

# A practical path guiding method for participating media

Hong Deng<sup>1</sup>, Beibei Wang<sup>1</sup> (✉), Rui Wang<sup>2</sup>, and Nicolas Holzschuch<sup>3</sup>

© The Author(s) 2020.

**Abstract** Rendering translucent materials is costly: light transport algorithms need to simulate a large number of scattering events inside the material before reaching convergence. The cost is especially high for materials with a large albedo or a small mean-free-path, where higher-order scattering effects dominate. In simple terms, the paths get lost in the medium. Path guiding has been proposed for surface rendering to make convergence faster by guiding the sampling process. In this paper, we introduce a path guiding solution for translucent materials. We learn an adaptive approximate representation of the radiance distribution in the volume and use it to sample the scattering direction, combining it with phase function sampling by resampled importance sampling. The proposed method significantly improves the performance of light transport simulation in participating media, especially for small lights and media with refractive boundaries. Our method can handle any homogeneous participating medium, with high or low scattering, with high or low absorption, and from isotropic to highly anisotropic.

**Keywords** Monte Carlo; ray tracing; path guiding; volume light transport; participating media; rendering

## 1 Introduction

Computing illumination simulation in scenes with participating media is still a costly process, even for algorithms specifically designed to handle them.

To render high-albedo materials, such as milk or skin, we need to simulate many scattering events before convergence. Path tracing, as a simple and powerful method, has recently been used to simulate these materials in industry. However, it takes a very long time to converge; it is even worse for small light sources. The main reason is that sampling the out-scattering direction only depends on importance sampling of the phase function, and has no knowledge of the location of the light. This results in many low contribution paths being followed.

Path guiding has been recently introduced for surface rendering [1–4]. The common goal of these works is to find an “optimal” distribution that can approximate the actual path integral to make convergence faster. In these prior works, the incoming radiance distribution of some samples is learned, and used by combining it with the bidirectional reflection distribution function (BRDF) with multiple importance sampling [2] or product importance sampling [3]. Besides path-space based path guiding methods, several works have focused on the primary sample space [5–7]; some of these works use machine learning for path guiding. Path guiding makes convergence faster than unaltered path tracing.

Zero-variance methods [8, 9] have been proposed for participating media, but only for isotropic homogeneous media. Recently, this approach has been further improved [10], using path guiding for arbitrary homogeneous media based on zero-variance random walk theory.

In this paper, we present a different path guiding method for arbitrary homogeneous participating media. As in Müller et al. [2], we use an SD-tree to represent the incoming radiance distribution in the scene. However, the distribution is sampled inside the volume rather than on the surface, and also learned in a bidirectional manner, considering difficult paths due

1 Nanjing University of Science and Technology, Nanjing 210094, China. E-mail: H. Deng, hong.deng.bot@njust.edu.cn; B. Wang, beibei.wang@njust.edu.cn (✉).

2 State Key Lab of CAD&CG, Zhejiang University, Hangzhou 310058, China. E-mail: rwang@cad.zju.edu.cn.

3 Univ. Grenoble-Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France. E-mail: Nicolas.Holzschuch@inria.fr.

Manuscript received: 2020-02-06; accepted: 2020-02-19

to the participating media boundaries. We propose a resampling method, combining this incoming radiance distribution with the phase function to better sample the product of two high frequency functions. Our resampling method only requires a couple of samples to approximate the outgoing radiance distribution. For more efficiency, we introduce a selective sampling strategy, choosing a sampling method based on the phase function. We use path guiding to select the direction in the medium. Our method results in faster convergence than unaltered path tracing, greatly reducing noise in the output images without impact on accuracy. In conjunction with our method, denoising methods can further remove the noise. Our method produces a better input for denoising compared to classical path tracing. To summarize, we make the following technical contributions:

- a modified learning method for the SD-tree used to represent incoming radiance distribution in the volume;
- a resampled importance sampling method to sample the product of the incoming radiance distribution function and the phase function;
- a selective sampling strategy to choose the sampling method based on phase function anisotropy of the medium.

In the next section, we review previous work on rendering participating media. In Section 4, we present path guiding in participating media. We present our results, compare them with previous work, and analyze performance in Section 5, and conclude in Section 6.

## 2 Previous work

### 2.1 Participating media rendering

#### 2.1.1 Path tracing

Simulating light transport in participating media requires simulating multiple scattering events in the medium. Early approaches used Monte Carlo integration to compute illumination [11–13].

In recent works, Georgiev et al. [14] used the product of geometry and scattering terms for importance sampling in rendering participating media. Unlike our method, this method is limited to low-order scattering, while we handle both high-order and low-order scattering.

#### 2.1.2 Photon mapping

Chandrasekhar [15] introduced the *radiative transfer equation*, describing radiation transport in participating media. Jensen and Christensen [16] applied this equation to light transport, and presented an algorithm based on photon mapping. Jarosz et al. [17–19] extended this algorithm using beams instead of photons for faster computation, with less noise. Křivánek et al. [20] automatically selected between beams, points, and paths in participating media using multiple importance sampling. Bitterli and Jarosz [21] further extended the idea by tracing photon planes and volumes.

Photon-mapping-based methods can provide high quality simulations of light transport in participating media, but they are usually biased. In contrast, our method provides an unbiased result.

#### 2.1.3 Radiance caching

Jarosz et al. [22] introduced radiance caching for participating media and decreased the rendering cost, by storing radiance at sample places in an octree. Marco et al. [23] improved this work by considering second-order occlusion. Our method also improves the performance, but in a different way. Furthermore, our method is unbiased, while theirs is biased.

### 2.2 Path guiding

Several path guiding methods have been proposed to drive importance sampling of outgoing ray directions based on learned information. Most focus on surface-to-surface interactions; one work presented path tracing for participating media.

#### 2.2.1 Path space

Jensen [24] proposed placing regular bins to collect photons traced from light. Histograms reconstructed from these bins guide sampling directions when tracing paths from the camera. Vorba et al. [1] learned the incoming radiance distribution at sampled points in the scene by shooting photons, and used the trained distribution to guide direction sampling in path tracing. The learned incoming radiance distribution is represented by a Gaussian mixture model. Müller et al. [2] used an adaptive binary tree in the spatial domain and a quadtree in the angular domain to store the incoming radiance distribution. They used multiple importance sampling of the BRDF and the learned distribution. Herholz et al. [3] used the product of the trained incoming light distribution

and the BRDF to sample the outgoing direction, resulting in higher sampling quality, at the cost of the product operator.

Reibold et al. [4] proposed using path guiding only for paths with high variance or difficult paths, and regular path tracing for other paths, as path guiding is more expensive compared to regular path tracing.

Deep learning has also been used in path guiding. Müller et al. [7] proposed a deep neural network model to determine the probability density function of samples. The learned model was leveraged for sampling the ray outgoing direction. The advantage of this work is its independence from the scene and local point parameters (e.g., including the textured BRDF). However, this method is biased and the sampling is very expensive.

### 2.2.2 Primary sample space

Other researchers have worked on path guiding in primary sample space (PSS) instead of path space. PSS is the space of random numbers that are used to generate paths. PSS has been widely used in Metropolis light transport algorithms [25]. Guo et al. [5] proposed a path guiding method in PSS. Random numbers are sampled from a learned structure, a  $k$ -D tree, to generate the paths. Müller et al. [7] and Zheng and Zwicker [6] provided PSS-based methods for path sampling using deep learning.

