (left) The network uses local information at the shading point to estimate cluster probabilities in a residual manner. The baseline cluster distribution, w , is precomputed using existing approaches and remains fixed throughout the process. (middle) A light is sampled in two stages: first, a cluster (one of the red nodes) is selected based on the estimated cluster distribution, and then the tree is stochastically traversed (green lines) to choose a specific light within the cluster. (right) Points are subsequently sampled within the selected light source, and paths are traced. These samples are used to train the network and refine the estimated cluster distribution. This entire process is repeated iteratively until the desired number of samples is traced. Since the network estimates the cluster PMF in a residual manner, it can guide sampling from the beginning of the rendering process.

so we can model the product of radiance and BSDF. We additionally provide normal as the input to the network as done in similar learning tasks [Dong et al. 2023; Müller et al. 2019], and encode them through one-blob encoding [Müller et al. 2019] using 32 bins.

4 Neural Light Cluster Sampling

The network described in the previous section directly estimates the probability of selecting each light source. While this approach is feasible when the number of light sources M in the scene is relatively small (e.g., order of tens), directly estimating the PMF for a large number of light sources (e.g., hundreds or thousands) requires a high-capacity network, which significantly increases computational complexity and makes the approach impractical.

Inspired by existing methods [Conty Estevez and Kulla 2018; Yuksel 2019], we address this issue by adopting a clustering strategy. Specifically, we group the lights into a smaller set of S non-overlapping clusters, \mathcal{L}_c , with $c \in 1, \dots, S$. This allows us to decompose the light selection PMF $p(y)$ into two distributions: the probability of selecting cluster c , $p(c)$, and the conditional probability of selecting the y^{th} light within cluster c , $p(y|c)$. We use the network to model the cluster selection PMF $p_\theta(c)$, while the light selection within each cluster is efficiently handled using existing hierarchy techniques [Conty Estevez and Kulla 2018; Yuksel 2019].

A key advantage of this clustering approach is that it allows us to leverage existing methods to construct the light hierarchy, which also provides initial cluster probabilities. Building on this, we propose a residual learning strategy, where the network learns corrections to these initial probabilities rather than predicting them from scratch, significantly improving training convergence.

The overview of our approach is illustrated in Fig. 2. In the following sections, we describe the light hierarchy and clustering, the network optimization process, and our residual learning strategy.

4.1 Light Hierarchy and Cluster Construction

Given a scene, we first construct a light hierarchy using an existing technique [Conty Estevez and Kulla 2018; Yuksel 2019]. The nodes at a specified level k of the tree are selected as light clusters (global; fixed for the entire scene), and our network estimates their probabilities. As shown in Fig. 3, the tree is typically unbalanced in practical scenarios, so the number of internal nodes at level k may not always

equal 2^k . In such cases, we include leaf nodes from higher levels as part of the clusters, allowing the network to directly estimate the probability of a leaf node that corresponds to a single light source. For instance, in the example shown in Fig. 3, the network outputs probabilities for the 7 cluster nodes (highlighted in red).

To sample a light source, we first select a cluster based on the cluster selection PMF $p_\theta(c)$. Next, we traverse the tree from the selected cluster to a leaf node, selecting one child node at each step based on their PMFs. These PMFs are computed using the same method used to construct the light hierarchy [Conty Estevez and Kulla 2018; Yuksel 2019]. The conditional light selection PMF $p(y|c)$ is then determined by multiplying the probabilities of all child nodes along the traversal path, as illustrated by the green lines in Fig. 3.

4.2 Optimization

We optimize the cluster selection network using the formulation described in Sec. 3.1. Consequently, the gradient of the KL-divergence loss can be approximated in a manner similar to Eq. 10. The key difference is that the light selection PMF is now expressed as the product of the cluster selection PMF, estimated by the network, and the light selection PMF within a cluster, computed using existing methods. Substituting $p_\theta(y)$ with $p(y|c)p_\theta(c)$ in Eq. 10, we derive the approximated gradient:

$$\langle \nabla_\theta D_{\text{KL}}(q, p_\theta) \rangle = -\frac{1}{N} \sum_{j=1}^N \left[\frac{F(\mathbf{x}, \mathbf{Z}_j, \omega_o)}{p(\mathbf{Z}_j|Y_j)p(Y_j|C_j)p_\theta(C_j)} \nabla_\theta \log \left(p(Y_j|C_j)p_\theta(C_j) \right) \right].$$

Since $p(y|c)$ is independent of θ , we can simplify the log term as:

$$\langle \nabla_\theta D_{\text{KL}}(q, p_\theta) \rangle = -\frac{1}{N} \sum_{j=1}^N \left[\frac{F(\mathbf{x}, \mathbf{Z}_j, \omega_o)}{p(\mathbf{Z}_j|Y_j)p(Y_j|C_j)p_\theta(C_j)} \nabla_\theta \log p_\theta(C_j) \right]. \quad (11)$$

By optimizing this objective, the network increases the likelihood of clusters that correspond to light samples with high contributions, while reducing the likelihood of clusters associated with light samples with low contributions.




Fig. 3. Given a light hierarchy constructed using an existing method [Conty Estevez and Kulla 2018; Yuksel 2019], we select nodes at a specific level ($k = 3$ in this example) as light clusters, with the network estimating their probabilities $p_\theta(c)$. The probability of selecting an individual light (e.g., the blue leaf node) is computed as the product of the cluster probability and the conditional probability of the light given the cluster, $p(y|c)$. The conditional probability is obtained by multiplying the probabilities of the child nodes along the tree traversal path (green lines). Note that the grayed-out nodes are used solely for defining the global clusters and are not involved in sampling or evaluating the PMF.

