Focal Path Guiding for Light Transport Simulation

Alexander Rath
Intel Corporation
Karlsruhe, Germany
DFKI
Saarbrücken, Germany
rath@cg.uni-saarland.de

Ömercan Yazici
Saarland University
Saarbrücken, Germany
yazici@cg.uni-saarland.de

Philipp Slusallek
Saarland University
Saarbrücken, Germany
DFKI
Saarbrücken, Germany
philipp.slusallek@dfki.de

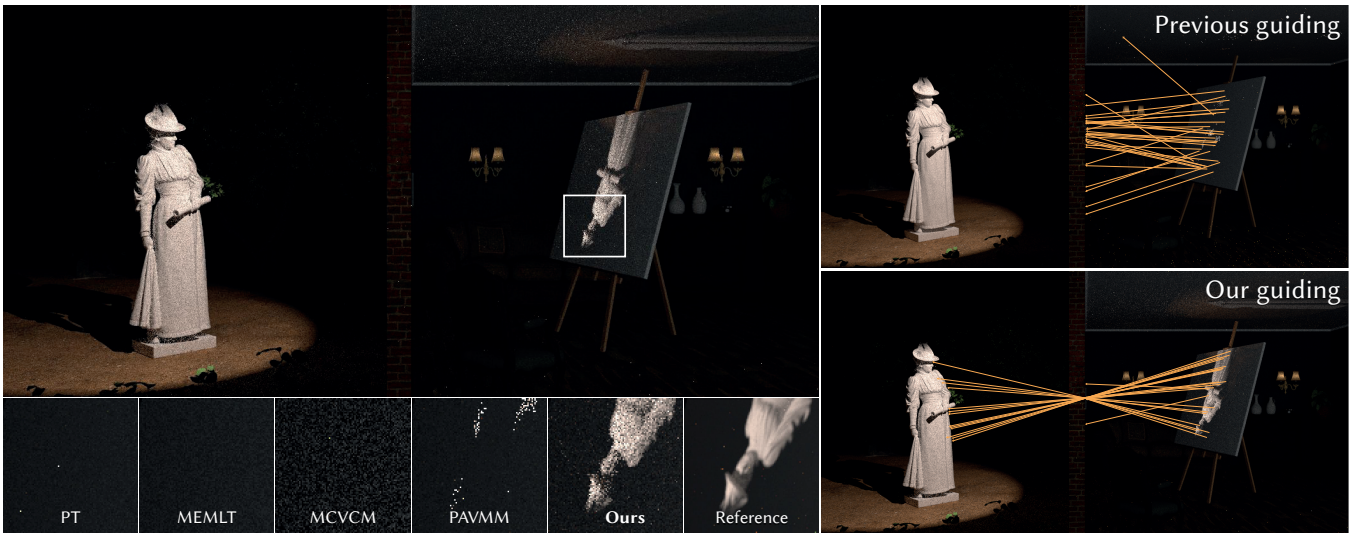

Figure 1: Focal points can emerge from various situations. Here we demonstrate the pinhole effect as part of a *camera obscura*, in which a focal point is created as rays are forced to converge in a small region of space. A small hole in the wall causes the brightly illuminated statue on the left to be projected onto the canvas on the right. This phenomenon is challenging even for sophisticated algorithms, as the location in which rays converge is not generally known beforehand, hindering the systematic exploration of such effects. Path guiding utilizes learning to react to unforeseen effects, but even state-of-the-art approaches such as PAVMM [Ruppert et al. 2020] fail to reconstruct focal points reliably. We propose a novel form of guiding, based on identifying and explicitly sampling focal points, which significantly accelerates convergence for various kinds of focal effects.

ABSTRACT

Focal points are fascinating effects that emerge from various constellations, for example when light passes through narrow gaps or when objects are seen through lenses or mirrors. These effects can be challenging to render, as paths need to pass through small regions that are not always known beforehand and can occur freely in space. Specialized algorithms exist for some effects, but many of them rely on Markov chain Monte Carlo integration, which is known to suffer from uneven convergence undesirable in practice. Path guiding methods are a promising alternative, but existing techniques only handle a subset of focal effects. We propose a novel form of guiding that is specifically tailored to identify focal points and sample them in accordance to their image contribution. Our technique is the first to unify all focal effects in a single framework and we demonstrate that it can render effects that previous state-of-the-art techniques are unable to handle.

CCS CONCEPTS

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

KEYWORDS

ray tracing, global illumination, Monte Carlo, multiple importance sampling

ACM Reference Format:
Alexander Rath, Ömercan Yazici, and Philipp Slusallek. 2023. Focal Path Guiding for Light Transport Simulation. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Proceedings (SIGGRAPH '23 Conference Proceedings), August 06–10, 2023, Los Angeles, CA, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3588432.3591543

SIGGRAPH '23 Conference Proceedings, August 06–10, 2023, Los Angeles, CA, USA

ACM ISBN 979-8-4007-0159-7/23/08.
https://doi.org/10.1145/3588432.3591543

1 INTRODUCTION

The desire to synthesize realistic images from virtual scenes is ubiquitous—be it in the entertainment industry, medical visualization, architecture or e-commerce. Despite decades of research, there is still no single algorithm that robustly handles all intricate ways in which light interacts with matter. A particularly challenging class of phenomena is that of focal points. These occur whenever light converges in small regions, for example when passing narrow gaps—like a keyhole in a closed door—or when beams of light are focused by lenses or reflectors. Even sophisticated methods struggle with such effects, as the location and extent of such regions is not always known beforehand.

A promising approach that can adapt to unforeseen light transport is path guiding, which uses information from earlier samples to learn from where to expect light. Once an important source of light is identified, the following samples thoroughly explore it to arrive at a noise-free render. While this approach has the potential to handle focal effects, we find that existing guiding methods struggle to learn accurate representations of focal regions. This is especially troublesome when focal points occur in free space, as demonstrated in Figure 1.

We propose a novel form of guiding, which is specifically tailored to handle focal effects. Instead of learning a directional or path space distribution like previous work, our method learns a spatial distribution that is not restricted to surfaces. Our method can robustly identify and sample focal points, even those in free space. We categorize focal effects into different classes, and show that our method is the first to unify all of them in a single framework. For many classes of focal points, our algorithm achieves substantial improvements over the respective previous state-of-the-art.

In summary, we make the following contributions:

- We identify and classify common sources of focal points and survey how they are handled by existing families of algorithms (Section 3)
- We introduce an iterative scheme which identifies relevant regions of focal light transport in a given scene (Section 4)
- We present a novel guiding approach based on sampling focal points and discuss how it can be implemented efficiently (Section 5)

Our implementation of the proposed algorithm is available at https://github.com/iRath96/focal-guiding.

2 BACKGROUND AND RELATED WORK

Light transport. Simulating light has been a highly active area of research for nearly 40 years. The common goal is to compute the rendering equation [Kajiya 1986], which describes how light propagates through virtual environments. The value of a pixel I_{px} is obtained by integrating the contribution function f_{px} over the space $\mathcal{P}$ of all light transport paths that connect the pixel to a light source [Veach 1998],

$$I_{\text{px}} = \int_{\mathcal{P}} f_{\text{px}}(\bar{x})\,\mathrm{d}\bar{x}. \tag{1}$$

This integral is challenging to compute as it is infinite-dimensional and features many discontinuities, strong peaks, and singularities. It is usually computed through Monte Carlo integration.