One problem of PSS-based methods is that they cannot handle long paths very well, as the random number count depends on the path length. We do not use them for participating media as these materials result in very long paths, due to the many scattering events.

## 2.3 Zero-variance based methods

Unlike path space and primary sample space path guiding methods, which use data-driven models or neural networks to guide sampling, zero-variance random walks rely on analytical models for sampling.

Křivánek and d'Eon [8] used zero-variance random walks in dense, isotropic, and highly scattering media, by fitting a slab to the bounding geometry, which helps to guide the random walks back to the surface. Later, Meng et al. [9] improved this method to handle thin geometry or back-lit cases.

Herholz et al. [10] used zero-variance path sampling theory and exploited the vMF model to represent the involved radiance distributions to guide all volume sampling including distance, direction, Russian

roulette, and splitting. Unlike their work, we use a different representation and only focus on direction sampling.

## 3 Background

### 3.1 Radiative transfer equation

We consider a scene containing objects with translucent materials. Each of these is assumed to be made of a homogeneous material, with index of refraction  $\eta$ , scattering coefficient  $\sigma_s$ , absorption coefficient  $\sigma_a$ , and phase function  $p(\vec{\omega}, \vec{\omega}_t)$ . We denote by  $\ell$  the mean-free path (mfp) inside the material, with  $1/\ell = \sigma_t = \sigma_s + \sigma_a$ .

Radiance leaving point  $x$  in the direction  $\vec{\omega}$  is the sum of exit radiance from the nearest surface along this direction and in-scattered radiance from the medium among the whole length of the ray [15] (see Fig. 1):

$$L(x, \vec{\omega}) = T_r(x \leftrightarrow x_s)L(x_s, \vec{\omega}) + \int_0^s T_r(x \leftrightarrow x_t)\sigma_s(x_t)L_i(x_t, \vec{\omega})dt \quad (1)$$

where  $T_r$  is the transmittance, defined as


$$T_r(x \leftrightarrow x_t) = \exp(-\sigma_t \|x - x_t\|) \quad (2)$$

$s$  is the distance through the medium to the nearest surface at  $x_s = x - s\vec{\omega}$ , and  $x_t = x - t\vec{\omega}$  with  $t \in (0, s)$ .  $L(x_s, \vec{\omega})$  is the exit radiance from the nearest surface, which is governed by the *rendering equation* [26].  $L_i(x_s, \vec{\omega})$  is the in-scattering radiance at  $x_t$  from all directions  $\vec{\omega}_t$  over the sphere of directions  $\Omega_{4\pi}$  using the phase function  $p$ , defined as

$$L_i(x_t, \vec{\omega}) = \int_{\Omega_{4\pi}} p(\vec{\omega}, \vec{\omega}_t)L(x_t, \vec{\omega}_t)d\vec{\omega}_t \quad (3)$$

### 3.2 SD-tree path guiding

Our method is based on the practical path guiding method proposed by Müller et al. [2], for which we now give a detailed review. They used a binary tree to represent spatial distribution and quadtrees to



**Fig. 1** Light transport in participating media. Symbols are defined in Eq. (1).

represent angular distribution of incoming radiance, so we call it *SD-tree path guiding*.

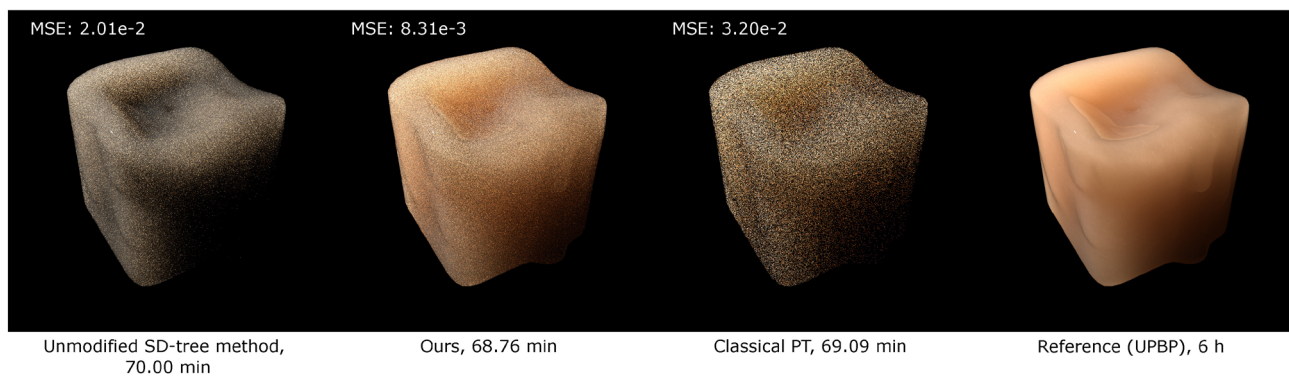
They learned an approximate representation of the scene's spatio-directional radiance field in an unbiased and iterative manner. To do so, they proposed an adaptive spatio-directional hybrid data structure, an SD-tree, for storing and sampling incident radiance. The SD-tree consists of an upper part, a binary tree that partitions the 3D spatial domain of the light field, and a lower part, a quadtree that partitions the 2D directional domain.

They used an iterative scheme similar to the one proposed by Vorba et al. [1]. They trained the incoming radiance sequence  $L_1, \dots, L_M$ , where  $L_1$  is estimated with just BSDF sampling, and for  $k > 1$ ,  $L_k$  is estimated by combining samples of  $L_{k-1}$  and the BSDF via multiple importance sampling. Two SD-trees are kept at the same time: one is used to represent  $L_{k-1}$  and the other is used for updating the incoming radiance distribution in the  $k$ th iteration.


For a given shading point, two steps are performed to sample the outgoing direction of the reflected ray. First, they descend spatially through the binary tree

to find the leaf node containing the shading point position. Next, they sample the outgoing direction from the quadtree contained in the spatial leaf node via hierarchical sample warping.


The SD-tree method [2] cannot be used for participating media directly, for two reasons. Firstly, the SD-tree is learned from the camera side, which leads to very high learning costs for participating media with boundaries. In Fig. 5, we use the unmodified SD-tree method [2] on the Candle scene. It produces a significantly darker result than the reference, due to its difficulty in finding the path to the light source from the camera side. Secondly, multiple importance sampling is not an efficient way to sample the product of two high-frequency functions (the incoming light distribution and the phase function). Figure 3 illustrates this issue: multiple importance sampling produces a darker image than resampling. This issue also appears without participating media, although it is less visible. Figure 4 shows a comparison of multiple importance sampling and resampling with high-frequency BRDFs and illumination: resampling reduces noise compared



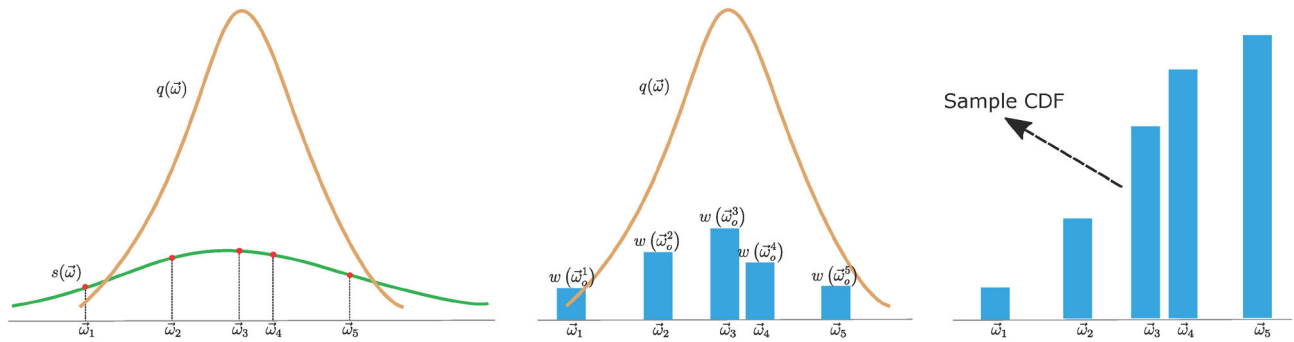
**Fig. 2** Comparison of SD-tree path guiding and our method, with equal time, on the Candle scene. The SD-tree method has difficulty in finding paths to the light source and produces a darker image.