4.3 Residual Learning

As discussed, training is performed in an online manner during the rendering process. We use the randomly initialized network to generate samples that are used as training data to update the network. This process is iteratively continued until a desired number of samples are traced. However, this process has two major issues: (1) the network requires a large number of iterations to converge to a meaningful PMF due to its random initialization, and (2) during this phase, the generated samples exhibit high variance, increasing the noise in final rendering and hindering effective training.

To address these issues, we observe that existing methods [Conty Estevez and Kulla 2018; Yuksel 2019] estimate cluster probabilities using carefully designed rules, which are significantly better than the initial PMF predicted by the network. Therefore, we propose to use these cluster probabilities as a baseline, with the network predicting the residual. When combining the network’s estimate with the baseline, two key requirements must be met: (1) all cluster probabilities must be positive, and (2) their sum must equal one.

We satisfy the first requirement by performing the combination in the log domain, ensuring that all values remain positive after applying the inverse log operation, while the second requirement is met by normalization. Specifically, given the importance weights (where probabilities are normalized importance weights) for all nodes in a cluster, $\mathbf{w} = w_1, \dots, w_S$, we compute the final cluster PMF by combining the network’s output with these importance weights as follows:

$$p_\theta(c) = \frac{e^{[\log(w_c) + f_\theta(\mathbf{x}, \omega_o)[c]]}}{\sum_{s=1}^S e^{[\log(w_s) + f_\theta(\mathbf{x}, \omega_o)[c]]}}, \quad (12)$$

where $f_\theta(\mathbf{x}, \omega_o)[c]$ is the estimated residual for the c^{th} cluster.

This approach addresses the aforementioned issues effectively. Initially, the cluster probabilities are close to those estimated by existing methods because the randomly initialized network produces outputs near zero. This ensures that the generated samples at the start are comparable to those from the baseline approach. As training progresses, the network refines its predictions, resulting in cluster PMFs that consistently improve upon the baseline.

Table 1. Equal-time comparisons measured in \mathbb{FLIP} [Andersson et al. 2020] with number of samples per pixel for eight representative scenes. We compare our approach against ATS [Conty Estevez and Kulla 2018], SLCRT [Lin and Yuksel 2020], ReSTIR [Bitterli et al. 2020], and VARL [Wang et al. 2021]. We color code the first, second, and third lowest numbers.

	ATS	SLCRT	ReSTIR	VARL	Ours
BATHROOM	0.1659 212	0.1413 220	0.1744 135	0.1346 139	0.0977 175
BEDROOM	0.3129 192	0.2442 196	0.2796 123	0.1960 123	0.0910 157
BISTRO	0.1548 163	0.2656 165	0.2837 74	0.2281 60	0.1126 168
LIVING ROOM	0.1247 194	0.0977 194	0.1314 131	0.1009 131	0.0671 146
SAN MIGUEL	0.2253 159	0.3015 158	0.2817 110	0.2472 107	0.1713 164
STAIRCASE	0.0763 198	0.0817 203	0.1059 129	0.0946 138	0.0529 160
STAIRCASE2	0.1020 233	0.0925 233	0.1132 159	0.0854 149	0.0646 191
ZERO DAY	0.1647 222	0.1455 223	0.1682 159	0.1338 161	0.1117 201

5 Results

In this section, we present visual and numerical comparisons to demonstrate the effectiveness of our approach. Specifically, we compare our method against ATS [Conty Estevez and Kulla 2018], SLCRT [Lin and Yuksel 2020], ReSTIR [Bitterli et al. 2020], and VARL [Wang et al. 2021]. The ATS implementation [Conty Estevez and Kulla 2018] is available in PBRT v4, while we have implemented the other three approaches. ReSTIR is implemented with $M = 32$ samples with $K = 3$ spatial reuse samples on a radius of 30 pixels. For VARL, which is designed for CPU, we convert their sequential update to a batch update to make it compatible with GPU. Our VARL implementation is significantly faster than the original version and produces better results (see supplementary material).

5.1 Comparisons

We perform equal-time and equal-sample comparisons on eight scenes that contain complex lighting. The number of lights per scene is: BATHROOM (4.7k), BEDROOM (7.9k), BISTRO (21.1k), LIVINGROOM (64), STAIRCASE (6.7k), STAIRCASE2 (2.1k) and ZERO DAY (8.9k). For SAN MIGUEL, we trace 139k virtual point lights (VPLs). For all the results, we compare direct illumination at first non-specular bounce. We perform all comparisons and ablations on an Intel i9-14900K CPU and an NVIDIA RTX4090 GPU. Here, we report quantitative results using \mathbb{FLIP} [Andersson et al. 2020]. Evaluation using additional metrics can be found in the supplementary. In all comparisons, we use ATS [Conty Estevez and Kulla 2018] to construct the light hierarchy, and we estimate node probabilities using our network at level $k = 6$. For optimization, we use the Adam [Kingma and Ba 2015] optimizer with default parameters and a fixed learning rate of 3×10^{-2} , and choose a training budget of 15% for our method.