Monte Carlo light transport. The Monte Carlo estimator

$$\langle I_{\text{px}}(\bar{x})\rangle = \frac{1}{n_{\text{spp}}}\sum_{i=1}^{n_{\text{spp}}}\frac{f_{\text{px}}(\bar{x}_i)}{p(\bar{x}_i)} \tag{2}$$

generates random path samples $\bar{x}$ distributed with probability density $p(\bar{x})$ and averages their contributions. This estimator converges to the desired integral with increasing sample count n_{spp}. The rate of convergence can be greatly improved through a well-chosen probability density function (PDF) $p(\bar{x})$, which we strive to achieve with our work.

Forward path tracing. Forward path tracing [Kajiya 1986] is the go-to solution for most production renderers at the moment. Paths are constructed through random walks initiated at the camera, and are usually augmented by *next event estimation*, which attempts direct connections to light sources at each intersection. Though enhancements such as *manifold next event estimation* [Hanika et al. 2015; Zeltner et al. 2020] are possible to capture, e.g., caustics and glints, it remains challenging to find all types of light transport with a-priori fixed forward sampling alone.

Bidirectional methods. Bidirectional methods additionally extend paths from the lights into the scene, either by connecting camera- and light subpaths [Lafortune and Willems 1993; Veach and Guibas 1995a], merging them at nearby vertices [Shirley et al. 1995; Jensen 1996; Walter et al. 1997] or a combination of the two [Georgiev et al. 2012; Hachisuka et al. 2012]. Thus, they can efficiently capture focused indirect illumination and caustics in small scenes. Large scenes, like outdoor settings, are problematic because many light paths never find the visible region. Further, bidirectional sampling can struggle when light has to pass through narrow gaps.

Path guiding. Path guiding methods construct adaptive sampling distributions on-the-fly to progressively improve rendering performance. They can operate with forward or bidirectional path tracing and help finding effects that are difficult to sample. Most methods focus on learning local directional distributions [Vorba et al. 2014; Müller et al. 2017], which struggle to reliably sample focal points as the parallax is left unaccounted for. Adjusting the directional distributions for parallax, as done by Ruppert et al. [2020], requires the distance to the focal point to be known beforehand, which is generally only the case for a limited subset of focal effects.

Recently, path space distributions are gaining popularity [Reibold et al. 2018; Schüßler et al. 2022; Li et al. 2022], which are capable of learning the correlation between consecutive path vertices. While this enables them to represent and sample arbitrary focal effects, their inherent local nature hinders them from reliably identifying focal points for all regions of the scene. Oftentimes, individual regions will fail to find some focal effects, resulting in renders that exhibit islands of high variance.

We present a novel path guiding approach that explicitly models focal regions and shares information globally, enabling it to robustly identify and sample focal effects.

MCMC. Markov chain Monte Carlo methods mutate paths such that they eventually achieve a desired distribution [Metropolis et al. 1953; Veach and Guibas 1997; Kelemen et al. 2002]. They are typically based on bidirectional path tracing, as they require good initial

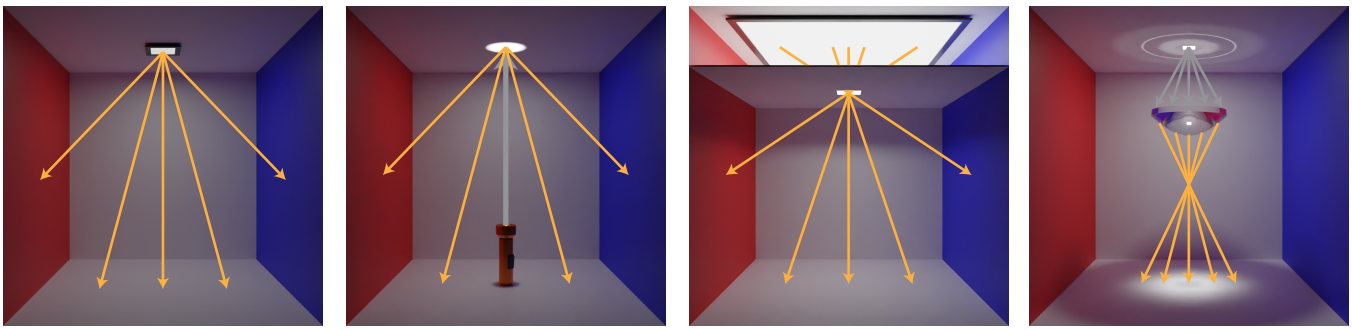

(a) (b) (c) (d)

Figure 2: We identify common causes of focal points. Surface-bound focal points include direct light sources (a) and indirect sources like caustics or spots (b). Focal points also occur in free space, for example when light must pass narrow gaps (c) or when lensing effects shift the apparent position of other focal points (d). Existing algorithms only handle subsets of such effects, often relying on smoothness assumptions or specific types of geometry. Our method does not rely on such assumptions and unifies all focal effects in a single framework.

samples to perform well. A main benefit of MCMC methods is that, through carefully designed mutations and acceptance probabilities, specialized solutions can be found that handle almost any problem. For instance, caustics can be captured through manifold exploration [Jakob and Marschner 2012], small gaps through geometry-aware mutations [Otsu et al. 2018], and so on. Unfortunately, MCMC methods exhibit uneven convergence and correlation artefacts, making them less popular in practice.

3 FOCAL EFFECTS

In the following, we outline the importance of focal effects in light transport and discuss how they are handled by existing techniques. To this end, we group focal effects in different classes and introduce terminology that we will use throughout the remainder of the paper.

Focal points arise where many light paths from different directions converge in a small region of space. Common causes for focal points are illustrated in Figure 2. We classify focal effects in the following three categories:

- *Direct focal points* are small light sources (*light focal points*) or the camera itself (*camera focal points*)
- *Indirect focal points* are caused when objects interact with light, either through diffuse reflection of brightly illuminated spots (*diffusing focal points*) or by forcing light to pass through narrow gaps (*occlusion focal points*)
- *Virtual images* are the apparent position of focal points seen through (potentially multiple) reflection or refraction events

We refer to focal points as *surface-bound* when their position lies on a surface (light, camera, and diffusing focal points) and as *free space* when they occur away from surfaces (occlusion focal points and virtual images).

Direct focal points. These focal effects are the simplest to handle. Since their location is known beforehand, they lend themselves well to explicit sampling, either by starting random walks in their position or by trying to connect to them directly (e.g., through next event estimation). A careful sampling of light sources may be necessary if the scene contains many of them [Walter et al. 2005; Cline et al. 2006; Vévoda et al. 2018; Guo et al. 2020a; Yuksel 2021].

Indirect focal points. More challenging are indirect sources of focal points, the location of which is not known a-priori. Diffusing focal points can be connected to, either by bidirectional methods or by learning their location through path guiding. Occlusion focal points cannot be connected to, as they do not occur at path vertices, but rather at an unknown point on the path in free space. While occlusion of direct light can be sampled explicitly through user intervention [Bitterli et al. 2015; Ogaki 2020], complex visibility of indirect light is still notoriously difficult and only explored well by specialized MCMC mutations [Otsu et al. 2018].