**Fig. 3** Comparison of multiple importance sampling and our RIS with equal SPP (400) on the Candle scene.



**Fig. 4** Comparison of multiple importance sampling and resampling on the Kitchen scene, with highly glossy surfaces. Resampling reduces noise even without participating media.



**Fig. 5** To sample the target function (orange), we first choose a function as the source function (green), and then sample the source function to get 5 outgoing direction candidates. In the second step, we evaluate the weight of each candidate, the target function value divided by the source function value. Finally, we compute the cumulative distribution function of the candidates' weights and use it to get the final sample.

to multiple importance sampling. In this paper, we extend this method to participating media and solve these issues.

### 3.3 Resampled importance sampling

Monte Carlo integration of a function  $f(\vec{\omega})$  is defined as

$$\hat{I} = \frac{1}{M} \sum_{j=1}^M \frac{f(\vec{\omega}_j)}{q(\vec{\omega}_j)}$$

where  $M$  is the sample count and  $q$  is the probability density function.

Ideally, the optimal sampling method is sampling  $q$  with the same distribution as  $f$ , which results in zero variance and error. However, this is impractical.  $q(\vec{\omega})$  may not be sampled. Talbot et al. [27] introduced resampled importance sampling (RIS) to approximately sample  $q(\vec{\omega})$  by firstly sampling another function  $s(\vec{\omega})$  which can be sampled. RIS is a robust variance reduction technique to sample a function.

We first define some notation.  $q(\vec{\omega})$  is called the target function, and  $s(\vec{\omega})$  is called the source function, which can be an arbitrary function which can be sampled.

The entire RIS algorithm runs as follows (see Fig. 5):

1. Sample  $N$  candidates ( $N \geq 1$ ) from  $s(\vec{\omega}_i)$ .

2. Calculate the weight for each candidate:  $w(\vec{\omega}_o^i) = q(\vec{\omega}_i)/s(\vec{\omega}_i)$ .
3. Choose the final sample from the candidates with probability proportional to the candidates' weights  $w(\vec{\omega}_o^i)$ .

A weighting term  $W$  is added to the standard Monte Carlo estimator for unbiasedness:

$$W = \frac{1}{N} \sum_{i=1}^N w(\vec{\omega}_o^i)$$


Incorporating this weighting term, we get the RIS Monte Carlo estimator:

$$\hat{I}_{\text{ris}} = \frac{1}{M} \sum_{j=1}^M \frac{f(\vec{\omega}_j)}{q(\vec{\omega}_j)} \frac{1}{N} \sum_{i=1}^N w(\vec{\omega}_o^i)$$

In Fig. 6, we compare importance resampling and multiple importance sampling to sample two function product. Given the analytic formula of two functions, we know the target function. By comparison, we can see RIS provides closer curve to the target function, thus producing higher sampling quality than MIS.

## 4 Algorithm

We now present our path guiding method for participating media. We use an SD-tree structure to represent the incident light field (Section 4.1). Our method has two steps: in the training step, the distribution is learned by shooting light rays;



**Fig. 6** Comparison of multiple importance sampling and RIS for the product  $f_1 f_2$  of two functions. The RIS produces a result much closer to the target product than MIS.

in the rendering step, we sample the out-scattering direction (Section 4.2) using importance resampling of the trained distribution and the phase function (Section 4.3). To further improve the efficiency of our method, we propose a method to optimally select between multiple importance sampling and importance resampling, depending on the anisotropy of the phase function (Section 4.4).

#### 4.1 Volume representation for radiance distribution

We propose a representation for incoming radiance distribution in participating media with SD-trees: a spatial tree for discrete samples in the volume and a directional tree in each leaf node for angular distribution.

For each medium, the spatial tree is built starting from an axis aligned bounding box (AABB) of the

scene. Nodes are divided into two child nodes when the contained sample count is larger than a threshold. The division dimension is chosen iteratively. The threshold depends on the iteration (see Section 4.2). The total count of STree nodes is bounded by the memory budget, or precisely by the total STree node count ( $N$ ).


When arriving at a leaf node, a directional node is built starting from a sphere, and subdivided into four child nodes. Each node is subdivided if its contained energy is larger than 1% of the entire energy in the entire directional tree.

#### 4.2 Learning and rendering

Our method has two steps: a learning step and rendering step. Figure 7 shows an overview of our method.

In the learning step, we learn the incoming radiance distribution at sampled positions in the volume from the light side. Participating media usually has refractive boundaries. If the path starts from the camera, it is difficult to reach the light sources, especially when the light sources are small. On the other hand, it is easier to work in a bidirectional manner. Furthermore, the learned distribution from the light side is view independent, so it can be reused when only the camera is modified.

We learn the incoming radiance distribution iteratively, use  $2048c$  photons in the first iteration, and increase the photon count by a factor of 2.  $c$  is a constant value related to the resolution of the DTree, and set to 12,000, following Ref. [2]. In each iteration, we sample the light sources and shoot rays from them. The rays are traced in the scene until refracted into the media and then we start a random walk in the medium. For each scattering event, the



**Fig. 7** Our algorithm has a training step and a rendering step. In the former, we learn the incoming radiance distribution iteratively. In each iteration, light rays are shot from light sources, scattered in the participating media, and stored in a spatial-directional hierarchy. In the rendering step, the camera rays are scattered in the media, while the outgoing direction sampling method is chosen by considering the anisotropy of the phase function. For highly anisotropic media, the outgoing directions are sampled via the resampling product of the learned incoming radiance and the phase function. For low anisotropic media, the outgoing directions are sampled via multiple importance sampling of these quantities.

outgoing direction is found by importance sampling the phase function. For the STree, the leaf node is subdivided, if the leaf node contains more than  $\sqrt{2^k c}$  items.  $k$  represents the iteration count. The DTree is also updated according to the rule described in Section 4.1. The maximum iteration count is set between 3 and 6 for all our test scenes.

In the rendering step, following classical path tracing, we shoot rays from the camera, and trace them in the scene. When the rays arrive at the volume, for each scattering event, we first traverse the spatial tree to find the leaf node that contains the scattering event. We then use the directional tree in the leaf node for direction sampling. The outgoing direction is computed by a combination of importance resampling and multiple importance sampling of the learned incoming direction and the phase function. We give details of our importance resampling method next.

### 4.3 Resampling the incoming radiance distribution and phase function

In the rendering step, we have two known functions: the phase function  $p(\vec{\omega}_i, \vec{\omega}_o)$  and the learned incoming radiance distribution function  $\hat{L}_i(\mathbf{x}, \vec{\omega}_o)$  (we use  $\vec{\omega}_o$  for value from the point view of the camera) at position  $\mathbf{x}$ . We use  $\hat{L}_i(\vec{\omega}_o)$  for short. Our goal is to importance sample the direction  $\vec{\omega}_o$  considering the target function, the product  $p(\vec{\omega}_i, \vec{\omega}_o)\hat{L}_i(\mathbf{x}, \vec{\omega}_o)$ .