Equal Time. All approaches are compared both quantitatively (Table 1) and qualitatively (Fig. 4) in an equal time setting with the time budget increasing with scene complexity and resolution. For BATHROOM, BEDROOM, LIVING ROOM, STAIRCASE, and STAIRCASE2, we allocate a budget of 5 seconds to each approach. For SAN MIGUEL and ZERO DAY, we allocate 10 seconds, while, for BISTRO, 20 seconds are allocated. As shown, our approach produces significantly better results than the others, both numerically and visually. In particular,

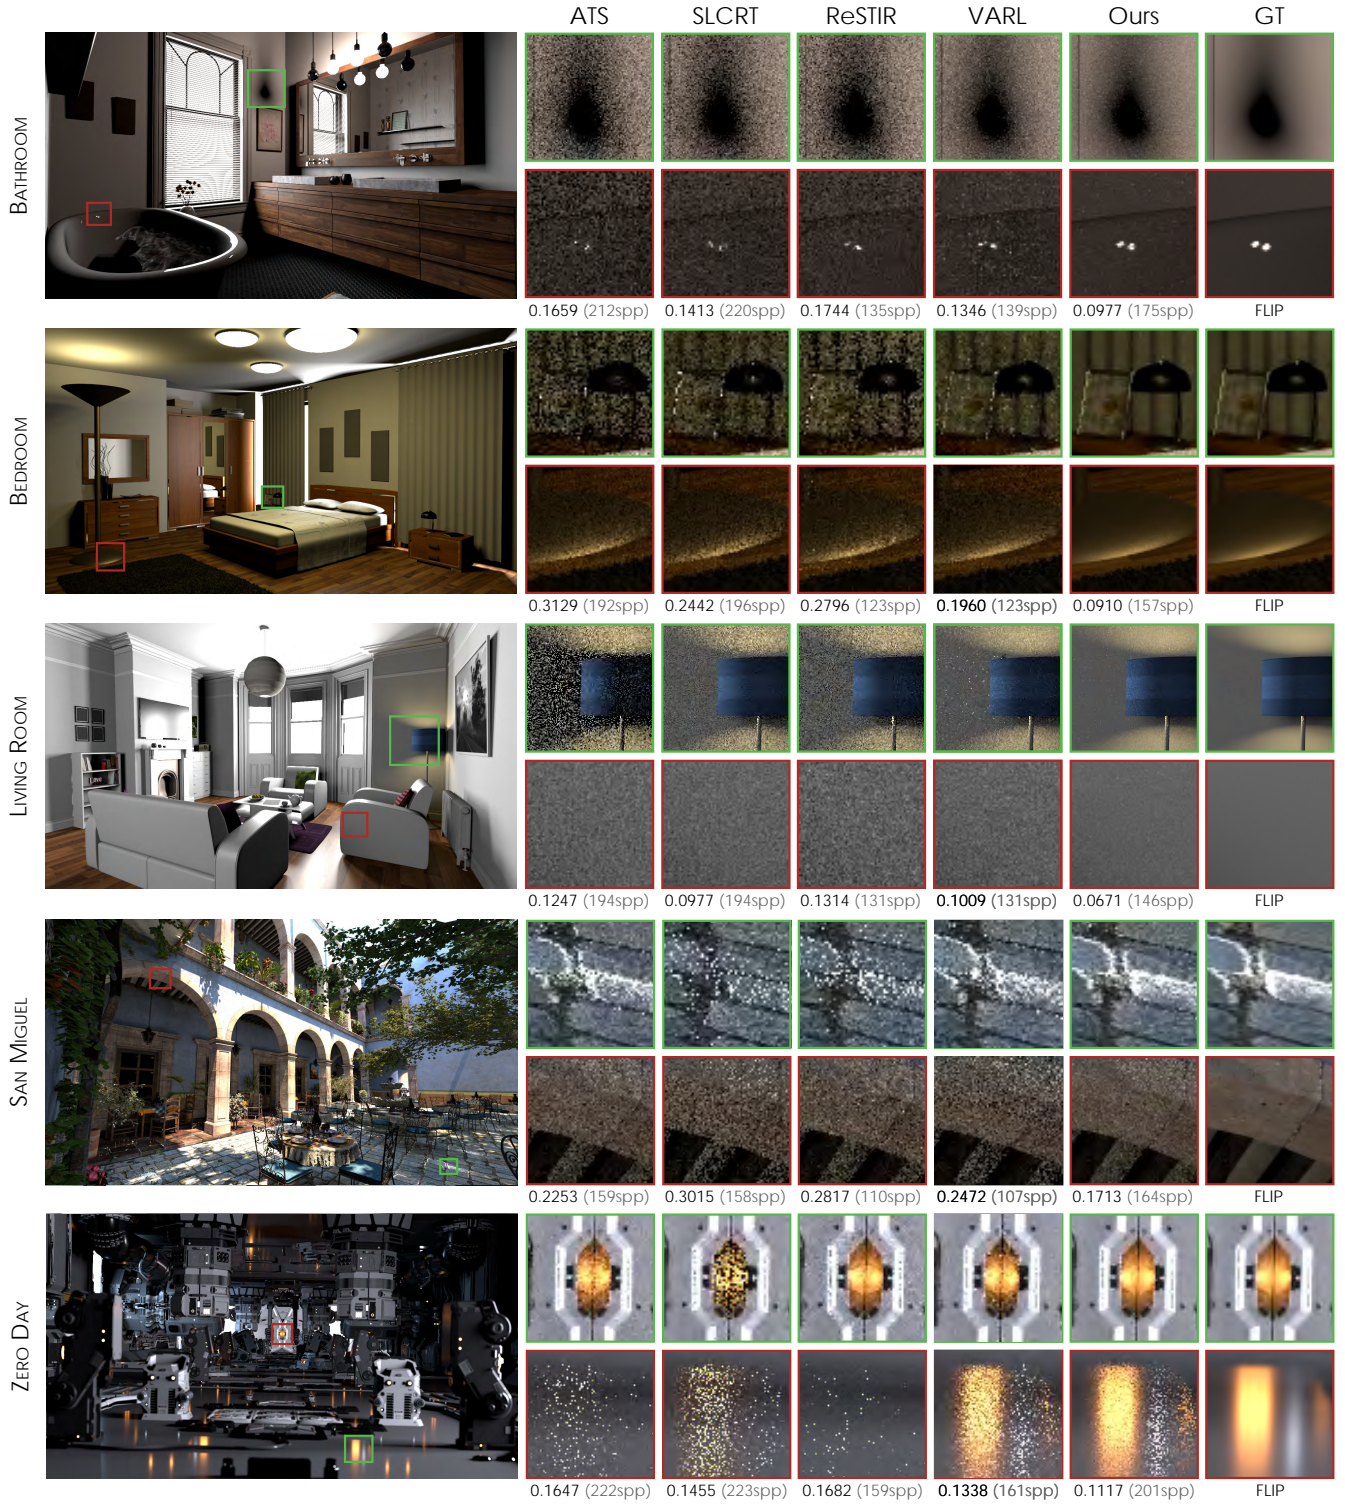


Fig. 4. Equal-time comparison of ATS [Conty Estevez and Kulla 2018], SLCRT [Lin and Yuksel 2020], ReSTIR [Bitterli et al. 2020], VARL [Wang et al. 2021], and our method. The time budget increases with scene complexity and resolution. For BATHROOM, BEDROOM and LIVING ROOM, the budget is 5 seconds; ZERO DAY and SAN MIGUEL 10 seconds.