Virtual images. The complexity of sampling virtual images depends greatly on the type of the original focal point and the size and smoothness of the surface producing the virtual image.

Direct virtual images. For virtual images of direct focal points, random walks initiated at the focal point work well if the surface causing the virtual image is large enough to be found by exploration (e.g., a camera pointed at a mirror is handled well by forward path tracing). Connections to virtual images of direct focal points are possible through specialized sampling techniques [Hanika et al. 2015; Zeltner et al. 2020; Jakob and Marschner 2012] (e.g., connecting to glass-enclosed light sources in forward path tracing), but require that the surface producing the virtual image is sufficiently large and smooth.

Indirect virtual images. Virtual images of indirect focal points are more challenging: While diffusing focal points can be handled by some specialized techniques [Li et al. 2022; Jakob and Marschner 2012], images of occlusion focal points remain extremely difficult to find and can only be explored through MCMC methods. Apart from virtual diffusing focal points [Li et al. 2022] and reflections of surface-bound focal points in flat mirrors [Ruppert et al. 2020], virtual images are currently not well explored by path guiding methods, as they are unable to reconstruct the location of free space focal points.

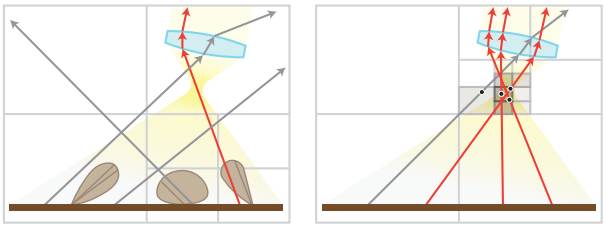

(a) Directional distributions **(b) Our spatial distribution**

Figure 3: (a) Most guiding algorithms partition space and learn directional distributions per region. Due to unaccounted parallax effects, many rays miss the focal point. Some regions never find the focal point during training and cause artefacts. (b) Our method samples the spatial regions themselves, inherently compensating for parallax and sharing information globally.

4 FOCAL GUIDING

We propose a novel form of path guiding, which reliably samples paths involving focal points that otherwise cause excessive variance even for sophisticated algorithms. Like many previous approaches, our guiding augments forward path tracing and iteratively refines its importance sampling density during rendering. Instead of learning local directional distributions directly, our method learns a global distribution over space and samples rays that pass through focal regions (see Figure 3). This has the benefit that focal points can be represented explicitly and shared globally. On the flipside, smoother and locally varying illumination cannot be captured. In the following, we detail the underlying model of our distribution and how focal regions are identified.

Spatial densities. The goal of path guiding in forward path tracing is to reconstruct a desired target density $p_d(\omega_i \mid x)$ from which a new direction ω_i should be sampled at each intersection x. Like many previous works, we drop the dependency on the outgoing direction ω_o, which eases training and reduces the computational cost. Our directional distribution p_d is implicitly defined through a spatial density $p_s(y)$, where the resulting direction ω_i points from the current intersection x towards y

$$\omega_i = \frac{y - x}{\|y - x\|}. \tag{3}$$

The corresponding PDF can be found by integrating along all y that could have produced ω_i

$$p_d(\omega_i \mid x) = \int_0^\infty p_s(x + t\omega_i) t^2 \, \mathrm{d}t, \tag{4}$$

where the t^2 results from the change of variables from y to ω_i. Note that while y lies on the sampled ray, the next vertex of the path is still found through ray tracing and does not need to coincide with y, it can also lie before or behind y.

Discretization. To represent our spatial density, we use adaptive trees inspired by Müller et al. [2017]. To this end, we use voxels ν to partition the space of interest V (e.g., the scene bounding box) into disjunct regions $V_\nu \subset V$. Each region is characterized by constant density $p_s(x) = p_{s,\nu} \forall x \in V_\nu$, leaving us with a piece-wise constant spatial density. All points outside the region of interest have zero sampling density. To sample from this density, we first sample a voxel ν with probability α_ν and then sample a point uniformly within its volume $p(y \mid \nu) = |V_\nu|^{-1}$. The voxel sampling probabilities are linked to the guiding density

$$\alpha_\nu = |V_\nu| \, p_{s,\nu}. \tag{5}$$

Discretized directional density. Our resulting directional density is found by summing the probabilities from all voxels

$$p_d(\omega_i \mid x) = \sum_\nu \alpha_\nu p_\nu(\omega_i \mid x), \tag{6}$$

where we analytically integrate Equation (4) within each voxel

$$p_\nu(\omega_i \mid x) = \begin{cases} \frac{t_1^3 - t_0^3}{3|V_\nu|} & \text{if intersected} \\ 0 & \text{else.} \end{cases} \tag{7}$$

Here, t_0 and t_1 denote the distances of entry and exit of the ray as it intersects the region V_ν of the voxel ν.

Our heuristic density. To find how the focal points are distributed, we will make use of their definition: A continuum of paths that participate in a focal effect must all meet in the same point to make a contribution to the image. We say a path meets in a point if the point lies on any of the lines that constitute the path, even if it lies behind an intersection or before the ray origin. To characterize focal points, we introduce the point-wise contribution integral

$$\mathscr{F}(p) = \frac{1}{n_{\text{px}}} \sum_{\text{px}} \int_{\mathcal{P}(p)} f_{\text{px}}(\bar{x}) \, \mathrm{d}\bar{x}, \tag{8}$$

where $\mathcal{P}(p) \subseteq \mathcal{P}$ is the space of all paths that meet in p. This integral is viewpoint-aware by design—we sample focal points proportional to their image contribution instead of their raw light intensity. This prioritizes focal points that have a high impact on the image and avoids those that are not visible at all. An example of this integral is illustrated in Figure 4a. Since many paths meet in focal points, they cause peaks in the integral $\mathscr{F}(p)$, making it a good initial guess for our guiding density

$$\alpha_\nu \propto \int_{V_\nu} \mathscr{F}(p) \, \mathrm{d}p. \tag{9}$$

Here, we set the selection probability of a voxel ν proportional to the contribution of all paths that pass through the voxel.

Estimating our density. An estimate $\widetilde{\alpha}$ of our density can be obtained simply by logging the contribution of paths as they are traced:

$$\widetilde{\alpha}_\nu \propto \sum_{\text{px}}^{n_{\text{px}}} \sum_{i=1}^{n_{\text{spp}}} \sum_{j=2}^{n(\bar{x}_i)} \begin{cases} (t_1 - t_0) \frac{f_{\text{px}}(\bar{x}_i)}{p(\bar{x}_i)} & \text{if intersected} \\ 0 & \text{else,} \end{cases} \tag{10}$$

where we collect the contribution from all path segments j that have been sampled. The values t_0 and t_1 denote the entry and exit distances of the ray $x_i^{j-1} \to x_i^j$ in voxel ν, and the multiplication with $t_1 - t_0$ results from the integral over the region of the voxel.