We use the importance resampling approach [27] to approximate the product sampling of the two functions (see Fig. 6) as follows:

1. using weighted addition of the phase function and the incoming radiance distribution function as the source function, sample the source function to get  $n$  outgoing direction candidates,  $\vec{\omega}_o^k$ ;
2. evaluate the weight of each candidate using Eq. (4),
3. compute the cumulative distribution function (CDF) of the weights of the candidate out-scattering directions and use it to importance sample the final out-scattered direction.
4. evaluate the scattering event for the chosen direction using Eq. (5).

The weight of each sample,  $1 \leq k \leq n$ , is defined as

$$w(\vec{\omega}_o^k) = \frac{q_{\hat{L}_i}(\vec{\omega}_o^k)q_p(\vec{\omega}_i, \vec{\omega}_o^k)}{\lambda q_p(\vec{\omega}_i, \vec{\omega}_o^k) + (1 - \lambda)q_{\hat{L}_i}(\vec{\omega}_o^k)} \quad (4)$$

where  $q_{\hat{L}_i}$  represents the pdf of the incoming radiance distribution function,  $q_p$  represents the pdf of the

phase function, and  $\lambda$  is the weight of the phase function (set as 0.9 for all of our test scenes).

In our implementation, with only 6 candidates, our method can provide convincing results.

After choosing the out-scattering direction  $\vec{\omega}_o$ , we evaluate the scattering event:

$$f(\vec{\omega}_i, \vec{\omega}_o) = \frac{p(\vec{\omega}_i, \vec{\omega}_o)t}{q_{\hat{L}_i}(\vec{\omega}_o)q_p(\vec{\omega}_i, \vec{\omega}_o)} \quad (5)$$

where  $t$  is the weight computed as

$$t = \frac{1}{n} \sum_{k=0}^n w(\vec{\omega}_o^k) \quad (6)$$

### 4.4 Selective importance sampling

In addition to the importance resampling method, we propose a selective importance sampling method, combining multiple importance sampling (MIS) and resampled importance sampling (RIS), based on the anisotropy of the phase function: we use MIS for media with low anisotropy ( $g < 0.5$ ) and RIS for highly anisotropic media ( $g > 0.5$ ).

For MIS, we sample the phase function with probability  $\mu$  and the incoming radiance distribution function with probability  $1 - \mu$ , where  $\mu$  is set to 0.5 for all test scenes. Details of RIS can be found in Section 4.3.

The main reason for selective importance sampling is that the MIS method can convergence faster than RIS for isotropic media but suffers from high variance for highly anisotropic media.

## 5 Results and discussion

### 5.1 Preliminaries

We have implemented our algorithm inside the Mitsuba Renderer [28]. We compared our algorithm to the original path tracing method and UPBP (Unified Points, Beams, and Paths) [20] for quality validation. We used UPBP as the reference in most cases and path tracing for the Bathroom scene and Candle scene with rough boundaries. All timings in this section were measured on a 2.20 GHz Intel i7 with 20 cores and 40 threads, with 32 GB of main memory.

All materials in our scenes are homogeneous, with Henyey–Greenstein phase functions and refractive boundaries. Material properties are taken from Krivánek et al. [20], Narasimhan et al. [29], and Holzschuch [30] (see Table 1).

**Table 1** Parameters for the materials used in this paper

Name	$\alpha$			$\ell$			$g$
	R	G	B	R	G	B	
Oil	0.004	0.454	0.100	9.71	11.63	2.74	0.9
Wax	0.980	0.962	0.750	0.65	0.63	0.59	0.8
Bumpy Sphere	0.955	0.677	0.457	4.55	3.23	2.17	0.8
Marble	0.999	0.998	0.997	0.46	0.38	0.33	0.0
Fog	0.973	0.973	0.973	12.16	12.16	12.16	0.3

We measure the convergence rate with mean square error (MSE).

## 5.2 Qualitative validation

We first compare our method with the reference solution and also classical path tracing given equal time. See Figs. 8, 9, 11–14. To demonstrate the ability to handle high frequency effects, we also provide a comparison between our method and path tracing for the Bumpy Sphere scene with single scattering only in Fig. 10. In all these test scenes, our method provides better quality (smaller MSE) than path tracing with equal time, thanks to path guiding.

We also compare the denoised results (using Intel Open Image Denoiser [31]) for our method and equal time path tracing. Both visual and quantitative

comparisons show that our method provides better input for denoising. The cost of denoising is negligible.

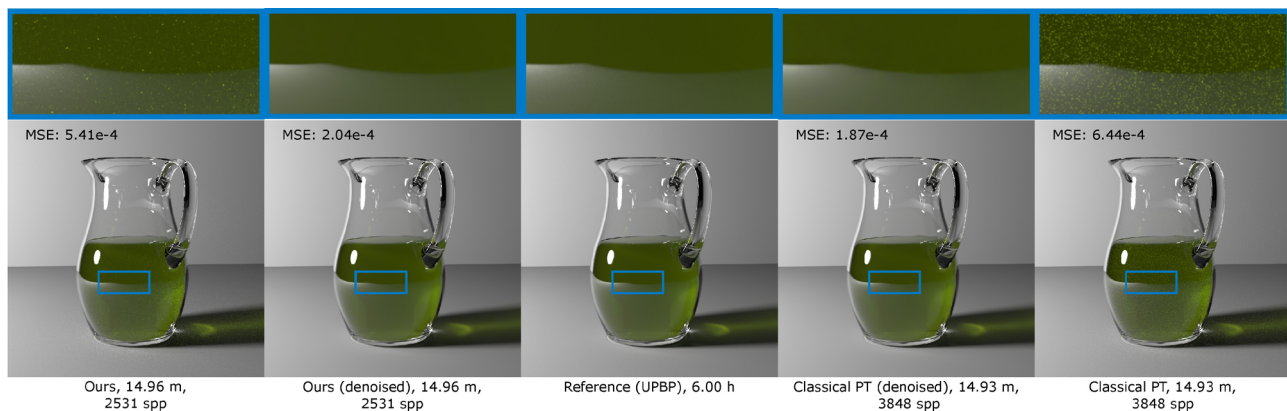
These scenes contain both highly anisotropic media, such as wax in Fig. 9 and Bumpy Sphere media in Fig. 11 and also isotropic media, such as marble in Fig. 12 and fog in Fig. 14.

## 5.3 Performance and timings


For all test scenes, we report the computation time of our method, path tracing with equal time and the reference method in Table 2. For our method, we also provide the cost of the training step.

Comparing our method and classical path tracing with equal time, our method has lower error (MSE), which confirms the superiority of our method. With equal time, our method has fewer samples than path tracing, due to the extra cost for the resampled importance sampling in path guiding.

Our method has an extra training cost to learn the incoming radiance distribution. The training costs depend on the scene complexity, media types, and number of learning iterations. In the equal time comparison, training time is included in the total time.



**Fig. 8** Comparison of our method and path tracing with equal time for the Oil scene.



**Fig. 9** Comparison of our method and path tracing with equal time for the Candle scene.






Fig. 10 Comparison of our method and path tracing with equal time for the Bumpy Sphere scene (single scattering only).