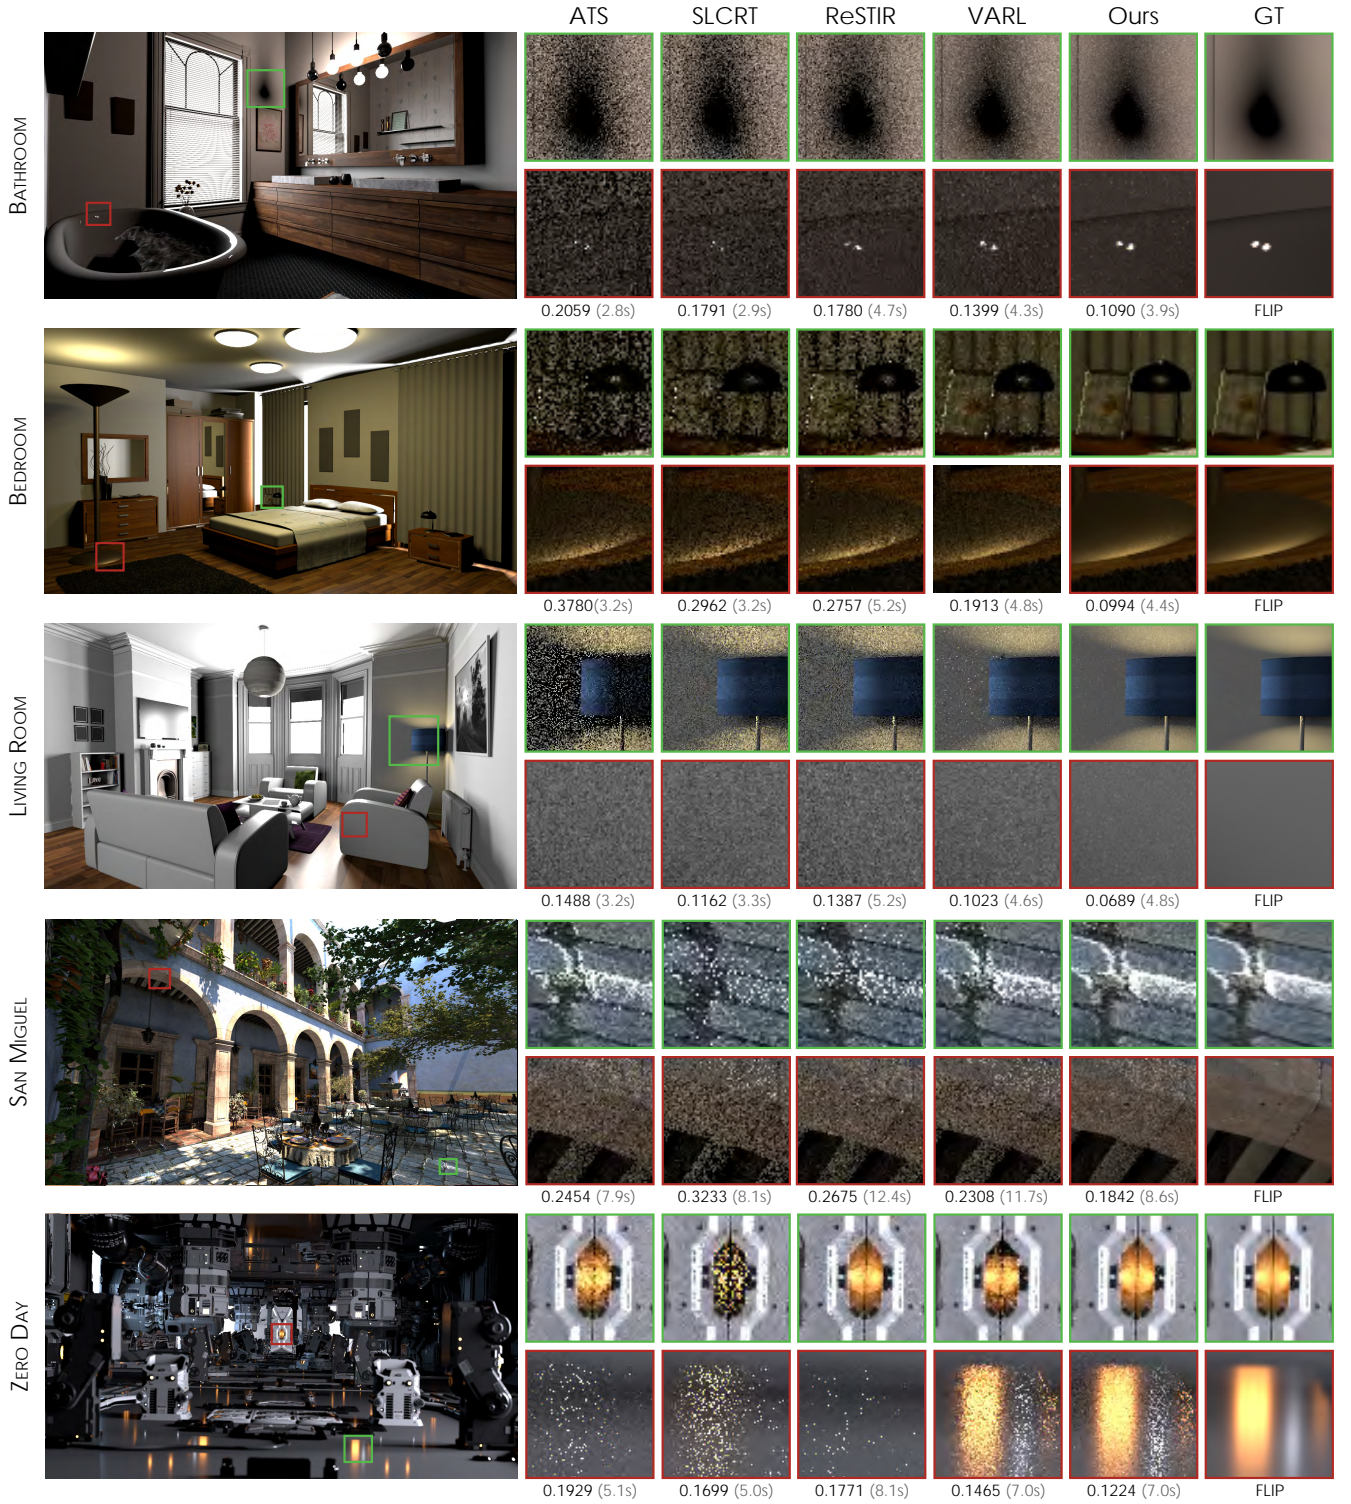


Fig. 5. Equal-sample comparison against several state-of-the-art methods. We use 128 spp for all the scenes.




Fig. 6. Equal-time convergence plots of all the approaches on the eight scenes. The time budget increases with scene complexity. For all scenes, the second vertical line is approximately where our method stops learning and uses the learned distributions to sample the remaining paths.




Fig. 7. Equal-sample convergence plots of all the approaches on the eight scenes from 8 to 128 spp. The second vertical line is approximately where our method stops learning and uses the learned distributions to sample the remaining paths.

Table 2. Equal-sample comparisons measured in FLIP [Andersson et al. 2020] with runtime in seconds. All the scenes are rendered using each approach for 128 samples per pixel.

	ATS		SLCRT		ReSTIR		VARL		Ours	
BATHROOM	0.2059	2.8	0.1791	2.9	0.1780	4.7	0.1399	4.3	0.1090	3.9
BEDROOM	0.3780	3.2	0.2962	3.2	0.2757	5.2	0.1913	4.8	0.0994	4.4
BISTRO	0.1704	15.9	0.2944	15.3	0.2316	36.8	0.1658	44.3	0.1180	19.6
LIVING ROOM	0.1488	3.2	0.1162	3.3	0.1387	5.2	0.1023	4.6	0.0689	4.8
SAN MIGUEL	0.2454	7.9	0.3233	8.1	0.2675	12.4	0.2308	11.7	0.1842	8.6
STAIRCASE	0.0923	3.3	0.1002	3.2	0.1060	4.9	0.0969	4.7	0.0553	4.5
STAIRCASE2	0.1246	2.7	0.1128	2.8	0.1206	4.0	0.0913	4.1	0.0721	3.8
ZERO DAY	0.1929	5.1	0.1699	5.0	0.1771	8.1	0.1465	7.0	0.1224	7.0

our method best handles specular reflections (see the red insets for BATHROOM and ZERO DAY in Fig. 4), and seamlessly learns local lighting variations such as shadows (see the green insets for BATHROOM and SAN MIGUEL).

Equal Sample. We compare all the approaches in an equal sample setting to evaluate their effectiveness regardless of speed. We render each scene using 128 samples per pixel, report the FLIP [Andersson et al. 2020] metric in Table 2, and compare all methods qualitatively in Fig. 5. Similar to equal-time results, our method ranks first in all scenes despite training for only 15% of the allocated samples. Additionally, our approach shows competitive runtime against fast tree-based approaches, highlighting the small overhead of our neural training and sampling strategies.

Convergence Analysis. To understand the convergence behavior of our approach, we plot the FLIP [Andersson et al. 2020] metrics of rendered images in equal-time (Fig. 6) and equal-sample (Fig. 7). Overall, our approach consistently produces better results than all




Fig. 8. Comparison of residual learning strategy using SLC [Yuksel 2019], SLCRT [Lin and Yuksel 2020], and ATS [Conty Estevez and Kulla 2018] baseline methods against the version of our technique where the cluster probabilities are directly estimated by the network (Direct). As shown, residual learning with any baseline method produces significantly better results than direct estimation with minimal timing overhead.

other methods, especially at higher spp counts. We note that, in some scenes, other methods produce better results initially in equal-time scenarios, but our performance improves significantly once the training phase ends and we can trace more samples.

5.2 Ablations

Residual Learning. Figure 8 compares our proposed residual learning approach using SLC [Yuksel 2019], SLCRT [Lin and Yuksel 2020], and ATS [Conty Estevez and Kulla 2018] baseline methods to direct network prediction. The results show that residual learning greatly speeds up convergence while adding only minimal runtime overhead, with ATS chosen as our baseline since it ranks first.

Input Encoding. Encodings for input position are evaluated in an equal-sample setting in Fig. 9. It compares identity, frequency [Mildenhall et al. 2020] (12 frequencies), one-blob [Müller et al. 2019] (32 bins), and the learnable dense grid [Dong et al. 2023] encodings, with the latter being chosen because it exhibits the least noise.