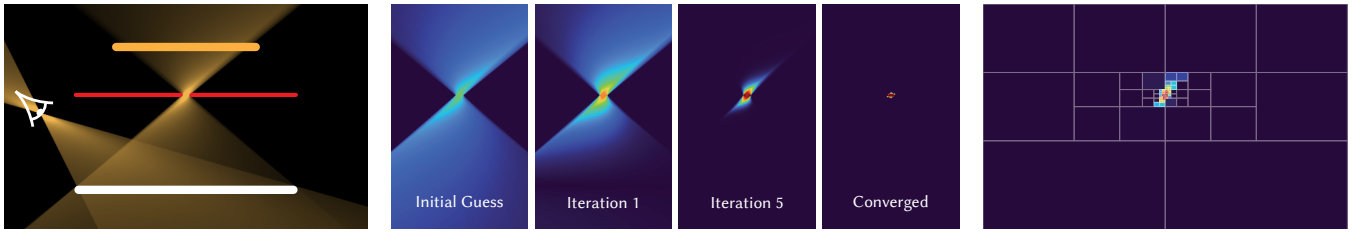

(a) Point-wise contribution integral **(b) Our iterative narrowing scheme** **(c) Adaptive octrees**

Figure 4: (a) We show the point-wise contribution integral of a diffuse surface (white line) illuminated by a large area light (orange line) through a narrow gap (occluder in red). The camera- and occlusion focal points cause peaks in the integral, as many paths meet in the same points. (b) Excluding segments that cannot be guided (e.g., camera segments), this serves as initial guess for our guiding density, which is then iteratively refined to avoid sampling points that are only relevant to subsets of the focal effect. (c) We represent the guiding density using adaptive octrees.

Spurious focal points. Focal effects are seldom singular points in practice, but rather small regions of finite extent. Unfortunately, this creates smearing in the point-wise contribution integral. Consider Figure 4a: The region within the narrow gap is sufficient to explore the full focal effect, as all paths need to cross it. But the points above and below the narrow gap also have a high value, as subsets of the paths pass them. We refer to these points as *spurious* focal points. They are detrimental to the quality of our sampling as they only benefit subsets of the focal effect, but take away probability mass that could be invested in true focal points benefiting all paths.

Iterative narrowing. To narrow our distribution down to true focal points, we introduce a scheme that we refer to as *iterative narrowing*. Given a previous distribution estimate $\widetilde{\alpha}$, we obtain a narrowed distribution $\widetilde{\alpha}^N$ by weighting the contributions of paths:

$$\widetilde{\alpha}_v^N \propto \sum_{\text{px}}^{n_{\text{px}}} \sum_{i=1}^{n_{\text{spp}}} \sum_{j=2}^{n(\bar{x}_i)} w_v(x_i^{j-1}, x_i^{j-1} \to x_i^j) \frac{f_{\text{px}}(\bar{x}_i)}{p(\bar{x}_i)}, \tag{11}$$

where the weight w_v is given by

$$w_v(x, \omega_i) = \frac{\widetilde{\alpha}_v p_v(\omega_i \mid x)}{\sum_{v'} \widetilde{\alpha}_{v'} p_{v'}(\omega_i \mid x)}. \tag{12}$$

Instead of equally contributing to all intersected voxels, iterative narrowing weights the contribution to voxels by how likely the voxel samples the path segment. Since spurious focal points only capture subsets of focal effects, fewer paths contribute to them than to the true focal point, leading to lower selection probabilities. Iterative narrowing then further lowers their value by assigning lower weights to the contribution from path segments. Applying this repeatedly effectively eliminates spurious focal points (Fig. 4b).

Diverging focal points. So far, we have glanced over the fact that focal points can also lie behind the surface we are currently sampling from. Supporting this is straightforward. We introduce additional voxels v^- that also partition the region of interest, but produce directions that point *away from* the sampled point y^-:

$$\omega_i = -\frac{y^- - x}{\|y^- - x\|}. \tag{13}$$

These voxels are considered intersected (for logging contributions or PDF computations) if they lie in negative ray direction.

5 IMPLEMENTATION

We implement our method as a guided forward path tracer in the Mitsuba renderer [Jakob 2010]. Similar to previous guiding works, our algorithm proceeds in iterations. Each iteration refines the density from the previous iteration. The last iteration then renders the final image. In the following, we detail the exact iteration schedule, data structure and parameters used by our approach.

Schedule. While it is possible to continue training throughout the render, there is usually a point where the diminishing returns of additional training no longer justify the overhead of logging contributions and updating data structures. Determining the optimal amount of training is still an open problem in path guiding, hence we follow the compromise of Müller et al. [2017] and split the total budget into 50% training and 50% final render. The training phase is further subdivided into iterations, which each train an improved density by sampling from the density learned in the previous iteration. We find that iterations of equal length [Ruppert et al. 2020] perform better than iterations of increasing duration [Müller et al. 2017], as they accelerate learning of rare effects and lend themselves nicely to our iterative narrowing scheme. The training budget is split equally into n_{iter} iterations, for which we find a choice of $n_{\text{iter}} = 15$ to work well across all tested scenes.

Burn-in period. Challenging focal effects can take several iterations before they are discovered. Before enabling iterative narrowing (Section 4), we need to be confident that all relevant focal points have been identified. Otherwise, once a valid path is finally discovered, its focal region might already be considered spurious due to its previously low value. While iterative narrowing can eventually recover from such cases, it can take several iterations for a low-scoring region to be considered focal again. To this end, we only enable iterative narrowing in the last five training iterations, which we find sufficient to dismiss most spurious focal points (see Figure 4b). A more sophisticated remedy to handling rare events could rely on Bayesian priors [Vévoda et al. 2018; Dodik et al. 2021], which we leave as future work.

Data structure. We use an adaptive octree to represent our guiding density (see Figure 4c). The region of the root node is set to the bounding box of the scene. Inspired by Müller et al. [2017], a node

is split when its selection probability α_v exceeds a given splitting threshold γ. This threshold is the main parameter of our technique: A lower value yields higher resolution, but requires more samples to fit and increases the cost of PDF computations. We observe that a value of $\gamma = 10^{-3}$ works well across all our test scenes, and a detailed analysis of this parameter is provided in Section 6.

Pruning. The increased resolution of splitting does not always justify the penalty on traversal cost, in particular when there is little variation within a node. To mitigate this, we use a simple pruning heuristic (see Figure 6): We collapse each node where the highest leaf density does not exceed twice the average density of the node. We notice that a fine subdivision is still beneficial for learning, so we prune only after the last training iteration is completed.

Sampling. We obtain spatial samples $y \sim p_s$ using hierarchical sample warping [McCool and Harwood 1997]. The direction vector ω_i follows from Eq. (3). For the probability density (Eq. (6)), it suffices to sum up the voxels intersected by the ray, since all other voxels have zero probability density of producing the sample. For this, we use the traversal algorithm by Revelles et al. [2000].

Multiple importance sampling. While our guiding density excels at sampling focal effects, its performance on other light transport can be poor. It is therefore advisable to combine it with other sampling strategies using multiple importance sampling [Veach and Guibas 1995b]. Like Müller et al. [2017], we perform mixture sampling of BSDF and guiding at a fixed ratio of $\lambda_B = 0.5$. Similarly, we also perform next event estimation. No guiding is performed on specular surfaces, and accordingly we do not log segments that cannot be guided.