Fig. 11 Comparison of our method and path tracing with equal time for the Bumpy Sphere scene (full solution).

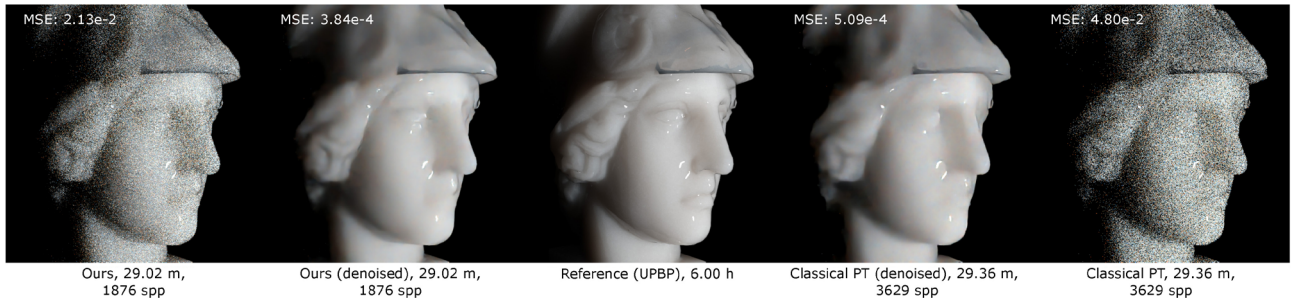


Fig. 12 Comparison of our method and path tracing with equal time for the Athena scene.

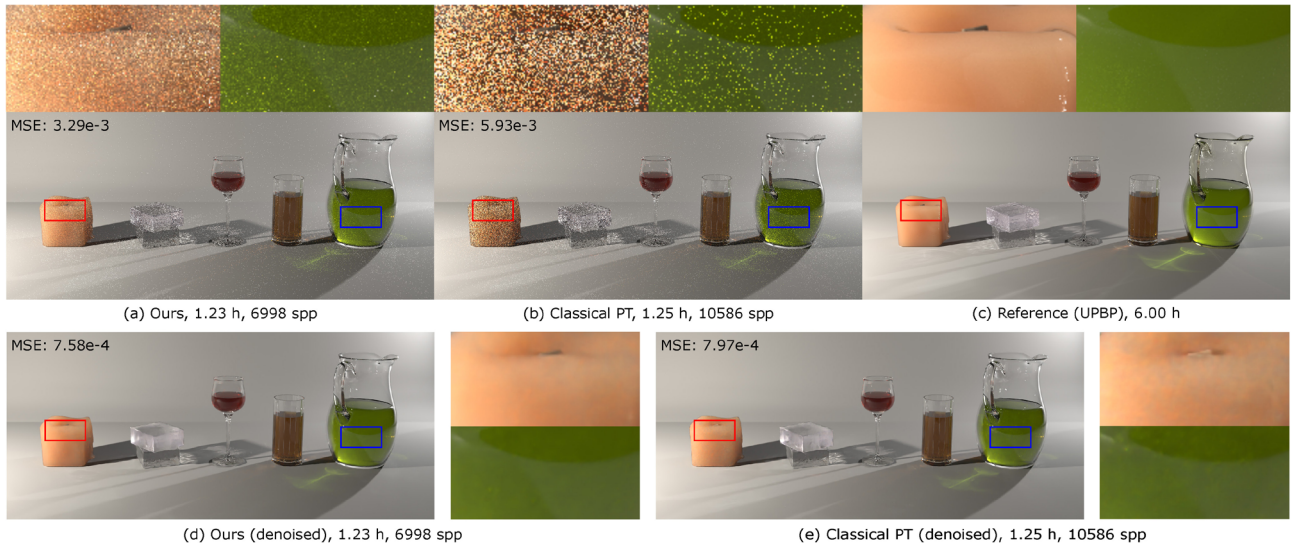



Fig. 13 Comparison of our method and path tracing with equal time for the Still Life scene.



**Fig. 14** Comparison of our method and path tracing with equal time for the Bathroom scene.

**Table 2** Computation time, memory costs, and error for our test scenes. #iter is number of training iterations

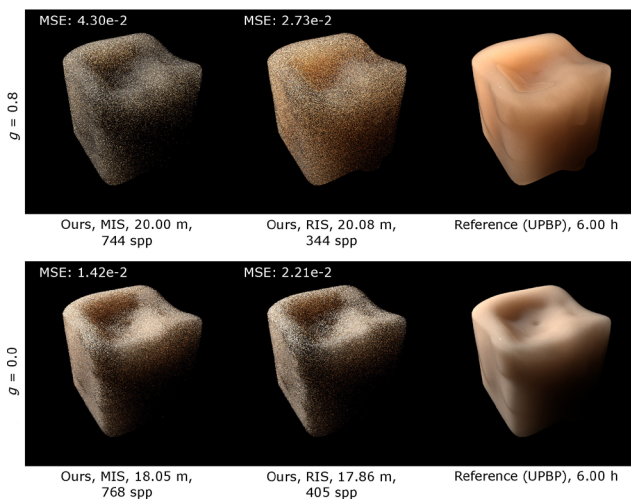
Scene	Our method						Path tracing (equal time)			Reference	
	#iter	Time (train) (m)	Memory (MB)	spp (render)	Time (render) (m)	Time (total) (m)	MSE	#sample	Time (m)	MSE	Time (h)
Bumpy Sphere	4	5	0.17	3397	9.98	14.98	6.04e-3	7989	15.20	6.48e-2	6
Candle	6	10	0.67	344	10.08	20.08	2.73e-2	1395	20.05	9.29e-2	6
Oil	5	5	0.07	2531	9.96	14.96	5.41e-4	3848	14.93	6.44e-4	6
Athena	6	11	7.25	1876	18.02	29.02	2.13e-2	3629	29.36	4.80e-2	6
Still Life	6	14	4.76	6998	59.80	73.80	3.29e-3	10586	75.00	5.93e-3	6
Bathroom	3	10	12.30	3954	20.22	30.22	3.73e-2	6081	30.37	7.74e-2	223.96

#### 5.4 Selective sampling method validation

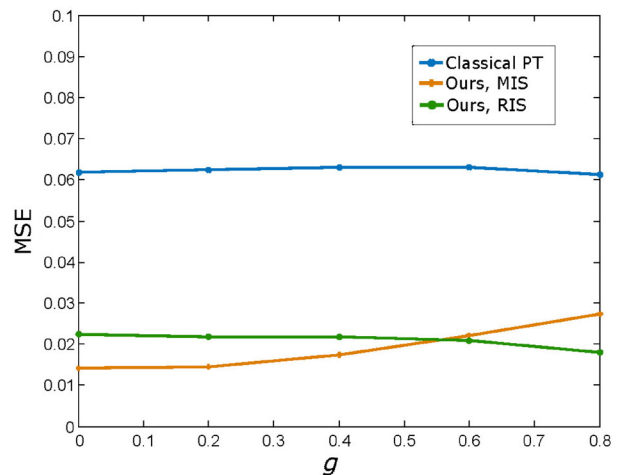
In Fig. 15, we show the results of two sampling methods: multiple importance sampling (MIS) and importance resampling (RIS). For isotropic media ( $g = 0.0$ ), MIS provides higher quality than RIS with equal rendering time. The cost of MIS is lower than RIS, so it can use more samples for a fixed amount of time. For highly anisotropic media ( $g = 0.8$ ), RIS has higher quality than MIS. With highly anisotropic media and a small light source, both the incoming light distribution and the phase function have high

frequency. Multiple importance sampling between these two functions can result in a very low probability density function for one of the functions. This is shown in Fig. 3. The same problem also exists in surface rendering, but it becomes even more obvious in participating media due to the high dimensionality of paths. On the other hand, resampling tries to sample the product of the two functions, and although it only uses 6 candidates, it still provides higher sampling quality.