Training Budget Ratio. Figure 10 shows that training budget ratios significantly influence quality in an equal-time setting. A 15%




Fig. 9. Evaluation of input encodings for the position in the equal-sample LIVING ROOM scene. As shown, the learnable dense grid [Dong et al. 2023] performs best with a small overhead.




Fig. 10. Impact of training budget ratios on the quality of the final render, measured in FLIP [Andersson et al. 2020], for the LIVING ROOM scene with 5 seconds of time budget. We choose the optimal ratio of 15%.

Table 3. Analysis of the effect of continuous representation and the cluster level. On the left, the results show that our approach is superior to a discrete version in all scenes. On the right, the comparisons indicate that the cluster level $k = 6$ performs best on most scenes. We use the FLIP [Andersson et al. 2020] metric and highlight the lowest (best) scores in **bold**.

	Effect of Continuous Representation		Cluster Level		
	Ours Discrete	Ours	k=4	k=6	k=8
BATHROOM	0.1138	0.1090	0.0969	0.0977	0.1081
BEDROOM	0.1027	0.0994	0.0924	0.0910	0.1021
BISTRO	0.1321	0.1217	0.1269	0.1157	0.1237
LIVING ROOM	0.0778	0.0689	0.0617	0.0665	0.0668
SAN MIGUEL	0.1891	0.1841	0.1856	0.1713	0.1782
STAIRCASE	0.0732	0.0695	0.0683	0.0630	0.0672
STAIRCASE2	0.0758	0.0721	0.0712	0.0646	0.0673
ZERO DAY	0.1387	0.1224	0.1230	0.1117	0.1377

training ratio achieves the highest quality in the 5-second LIVING ROOM scene, and similar trends are observed across all scenes.

Effect of Continuous Representation. Table 3 (left) evaluates the impact of our continuous representation versus a discrete alternative. The discrete version takes as input the center of the nearest uniform grid cell with resolution $32 \times 32 \times 32$ and the normalized outgoing ray direction discretized into 8 buckets per dimension [Wang et al. 2021]. As seen, our method benefits from continuous input in all scenes, demonstrating its ability to encode local variations.

Cluster Level. We assess how varying cluster levels (see Fig. 3) affects the final result. Table 3 (right) shows numerical comparisons for three cluster levels ($k = 4, 6, 8$) on all scenes. The comparisons indicate that the chosen $k = 6$ performs best in most scenes.

6 Conclusion, Limitations, and Future Work

We have presented a neural light importance sampling approach for rendering scenes with many lights. Our method employs a neural network to predict spatially varying light selection distributions using local information at each shading point. The network is trained




Fig. 11. Comparison of learning rates in the equal-time LIVING ROOM scene. Using an aggressive learning rate of 3×10^{-2} yields lower noise levels than a more stable learning rate of 3×10^{-3} at the cost of outlier fireflies.

online by minimizing the KL-divergence between the learned and target distributions. To efficiently handle scenes with a large number of lights, we integrate our neural approach with existing light hierarchy techniques. Additionally, we introduce a residual training strategy to improve convergence. Extensive experiments demonstrate that our approach outperforms existing methods on challenging scenes with complex lighting.

Despite its superior performance, our approach has some limitations. For instance, in some cases, it produces outlier fireflies due to our aggressive optimization strategy. During training, the network may assign low probability (in initial iterations) to important, unobserved light clusters. If a light with high power from such a cluster is sampled, it significantly contributes to the final pixel color, causing the outlier. Using a more conservative optimization strategy produces fewer outliers but results in slower convergence (Fig. 11).

Furthermore, we use a fixed global cut for the entire scene and rely on existing methods to sample lights within a cluster. Consequently, if the baseline method produces poor light sampling probabilities, our approach cannot fully correct these errors. One way to address this limitation is to locally adapt the cut, as demonstrated in existing methods [Pantaleoni 2020; Wang et al. 2021]. We leave the combination of our approach with this strategy for future work.

Moreover, we have focused solely on learning the light sampling distributions. However, sampling points within light sources also significantly impacts the variance of the Monte Carlo (MC) estimator, particularly for larger light sources. Learning the sampling distribution within a light source is another avenue for future work. Finally, we are interested in combining our approach with Zhu et al.’s method [2021a] to effectively handle scenes with a large number of complex luminaires. The key challenge lies in formulating a framework that seamlessly optimizes both our network and the luminaire network, which we leave for future exploration.

Acknowledgments

This project was funded in part by the NSF CAREER Award #2238193. We thank Ryusuke Villemin for the valuable discussions, Vincent Serritella for generating the experimental test data, and Magnus Wrenninge and the Aurora simulation team for support. We are grateful to Wang et al. [2021] 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: Marek (BATHROOM), SlykDrako (BEDROOM), Amazon Lumberyard [2017] (BISTRO), Jay-Artist (LIVING ROOM), Guillermo M. Leal Llaguno (SAN MIGUEL), Wig42 (STAIRCASE), NewSee2l035 (STAIRCASE2), and Mike Winkelmann [2019] (ZERO DAY).