MIS compensation. Path guiding works best when it learns to augment rather than replace other sampling strategies. This is known as MIS compensation [Karlík et al. 2019]: We want guiding to only learn effects not handled well by other sampling strategies in our MIS combination. To compensate for BSDF sampling, we use the scheme proposed by Rath et al. [2020], which—similar to our iterative narrowing scheme—iteratively weights the guiding distribution by how likely it is to sample an effect. Applied to our iterative narrowing scheme (Equation (11)), our weights become

$$w_v(x, \omega_i) = \frac{(1-\lambda_B)\alpha_v p_v(\omega_i \mid x)}{\lambda_B p_B(\omega_i \mid x, \omega_o) + (1-\lambda_B)\sum_{v'} \alpha_{v'} p_{v'}(\omega_i \mid x)}, \quad (14)$$

where p_B is the density of BSDF sampling and λ_B its selection probability. We additionally multiply the contribution by the next event estimation MIS weight, as done by Ruppert et al. [2020].

6 EVALUATION

We compare the performance of our method for various focal effects against the following methods (Figure 5):

- PT: Forward path tracing with next event estimation [Kajiya 1986]
- MEMLT: Manifold-exploration MLT [Jakob and Marschner 2012]
- MCVCM: Metropolised vertex connection and merging [Šik et al. 2016]
- PAVMM: Parallax-aware mixtures for path guiding [Ruppert et al. 2020]

All techniques are implemented in Mitsuba [Jakob 2010]. PAVMM and our approach use half of the time budget for learning. Russian roulette is an orthogonal problem in path guiding, as guided paths often have low throughput until they reach a light source [Vorba and Křivánek 2016; Rath et al. 2022], hence we disable Russian roulette in all techniques for comparability. In addition to our renders, we demonstrate the benefits of our approach over path space guiding [Reibold et al. 2018] in a 2-D experiment.

All evaluations were run on a 16-core AMD Ryzen™ 9 3950X processor with 64 GB of memory. All renders are rendered for approximately three minutes. We provide the mean squared error (MSE), for which we clamp outliers at the 99.9% percentile to arrive at a stable value, and ꟻLIP metric [Andersson et al. 2021].

Splitting threshold. We begin by analyzing the impact of the main parameter of our approach: the splitting threshold γ, which determines the resolution of our adaptive octree. We evaluate its impact on computational overhead and memory usage, averaged over three-minute renders of our five test scenes, in Figure 7. Finer resolutions result in more accurate sampling, at the expense of higher sample cost. The best efficiency across all scenes is achieved at a value of $\gamma = 10^{-3}$, but the performance of our algorithm is not particularly sensitive to its exact value. At this threshold, samples are roughly twice as expensive as unguided sampling, and the memory overhead of our method amounts to a mere 76 KiB.

6.1 Camera Obscura

The Camera Obscura scene features an occlusion focal point in form of a pinhole in the wall, which projects the strongly illuminated statue on the left onto the canvas on the right. Path tracing struggles with this scene, as uninformed sampling is unlikely to find the pinhole. MEMLT fails to find a seed path that passes the pinhole and hence completely misses the projection of the statue. The metropolised light tracing of MCVCM fails for the same reason, but as with PT, a few rare eye paths manage to pass through the pinhole. PAVMM finds the pinhole, but only in small patches of the scene. Parallax compensation is of little help here, as it aims rays at surface points of the statue rather than the pinhole that is causing the focal effect. As evident from our density visualization, our approach identifies the pinhole as focal region and systematically constructs paths that pass through it, resulting in significantly faster convergence than previous state-of-the-art algorithms.

6.2 Dining Room

The parabolic lampshades in our version of the Dining Room scene create a virtual image of a tiny light source pointing upwards. Path tracing struggles to find paths that pass the focal region. Based on bidirectional path tracing, both MEMLT and MCVCM successfully construct light subpaths that pass through the virtual light focal point. The specular jug on the table, however, introduces a specular-diffuse-specular (SDS) path that prohibits the connection of eye and light subpaths. MCVCM solves this issue through vertex merging, introducing bias in the process. MEMLT can reliably explore SDS paths, but again fails to find a valid subpath to initiate the exploration. Without light tracing (PT, PAVMM, and our approach), paths that pass through the virtual image need to be found explicitly, which only our method can sample systematically.

6.3 Funky Living Room

The Funky Living Room scene includes a glossy disco ball illuminated by five spot lights. Path tracing struggles to find the correct facet on the disco ball connecting the path to a spot light. MEMLT produces noticeable artefacts due to uneven convergence, while MCVCM exhibits density estimation artefacts. PAVMM learns to guide towards the disco ball, but mispredicts the parallax distance of the spotlights. Our approach learns that it is best to aim rays at the interior of the disco ball, in which all spot lights converge.

6.4 Modern Hall

The Modern Hall scene is illuminated by four tiny area lights enclosed in translucent boxes, which constitute diffusing focal points. This effectively prevents next event estimation and makes the technique difficult for forward unidirectional techniques. For bidirectional methods like MEMLT and MCVCM, however, this scene is fairly simple to render. As far as unidirectional methods go, PAVMM performs exceptionally well in this scene, as it can correctly compensate the parallax to the light source given the surface-bound nature of the focal point. While our method significantly improves over PT, it falls short of the performance of PAVMMs as it cannot account for occlusion or the falloff of light over distance.

6.5 Modern Living Room

The Modern Living Room is a failure case of our approach. This scene is illuminated by a large light source just out of frame, which is handled very well by next event estimation. The path tracer remains unbeaten in this scene, as the overhead of more sophisticated methods does not pay off in this simple scene. Our approach faces an additional problem: Since it cannot represent smoothly varying illumination, our guiding degrades the quality of sampling in this scene. This is expected, as our guiding is specifically tailored to handled focal effects, which are absent from this scene.

6.6 Comparison with Path Space Guiding

In Figure 8 we compare the behavior of path space guiding [Reibold et al. 2018] and our approach in a simplified 2-D experiment. Paths originating at a plane need to pass through an obstacle with three narrow gaps to reach a large area light at the top. Both methods were trained for 6 iterations with 4096 samples each.

In path space guiding (Fig. 8a), directions are sampled according to statistics over earlier paths in the vicinity of the current vertex. This makes the method inherently local, as individual focal points must be re-discovered for each region of the floor, which is prone to overlooking parts of focal effects due to the sparsity of constructing valid focal paths through exploration. In our example, the right side of the floor has not found sufficient guide paths to reliably sample the center focal point, which manifests itself in residual outliers and islands of high variance.

Our approach shares information globally (Fig. 8b), which allows the entire floor to benefit from robust sampling once a focal region has been identified. The price for this is that local effects are handled poorly by our approach: Occlusion, light falloff over distance, and anisotropic effects are handled much better by local methods, which can learn different distributions for different regions of space.

7 DISCUSSION AND FUTURE WORK

We have shown a simple method to identify and sample focal points in a forward path tracer. We believe that there are many promising ways in which our work could be continued:

Augmenting other techniques. While our technique excels at focal effects, its performance on general light transport is limited. We believe its best use is to augment other techniques. For instance, it could identify light portals [Bitterli et al. 2015; Ogaki 2020] or learn distances for parallax compensation [Ruppert et al. 2020].

Region of interest. We currently assume that focal points lie within the bounding box of the scene, which is not necessarily true for virtual images. It could be interesting to automatically extend the region of interest when necessary. An alternative approach could be to use a representation that covers all of space, potentially with reduced precision for increasing distances.