Figure 16 displays the impact of media anisotropy



**Fig. 15** Comparison of our method with multiple importance sampling (MIS) and our method with importance resampling (RIS) on the Candle scene. Materials:  $g = 0.0$  and  $g = 0.8$ .



**Fig. 16** Error (MSE) for our method with multiple importance sampling, our method with resampled importance sampling, and path tracing with equal time (about 20 minutes), as a function of media anisotropy, for the Candle scene.

on path tracing and our two different sampling methods: MIS and RIS. With low anisotropy (from 0.0 to about 0.5), MIS provides better quality than RIS. With higher anisotropy, RIS provides better quality than MIS, thanks to the better sampling quality. Thus, in our implementation, we use 0.5 as a threshold to choose between the sampling methods.

### 5.5 Parameter analysis

Figure 17 shows the convergence (error) of our method and path tracing with varying sample counts. Increasing the sample count, both methods have decreasing error. Our method consistently has lower error than path tracing.

Figures 18 and 19 show the impact of several parameters on our algorithm for the Candle scene in Fig. 9. Figure 18 displays error (MSE) as a function of resampling sample count with the same number of samples per pixel (spp) and the same rendering time respectively. With the same spp, the difference between our result and the reference decreases, as the resampling sample count increases. With the same rendering time, the difference between our result and the reference decreases at first and then increases after 6. More sample candidates for resampling produce higher quality samples, but the cost becomes very expensive with many candidates. With equal time, fewer spp can be used, so rendering quality decreases with more than 6 resampling candidates. In practice, we use 6 resampling candidates for all the test scenes. Figure 19 displays the difference (MSE) between our method and the reference as a function of the number

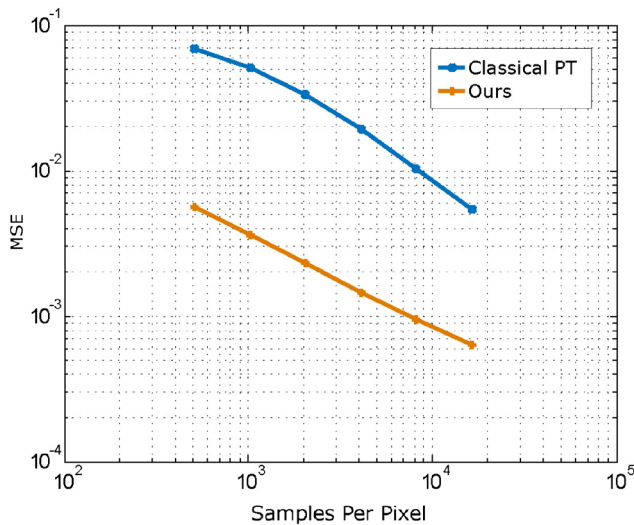


Fig. 17 Error (MSE) for our method and path tracing, as a function of the sample count per pixel, for the Candle scene.

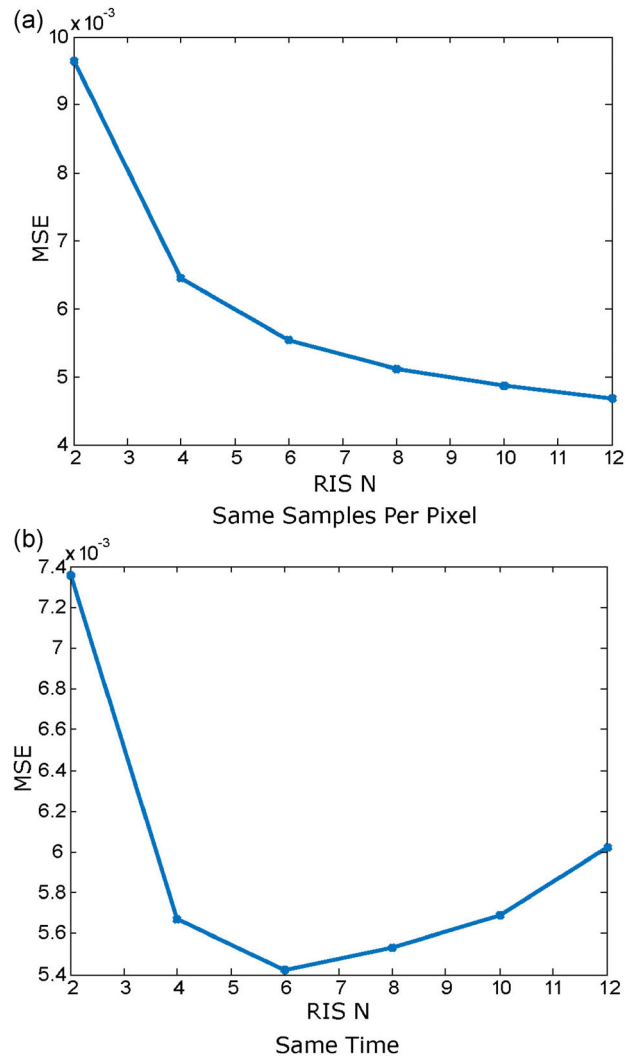
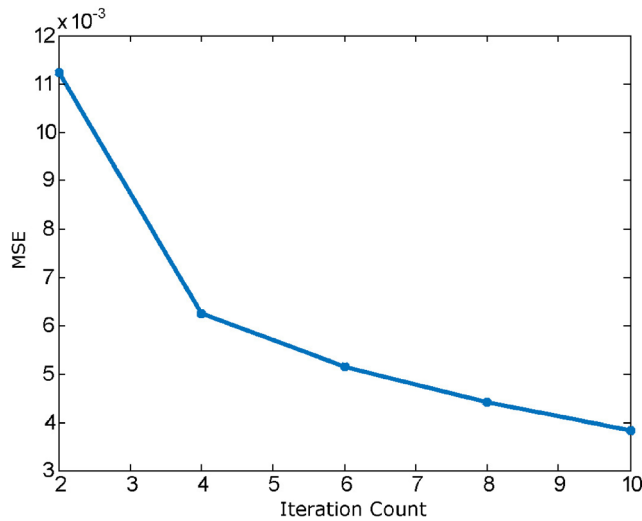


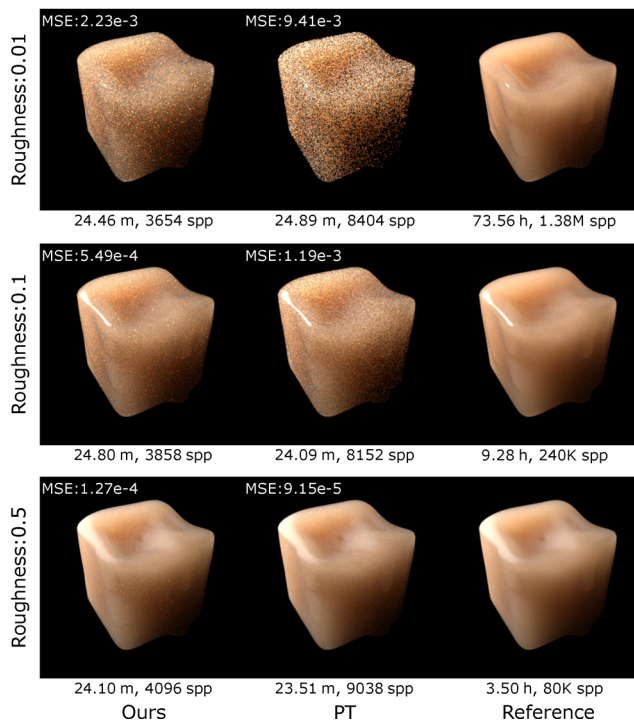
Fig. 18 Error (MSE) as a function of resampling sample count with equal samples per pixel (spp: 1029) (a) and equal time (30 minutes) (b) for the Candle scene.

of training iterations; increasing them decreases the error. The rate of decrease becomes slower after 4 iterations. With more iterations, we can get higher quality, but the training time increases by a factor of 2. We thus use 3–6 iterations for all of our test scenes.