References

- Pontus Andersson, Jim Nilsson, Tomas Akenine-Möller, Magnus Oskarsson, Kalle Åström, and Mark D. Fairchild. 2020. FLIP: A Difference Evaluator for Alternating Images. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 3, 2 (2020), 15:1–15:23.
- Steve Bako, Mark Meyer, Tony DeRose, and Pradeep Sen. 2019. Offline Deep Importance Sampling for Monte Carlo Path Tracing. *Computer Graphics Forum* 38, 7 (2019), 527–542.
- Benedikt Bitterli, Chris Wyman, Matt Pharr, Peter Shirley, Aaron Lefohn, and Wojciech Jarosz. 2020. Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting. *ACM Trans. Graph.* 39, 4, Article 148 (Aug. 2020), 17 pages. doi:10.1145/3386569.3392481
- Alejandro Conty Estevez and Christopher Kulla. 2018. Importance Sampling of Many Lights with Adaptive Tree Splitting. *Proc. ACM Comput. Graph. Interact. Tech.* 1, 2, Article 25 (Aug. 2018), 17 pages. doi:10.1145/3233305
- Honghao Dong, Guoping Wang, and Sheng Li. 2023. Neural Parametric Mixtures for Path Guiding. In *ACM SIGGRAPH 2023 Conference Proceedings (SIGGRAPH '23)*. Association for Computing Machinery, New York, NY, USA, Article 29, 10 pages.
- Michael Donikian, Bruce Walter, Kavita Bala, Sebastian Fernandez, and Donald P. Greenberg. 2006. Accurate Direct Illumination Using Iterative Adaptive Sampling. *IEEE Transactions on Visualization and Computer Graphics* 12, 3 (May 2006), 353–364. doi:10.1109/TVCG.2006.41
- Sebastian Fernandez, Kavita Bala, and Donald P. Greenberg. 2002. Local illumination environments for direct lighting acceleration. In *Proceedings of the 13th Eurographics Workshop on Rendering (Pisa, Italy) (EGRW '02)*. Eurographics Association, Goslar, DEU, 7–14.
- Iliyan Georgiev, Jaroslav Krivánek, Stefan Popov, and Philipp Slusallek. 2012. Importance Caching for Complex Illumination. *Computer Graphics Forum* 31, 2pt3 (2012), 701–710. doi:10.1111/j.1467-8659.2012.03049.x
- D. Hart, M. Pharr, T. Müller, W. Lopes, M. McGuire, and P. Shirley. 2020. Practical Product Sampling by Fitting and Composing Warps. *Computer Graphics Forum* 39, 4 (2020), 149–158.
- Miloš Hašan, Fabio Pellacini, and Kavita Bala. 2007. Matrix row-column sampling for the many-light problem. *ACM Trans. Graph.* 26, 3 (July 2007), 26–es. doi:10.1145/1276377.1276410
- Miloš Hašan, Edgar Velázquez-Armendáriz, Fabio Pellacini, and Kavita Bala. 2008. Tensor Clustering for Rendering Many-Light Animations. *Computer Graphics Forum* 27, 4 (2008), 1105–1114. doi:10.1111/j.1467-8659.2008.01248.x
- Jiawei Huang, Akito Iizuka, Hajime Tanaka, Taku Komura, and Yoshifumi Kitamura. 2024. Online Neural Path Guiding with Normalized Anisotropic Spherical Gaussians. *ACM Trans. Graph.* 43, 3, Article 26 (April 2024), 18 pages.
- Yuchi Huo, Rui Wang, Shihao Jin, Xinguo Liu, and Hujun Bao. 2015. A matrix sampling-and-recovery approach for many-lights rendering. *ACM Trans. Graph.* 34, 6, Article 210 (Nov. 2015), 12 pages. doi:10.1145/2816795.2818120
- Yuchi Huo, Rui Wang, Ruzhang Zheng, Hualin Xu, Hujun Bao, and Sung-Eui Yoon. 2020. Adaptive Incident Radiance Field Sampling and Reconstruction Using Deep Reinforcement Learning. *ACM Trans. Graph.* 39, 1, Article 6 (jan 2020), 17 pages.
- Diederik Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*.
- Arjan J. F. Kok and Frederik W. Jansen. 1994. Source Selection for the Direct Lighting Computation in Global Illumination. In *Photorealistic Rendering in Computer Graphics*, P. Brunet and F. W. Jansen (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 75–82.
- Daqi Lin, Markus Kettunen, Benedikt Bitterli, Jacopo Pantaleoni, Cem Yuksel, and Chris Wyman. 2022. Generalized resampled importance sampling: foundations of ReSTIR. *ACM Trans. Graph.* 41, 4, Article 75 (July 2022), 23 pages. doi:10.1145/3528223.3530158
- Daqi Lin and Cem Yuksel. 2020. Real-time stochastic lightcuts. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 3, 1 (2020), 1–18.
- Yifan Liu, Kun Xu, and Ling-Qi Yan. 2019. Adaptive BRDF-Oriented Multiple Importance Sampling of Many Lights. *Computer Graphics Forum* 38, 4 (2019), 123–133. doi:10.1111/cgf.13776 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13776
- Haolin Lu, Wesley Chang, Trevor Hedstrom, and Tzu-Mao Li. 2024. Real-Time Path Guiding Using Bounding Voxel Sampling. *ACM Trans. Graph.* 43, 4, Article 125 (July 2024), 14 pages. doi:10.1145/3658203