Handling local effects. Our sampling scheme assumes that focal points contribute equally throughout the scene, ignoring both visibility and the falloff of light intensity over distance. A promising remedy lending itself well to our approach is the voxel-based learning of visibility by Guo et al. [2020b].

Alternative representations. Our implementation relies on adaptive trees [Müller et al. 2017] to represent the focal guiding density. An interesting alternative would be to fit parametric mixture models [Vorba et al. 2014; Ruppert et al. 2020; Dodik et al. 2021] or neural networks [Müller et al. 2019], which could enable product sampling (to handle local effects) or reduce sample overhead.

Variance-aware target density. Crafting densities that explicitly minimize image variance can greatly increase the rate of convergence [Rath et al. 2020]. Our guiding scheme learns selection probabilities of a mixture density, in which each voxel is a unique sampling strategy. Unfortunately, finding variance-optimal selection probabilities remains an unsolved problem [Lu et al. 2013; Sbert et al. 2018]. Stochastic gradient descent could be a promising tool to optimize our voxel selection probabilities [Müller et al. 2017].

8 CONCLUSION

Light interacts with our world in fascinating ways. While not all of them can be simulated yet, this work takes us one step closer by addressing a class of effects that has received little attention so far: *focal effects*. We have classified common causes of focal effects and surveyed how they are handled by existing families of algorithms. To render them robustly, we introduced a novel form of guiding, based on locating and sampling focal regions. Our technique unifies all focal effects in a single framework and can render effects that previous state-of-the-art techniques are unable to handle.

ACKNOWLEDGMENTS

This project has received funding from Velux Stiftung project 1350. We thank Florian Reibold, Pascal Grittmann and Sebastian Herholz for insightful discussions, as well as our anonymous reviewers for their valuable feedback. Our test scenes are based on the works of Bitterli [2016], Wig42 (Modern Living Room, Dining Room), Jay-Artist (Living Room), NewSee2l035 (Modern Hall), Statens Museum for Kunst, and Polyhaven (Camera Obscura).

REFERENCES

Pontus Andersson, Jim Nilsson, Peter Shirley, and Tomas Akenine-Möller. 2021. Visualizing the Error in Rendered High Dynamic Range Images. In *Eurographics Short Papers*. https://doi.org/10.2312/egs.20211015

Benedikt Bitterli. 2016. Rendering resources. https://benedikt-bitterli.me/resources/

Benedikt Bitterli, Jan Novák, and Wojciech Jarosz. 2015. Portal-masked environment map sampling. In *Computer Graphics Forum*, Vol. 34. Wiley Online Library, 13–19.

David Cline, Parris K. Egbert, Justin F. Talbot, and David L. Cardon. 2006. Two Stage Importance Sampling for Direct Lighting. In *Proceedings of the 17th Eurographics Conference on Rendering Techniques* (Nicosia, Cyprus) *(EGSR '06)*. Eurographics Association, Goslar, DEU, 103–113.

Ana Dodik, Marios Papas, C. Öztireli, and T. Müller. 2021. Path Guiding Using Spatio-Directional Mixture Models. *Comput. Graph. Forum* (2021). https://doi.org/10.1111/cgf.14428

Iliyan Georgiev, Jaroslav Krivánek, Tomas Davidovic, and Philipp Slusallek. 2012. Light transport simulation with vertex connection and merging. *ACM Trans. Graph.* 31, 6 (2012), 192–1.

Jerry Jinfeng Guo, Martin Eisemann, and Elmar Eisemann. 2020a. Next Event Estimation++: Visibility Mapping for Efficient Light Transport Simulation. *Computer Graphics Forum* 39, 7 (2020), 205–217. https://doi.org/10.1111/cgf.14138 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14138

Jerry Jinfeng Guo, Martin Eisemann, and Elmar Eisemann. 2020b. Next Event Estimation++: Visibility Mapping for Efficient Light Transport Simulation. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 205–217.

Toshiya Hachisuka, Jacopo Pantaleoni, and Henrik Wann Jensen. 2012. A path space extension for robust light transport simulation. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 1–10.

Johannes Hanika, Marc Droske, and Luca Fascione. 2015. Manifold Next Event Estimation. *Comput. Graph. Forum* 34, 4 (jul 2015), 87–97.

Wenzel Jakob. 2010. *Mitsuba renderer*. https://www.mitsuba-renderer.org

Wenzel Jakob and Steve Marschner. 2012. Manifold exploration: A markov chain monte carlo technique for rendering scenes with difficult specular transport. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–13.

Henrik Wann Jensen. 1996. Global illumination using photon maps. In *Eurographics workshop on Rendering techniques*. Springer, 21–30.

James T. Kajiya. 1986. The Rendering Equation. *SIGGRAPH Comput. Graph.* 20, 4 (Aug. 1986), 143–150.

Ondřej Karlík, Martin Šik, Petr Vévoda, Tomáš Skřivan, and Jaroslav Křivánek. 2019. MIS compensation: optimizing sampling techniques in multiple importance sampling. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–12.

Csaba Kelemen, László Szirmay-Kalos, György Antal, and Ferenc Csonka. 2002. A simple and robust mutation strategy for the metropolis light transport algorithm. In *Computer Graphics Forum*, Vol. 21. Wiley Online Library, 531–540.

Eric P Lafortune and Yves D Willems. 1993. Bi-directional path tracing. (1993).

He Li, Beibei Wang, Changhe Tu, Kun Xu, Nicolas Holzschuch, and Ling-Qi Yan. 2022. Unbiased Caustics Rendering Guided by Representative Specular Paths. In *SIGGRAPH Asia 2022 Conference Papers* (Daegu, Republic of Korea) *(SA '22)*. Association for Computing Machinery, New York, NY, USA, Article 41, 8 pages. https://doi.org/10.1145/3550469.3555381

Heqi Lu, Romain Pacanowski, and Xavier Granier. 2013. Second-Order Approximation for Variance Reduction in Multiple Importance Sampling. In *Computer Graphics Forum*, Vol. 32. Wiley Online Library, 131–136.

Michael D McCool and Peter K Harwood. 1997. Probability trees. In *Graphics Interface*, Vol. 97. 37–46.

Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. 1953. Equation of state calculations by fast computing machines. *The journal of chemical physics* 21, 6 (1953), 1087–1092.

Thomas Müller, Markus Gross, and Jan Novák. 2017. Practical path guiding for efficient light-transport simulation. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 91–100.

Thomas Müller, Brian McWilliams, Fabrice Rousselle, Markus Gross, and Jan Novák. 2019. Neural importance sampling. *ACM Transactions on Graphics (TOG)* 38, 5 (2019), 1–19.

Shinji Ogaki. 2020. Generalized Light Portals. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 3, 2 (2020), 1–19.

Hisanari Otsu, Johannes Hanika, Toshiya Hachisuka, and Carsten Dachsbacher. 2018. Geometry-Aware Metropolis Light Transport. *ACM Transactions on Graphics (Proc. of SIGGRAPH Asia)* 37, 6, Article 278 (2018), 278:1–278:11 pages.

Alexander Rath, Pascal Grittmann, Sebastian Herholz, Petr Vévoda, Philipp Slusallek, and Jaroslav Křivánek. 2020. Variance-aware path guiding. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 151–1.