Figures 20 and 21 show the impact of surface and media parameters on our algorithm for the Candle scene (Fig. 9). Figure 20 shows the impact of surface roughness on our algorithm. We compare our method and path tracing with varying surface roughness (0.01, 0.1, and 0.5). With low roughness (0.01 or 0.1), our method has higher quality than path tracing with equal time. With high roughness (0.5), path tracing produces higher quality results than our method. When the surface is rough, the

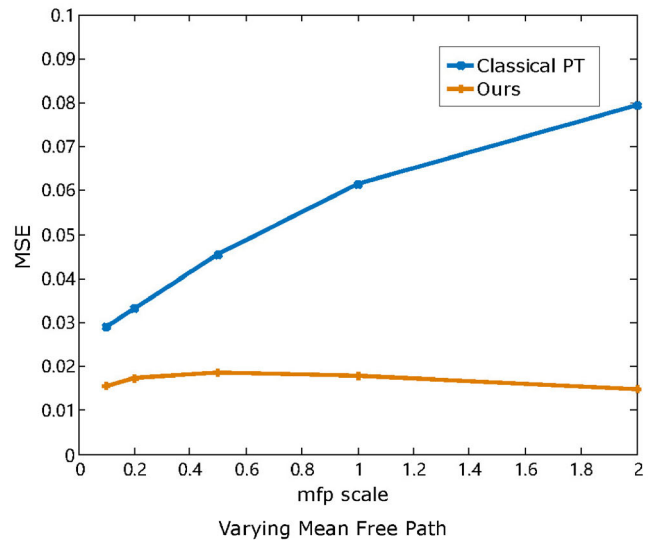
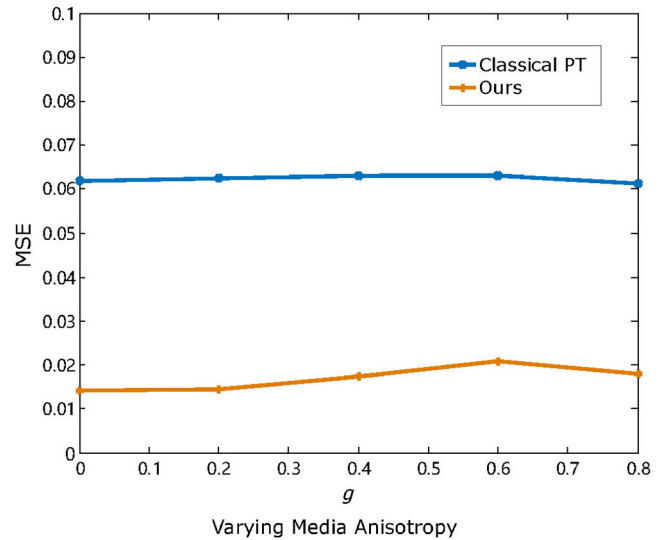


**Fig. 19** Error (MSE) as a function of training iterations with varying training time and the same rendering time (30 minutes) for the Candle scene.



**Fig. 20** Comparison between our method and path tracing with equal time on the Candle scene with varying surface roughness.

rays for the camera have a higher chance of reaching the light source, so our method loses the benefit of guiding the camera rays to the light source. With the extra cost for guiding, our method results in slower convergence for rough surfaces. Figure 21 displays the impact of media anisotropy and mean free path on our algorithm. From isotropic to highly anisotropic media, and from small to large mean free path, our



**Fig. 21** Error (MSE) for our method and path tracing with equal time (about 20 minutes), as a function of media anisotropy and the media mean free path, for the Candle scene.

method consistently produces better quality results than equal-time path tracing. The rendered images can be found in the Electronic Supplemental Material.

### 5.6 Discussion and limitations

In Fig. 2, we compare our method with Müller et al. [2] on the Candle scene. Müller et al. produce darker results, as their method has difficulties to find paths to the light from the camera side due to the refractive boundaries; the multiple importance sampling used in their method leads the path to exit the participating media too early. Thus, we believe the method of Müller et al. [2] cannot be used for participating media.

Our method shows an advantage for media with

refractive boundaries. See Fig. 20: as the surface roughness increases, our method loses its benefits due to the expensive sampling method. Automatic switching between the proposed sampling methods and simple phase sampling would be helpful.

We use path tracing as our comparison method, as our method is based on path tracing. We would like to keep the comparison fair.

Our method does not rely on denoising method. We provide denoised results to show that our method can produce better inputs than classical path tracing.

We only consider out scattering direction sampling in our paper, but distance sampling is also very important for media simulation. Including distance sampling would further improve the rate of convergence.

## 6 Conclusions

We have presented a practical path guiding method for participating media. We use an SD-tree to represent the incoming radiance distribution in the volume. We propose a novel product importance sampling method, combining the incoming radiance distribution and the phase function, with only a few samples. The proposed method converges faster than classical path tracing.

We also analyse the light and phase function distributions, and provide a selective importance sampling strategy, to avoid expensive product importance sampling where it is not required. This results in faster convergence.

Our algorithm is suitable for any media, from low frequency effects where multiple scattering is dominant to high frequency effects where single scattering is dominant.

In future, we will consider introducing our method to heterogeneous participating media and we also wish to combine the volume path guiding representation and the surface path guiding representation. As for non-exponentially correlated media [32], changing exponential distance sampling to non-exponential sampling will require further consideration of the volumetric path guiding domain.

## Acknowledgements

We thank the reviewers for their valuable comments. This work was partially supported by the National

Key R&D Program of China under Grant No. 2017YFB0203000, the National Natural Science Foundation of China under Grant Nos. 61802187 and 61872223, the Natural Science Foundation of Jiangsu under Grant No. BK20170857, the fundamental research funds for the central universities No. 30918011320, and ANR project ANR-15-CE38-0005 “Materials”.

**Electronic Supplementary Material** Electronic supplementary material is available in the online version of this article at <https://doi.org/10.1007/s41095-020-0160-1>.