- Amazon Lumberyard. 2017. Amazon Lumberyard Bistro, Open Research Content Archive (ORCA). <http://developer.nvidia.com/orca/amazon-lumberyard-bistro> <http://developer.nvidia.com/orca/amazon-lumberyard-bistro>.
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.
- P. Moreau, M. Pharr, and P. Clarberg. 2022. Dynamic many-light sampling for real-time ray tracing. In *Proceedings of the Conference on High-Performance Graphics (Strasbourg, France) (HPG '19)*. Eurographics Association, Goslar, DEU, 21–26. doi:10.2312/hpg.20191191
- Thomas Müller. 2021. *tiny-cuda-nn*. <https://github.com/NVlabs/tiny-cuda-nn>
- Thomas Müller, Brian McWilliams, Fabrice Rousselle, Markus Gross, and Jan Novák. 2019. Neural Importance Sampling. *ACM Trans. Graph.* 38, 5, Article 145 (oct 2019), 19 pages.
- Thomas Müller, Markus Gross, and Jan Novák. 2017. Practical Path Guiding for Efficient Light-Transport Simulation. *Computer Graphics Forum* 36, 4 (2017), 91–100.
- Jiawei Ou and Fabio Pellacini. 2011. LightSlice: matrix slice sampling for the many-lights problem. *ACM Trans. Graph.* 30, 6 (Dec. 2011), 1–8. doi:10.1145/2070781.2024213
- Y. Ouyang, S. Liu, M. Kettunen, M. Pharr, and J. Pantaleoni. 2021. ReSTIR GI: Path Resampling for Real-Time Path Tracing. *Computer Graphics Forum* 40, 8 (2021), 17–29. doi:10.1111/cgf.14378
- Jacopo Pantaleoni. 2019. Importance Sampling of Many Lights with Reinforcement Lightcuts Learning. arXiv:1911.10217 [cs.GR] <https://arxiv.org/abs/1911.10217>
- Jacopo Pantaleoni. 2020. Online path sampling control with progressive spatio-temporal filtering. arXiv:2005.07547 [cs.GR]
- Eric Paquette, Pierre Poulin, and George Drettakis. 1998. A Light Hierarchy for Fast Rendering of Scenes with Many Lights. *Computer Graphics Forum* 17, 3 (1998), 63–74. doi:10.1111/1467-8659.00254
- Alexander Rath, Pascal Grittmann, Sebastian Herholz, Petr Vévoda, Philipp Slusallek, and Jaroslav Krivánek. 2020. Variance-aware path guiding. *ACM Trans. Graph.* 39, 4, Article 151 (aug 2020), 12 pages.
- Peter Shirley, Changyaw Wang, and Kurt Zimmerman. 1996. Monte Carlo techniques for direct lighting calculations. *ACM Trans. Graph.* 15, 1 (Jan. 1996), 1–36. doi:10.1145/226150.226151
- Justin F. Talbot, David Cline, and Parris Egbert. 2005. Importance resampling for global illumination. In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques (Konstanz, Germany) (EGSR '05)*. Eurographics Association, Goslar, DEU, 139–146.
- Yusuke Tokuyoshi, Sho Ikeda, Paritosh Kulkarni, and Takahiro Harada. 2024. Hierarchical Light Sampling with Accurate Spherical Gaussian Lighting. In *SIGGRAPH Asia 2024 Conference Papers (SA '24)*. Association for Computing Machinery, New York, NY, USA, Article 82, 11 pages. doi:10.1145/3680528.3687647
- Petr Vévoda, Ivo Kondapaneni, and Jaroslav Krivánek. 2018. Bayesian online regression for adaptive direct illumination sampling. *ACM Trans. Graph.* 37, 4, Article 125 (July 2018), 12 pages. doi:10.1145/3197517.3201340
- Jiří Vorba, Johannes Hanika, Sebastian Herholz, Thomas Müller, Jaroslav Krivánek, and Alexander Keller. 2019. Path Guiding in Production. In *ACM SIGGRAPH 2019 Courses (Los Angeles, California) (SIGGRAPH '19)*. ACM, New York, NY, USA, Article 18, 77 pages.
- Bruce Walter, Sebastian Fernandez, Adam Arbree, Kavita Bala, Michael Donikian, and Donald P. Greenberg. 2005. Lightcuts: a scalable approach to illumination. *ACM Trans. Graph.* 24, 3 (July 2005), 1098–1107. doi:10.1145/1073204.1073318
- Yu-Chen Wang, Yu-Ting Wu, Tzu-Mao Li, and Yung-Yu Chuang. 2021. Learning to cluster for rendering with many lights. *ACM Trans. Graph.* 40, 6, Article 277 (Dec. 2021), 10 pages. doi:10.1145/3478513.3480561
- Gregory J. Ward. 1994. Adaptive Shadow Testing for Ray Tracing. In *Photorealistic Rendering in Computer Graphics*, P. Brunet and F. W. Jansen (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 11–20.
- Mike Winkelmann. 2019. Zero-Day, Open Research Content Archive (ORCA). <https://developer.nvidia.com/orca/beeple-zero-day> <https://developer.nvidia.com/orca/beeple-zero-day>.
- Yu-Ting Wu and Yung-Yu Chuang. 2013. VisibilityCluster: Average Directional Visibility for Many-Light Rendering. *IEEE Transactions on Visualization and Computer Graphics* 19, 9 (2013), 1566–1578. doi:10.1109/TVCG.