Alexander Rath, Pascal Grittmann, Sebastian Herholz, Philippe Weier, and Philipp Slusallek. 2022. EARS: efficiency-aware russian roulette and splitting. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–14.

Florian Reibold, Johannes Hanika, Alisa Jung, and Carsten Dachsbacher. 2018. Selective Guided Sampling with Complete Light Transport Paths. *ACM Trans. Graph.* 37, 6, Article 223 (dec 2018), 14 pages. https://doi.org/10.1145/3272127.3275030

Jorge Revelles, Carlos Urena, and Miguel Lastra. 2000. An efficient parametric algorithm for octree traversal. (2000).

Lukas Ruppert, Sebastian Herholz, and Hendrik P. A. Lensch. 2020. Robust Fitting of Parallax-Aware Mixtures for Path Guiding. *ACM Trans. Graph.* 39, 4, Article 147 (aug 2020), 15 pages. https://doi.org/10.1145/3386569.3392421

Mateu Sbert, Vlastimil Havran, and Laszlo Szirmay-Kalos. 2018. Multiple importance sampling revisited: breaking the bounds. *EURASIP Journal on Advances in Signal Processing* 2018, 1 (2018), 1–15.

Vincent Schüßler, Johannes Hanika, Alisa Jung, and Carsten Dachsbacher. 2022. Path Guiding with Vertex Triplet Distributions. *Computer Graphics Forum (Proceedings of Eurographics Symposium on Rendering)* 41, 4 (2022). https://doi.org/10.1111/cgf.14582

Peter Shirley, Bretton Wade, Philip M Hubbard, David Zareski, Bruce Walter, and Donald P Greenberg. 1995. Global illumination via density-estimation. In *Eurographics Workshop on Rendering Techniques*. Springer, 219–230.

Eric Veach. 1998. *Robust Monte Carlo methods for light transport simulation*. Stanford University.

Eric Veach and Leonidas Guibas. 1995a. Bidirectional estimators for light transport. In *Photorealistic Rendering Techniques*. Springer, 145–167.

Eric Veach and Leonidas J Guibas. 1995b. Optimally combining sampling techniques for Monte Carlo rendering. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. 419–428.

Eric Veach and Leonidas J Guibas. 1997. Metropolis light transport. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 65–76.

Petr Vévoda, Ivo Kondapaneni, and Jaroslav Křivánek. 2018. Bayesian online regression for adaptive direct illumination sampling. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–12.

Jiří Vorba, Ondřej Karlík, Martin Šik, Tobias Ritschel, and Jaroslav Křivánek. 2014. On-line learning of parametric mixture models for light transport simulation. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–11.

Jiří Vorba and Jaroslav Křivánek. 2016. Adjoint-driven Russian roulette and splitting in light transport simulation. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–11.

Martin Šik, Hisanari Otsu, Toshiya Hachisuka, and Jaroslav Křivánek. 2016. Robust Light Transport Simulation via Metropolised Bidirectional Estimators. *ACM Trans. Graph.* 35, 6, Article 245 (dec 2016), 12 pages. https://doi.org/10.1145/2980179.2982411

Bruce Walter, Sebastian Fernandez, Adam Arbree, Kavita Bala, Michael Donikian, and Donald P Greenberg. 2005. Lightcuts: a scalable approach to illumination. In *ACM SIGGRAPH 2005 Papers*. 1098–1107.

Bruce Walter, Philip M Hubbard, Peter Shirley, and Donald P Greenberg. 1997. Global illumination using local linear density estimation. *ACM Transactions on Graphics (TOG)* 16, 3 (1997), 217–259.

Cem Yuksel. 2021. Stochastic Lightcuts for Sampling Many Lights. *IEEE Transactions on Visualization and Computer Graphics* 27, 10 (2021), 4049–4059. https://doi.org/10.1109/TVCG.2020.3001271

Tizian Zeltner, Iliyan Georgiev, and Wenzel Jakob. 2020. Specular Manifold Sampling for Rendering High-Frequency Caustics and Glints. *ACM Trans. Graph.* 39, 4, Article 149 (aug 2020), 15 pages. https://doi.org/10.1145/3386569.3392408

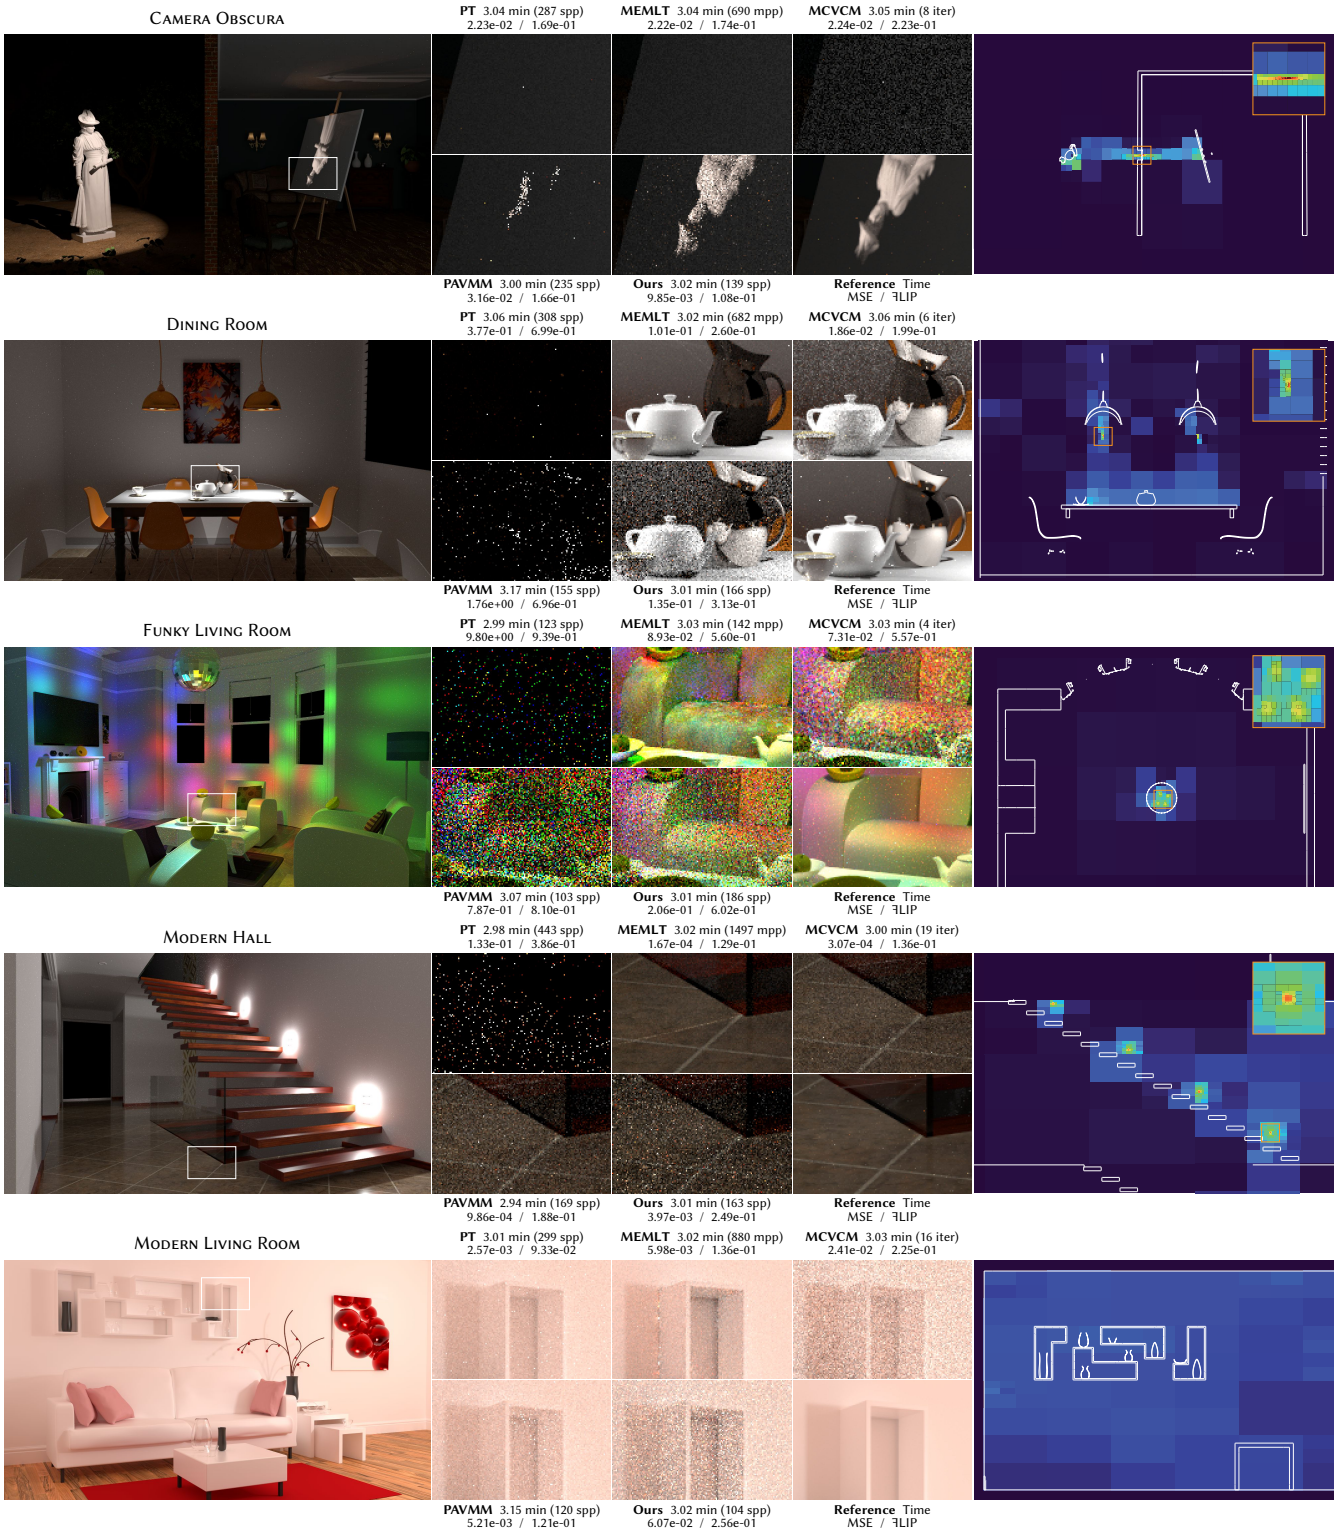

Figure 5: We compare the performance of our approach against previous works on five challenging test scenes. Each method renders for approximately three minutes. The first four scenes demonstrate various focal effects, while the last scene exhibits primarily afocal light transport. A slice through our spatial guiding density is displayed on the right side, overlaid by the corresponding slice of the scene geometry for orientation (white lines).

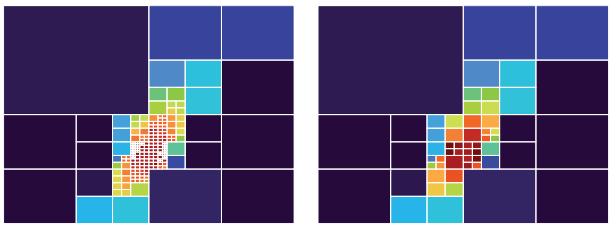

(a) Before pruning (b) After pruning

Figure 6: After the last training iteration, we collapse octree nodes that have little variation. This reduces the cost of the traversal required to compute probability densities.

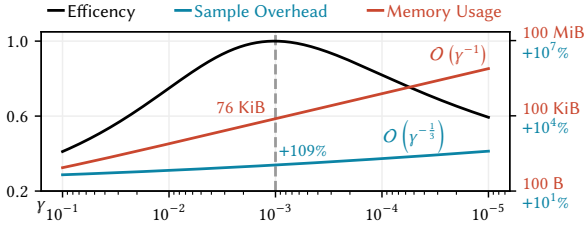

Figure 7: We evaluate our splitting threshold γ. Finer resolutions (i.e., low values of γ) yield more accurate sampling, but incur higher memory usage and sample overhead relative to an unguided path tracer. The efficiency (MSE error of equal-time renders) peaks at $\gamma = 10^{-3}$.

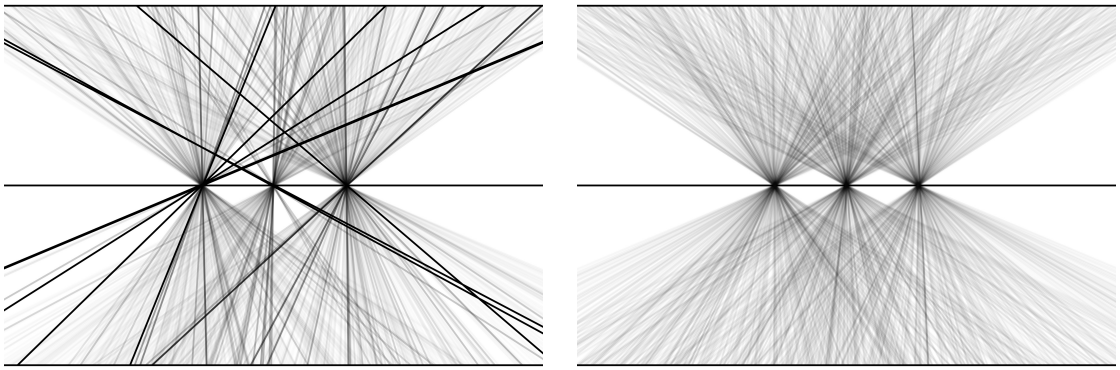

(a) Sampling complete light paths [Reibold et al. 2018] (b) Sampling a global spatial density (our approach)

Figure 8: We compare the performance of local versus global sampling distributions in a simplified 2-D experiment. Paths originating at the floor need to pass an obstacle with three pinholes to reach the light source at the top. The opacity of paths is proportional to the value of its estimator, i.e., darker paths represent outliers which are sampled insufficiently. While both approaches can reliably sample focal points once they are identified, local methods (a) are prone to overlooking parts of focal effects, such as the contribution from the center focal point on the right side of the floor in this example. This results in uneven sampling, which manifests itself as islands of high variance in the rendered image. Global distributions like ours (b) do not suffer from this issue: Once focal points are identified, all regions directly exploit this information. This results in more even sampling, eliminating residual outliers.