2013.21

- Cem Yuksel. 2019. Stochastic Lightcuts. In *High-Performance Graphics (HPG 2019)* (Strasbourg, France). The Eurographics Association, 27–32. doi:10.2312/hpg.20191192
- Junqiu Zhu, Yaoyi Bai, Zilin Xu, Steve Bako, Edgar Velázquez-Armendáriz, Lu Wang, Pradeep Sen, Miloš Hašan, and Ling-Qi Yan. 2021a. Neural complex luminaires: representation and rendering. *ACM Trans. Graph.* 40, 4, Article 57 (July 2021), 12 pages.
- Shilin Zhu, Zexiang Xu, Tiancheng Sun, Alexandr Kuznetsov, Mark Meyer, Henrik Wann Jensen, Hao Su, and Ravi Ramamoorthi. 2021b. Hierarchical neural reconstruction for path guiding using hybrid path and photon samples. *ACM Trans. Graph.* 40, 4, Article 35 (jul 2021), 16 pages.
- Shilin Zhu, Zexiang Xu, Tiancheng Sun, Alexandr Kuznetsov, Mark Meyer, Henrik Wann Jensen, Hao Su, and Ravi Ramamoorthi. 2021c. Photon-Driven Neural Reconstruction for Path Guiding. *ACM Trans. Graph.* 41, 1, Article 7 (Nov. 2021), 15 pages.