## References

- [1] Vorba, J.; Karlík, O.; Šik, M.; Ritschel, T.; Křivánek, J. On-line learning of parametric mixture models for light transport simulation. *ACM Transactions on Graphics* Vol. 33, No. 4, Article No. 101, 2014.
- [2] Müller, T.; Gross, M.; Novák, J. Practical path guiding for efficient light-transport simulation. *Computer Graphics Forum* Vol. 36, No. 4, 91–100, 2017.
- [3] Herholz, S.; Elek, O.; Vorba, J.; Lensch, H.; Křivánek, J. Product importance sampling for light transport path guiding. *Computer Graphics Forum* Vol. 35, No. 4, 67–77, 2016.
- [4] Reibold, F.; Hanika, J.; Jung, A.; Dachsbacher, C. Selective guided sampling with complete light transport paths. *ACM Transactions on Graphics* Vol. 37, No. 6, Article No. 223, 2019.
- [5] Guo, J.; Bauszat, P.; Bikker, J.; Eisemann, E. Primary sample space path guiding. In: *Proceedings of the Eurographics Symposium on Rendering*, 73–82, 2018.
- [6] Zheng, Q.; Zwicker, M. Learning to importance sample in primary sample space. *Computer Graphics Forum* Vol. 38, No. 2, 169–179, 2019.
- [7] Müller, T.; McWilliams, B.; Rousselle, F.; Gross, M.; Novák, J. Neural importance sampling. *arXiv preprint arXiv:1808.03856*, 2018.
- [8] Křivánek, J.; d’Eon, E. A zero-variance-based sampling scheme for Monte Carlo subsurface scattering. In: *Proceedings of the ACM SIGGRAPH 2014 Talks*, Article No. 66, 2014.
- [9] Meng, J.; Hanika, J.; Dachsbacher, C. Improving the dwivedi sampling scheme. *Computer Graphics Forum* Vol. 35, No. 4, 37–44, 2016.
- [10] Herholz, S.; Zhao, Y. Y.; Elek, O.; Nowrouzezahrai, D.; Lensch, H. P. A.; Křivánek, J. Volume path guiding based on zero-variance random walk theory. *ACM Transactions on Graphics* Vol. 38, No. 3, Article No. 25, 2019.

- [11] Pattanaik, S. N.; Mudur, S. P. Computation of global illumination in a participating medium by Monte Carlo simulation. *The Journal of Visualization and Computer Animation* Vol. 4, No. 3, 133–152, 1993.
- [12] Lafortune, E. P.; Willems, Y. D. Rendering participating media with bidirectional path tracing. In: *Rendering Techniques '96*. Pueyo, X.; Schröder, P. Eds. Springer Vienna, 91–100, 1996.
- [13] Pauly, M.; Kollig, T.; Keller, A. Metropolis light transport for participating media. In: *Rendering Techniques 2000*. Péroche, B.; Rushmeier, H. Eds. Springer Vienna, 11–22, 2000.
- [14] Georgiev, I.; Křivánek, J.; Hachisuka, T.; Nowrouzezahrai, D.; Jarosz, W. Joint importance sampling of low-order volumetric scattering. *ACM Transactions on Graphics* Vol. 32, No. 6, Article No. 164, 2013.
- [15] Chandrasekhar, S. *Radiative Transfer*. New York: Dover Publications, 1960.
- [16] Jensen H. W.; Christensen, P. H. Efficient simulation of light transport in scenes with participating media using photon maps. In: Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, 311–320, 1998
- [17] Jarosz, W.; Zwicker, M.; Jensen, H. W. The beam radiance estimate for volumetric photon mapping. *Computer Graphics Forum* Vol. 27, No. 2, 557–566, 2008.
- [18] Jarosz, W.; Nowrouzezahrai, D.; Sadeghi, I.; Jensen, H. W. A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM Transactions on Graphics* Vol. 30, No. 1, Article No. 5, 2011.
- [19] Jarosz, W.; Nowrouzezahrai, D.; Thomas, R.; Sloan, P.-P.; Zwicker, M. Progressive photon beams. In: Proceedings of the SIGGRAPH Aisa Conference, 2011.
- [20] Křivánek, J.; Georgiev, I.; Hachisuka, T.; Vévoda, P.; Šik, M.; Nowrouzezahrai, D.; Jarosz, W. Unifying points, beams, and paths in volumetric light transport simulation. *ACM Transactions on Graphics* Vol. 33, No. 4, Article No. 103, 2014.
- [21] Bitterli, B.; Jarosz, W. Beyond points and beams: Higher-dimensional photon samples for volumetric light transport. *ACM Transactions on Graphics* Vol. 36, No. 4, Article No. 112, 2017.
- [22] Jarosz, W.; Donner, C.; Zwicker, M.; Jensen, H. W. Radiance caching for participating media. *ACM Transactions on Graphics* Vol. 27, No. 1, Article No. 7, 2008.
- [23] Marco, J.; Jarabo, A.; Jarosz, W.; Gutierrez, D. Second-order occlusion-aware volumetric radiance caching. *ACM Transactions on Graphics* Vol. 37, No. 2, Article No. 20, 2018.
- [24] Jensen, H. W. Importance driven path tracing using the photon map. In: *Rendering Techniques '95*. Hanrahan, P. M.; Purgathofer, W. Eds. Springer Vienna, 326–335, 1995.
- [25] Veach, E.; Guibas, L. J. Metropolis light transport. In: Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques, 65–76, 1997.
- [26] Kajiya, J. T. The rendering equation. *ACM SIGGRAPH Computer Graphics* Vol. 20, No. 4, 143–150, 1986.
- [27] Talbot, J. F.; Cline, D.; Egbert, P. Importance resampling for global illumination. In: Proceedings of the 16th Eurographics Conference on Rendering Techniques, 139–146, 2005.
- [28] Jakob, W. Mitsuba renderer. 2010. Available at <http://www.mitsubarenderer.org/>.
- [29] Narasimhan, S. G.; Gupta, M.; Donner, C.; Ramamoorthi, R.; Nayar, S. K.; Jensen, H. W. Acquiring scattering properties of participating media by dilution. *ACM Transactions on Graphics* Vol. 25, No. 3, 1003–1012, 2006.
- [30] Holzschuch, N. Accurate computation of single scattering in participating media with refractive boundaries. *Computer Graphics Forum* Vol. 34, No. 6, 48–59, 2015.
- [31] Intel. Intel Open Image Denoise. Available at <https://openimagedenoise.github.io/index.html>.
- [32] Jarabo, A.; Aliaga, C.; Gutierrez, D. A radiative transfer framework for spatially-correlated materials. *ACM Transactions on Graphics* Vol. 37, No. 4, Article No. 83, 2018.



**Hong Deng** is currently a master student of computer science at Nanjing University of Science and Technology. His main areas of interest are physically-based rendering, game development, and real-time graphics.



**Beibei Wang** is an associate professor at Nanjing University of Science and Technology. She received her Ph.D. degree from Shandong University in 2014 and visited Telecom ParisTech from 2012 to 2014. She worked as a postdoc in Inria from 2015 to 2017. She joined NJUST in March 2017. Her research interests

include rendering.



**Rui Wang** is a professor at the State Key Lab of CAD&CG, Zhejiang University. He is working in the area of computer graphics. His current interests are mainly in real-time rendering, realistic rendering, GPU-based computation, and 3D display techniques. He received his Ph.D. degree in mathematics in 2007 and bachelor degree in computer science from Zhejiang University in 2001. He was a visiting associate professor hosted by Prof. Kavita Bala at Cornell University in 2012–2014.



**Nicolas Holzschuch** is a senior researcher at INRIA Grenoble Rhône-Alpes, and the scientific leader of the MAVERICK research team. He received his Ph.D. degree from Grenoble University in 1996 and his habilitation in 2007. He joined INRIA in 1997. His research interests include photorealistic rendering and real-time rendering, with an emphasis on material models and participating media.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made.

The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Other papers from this open access journal are available free of charge from <http://www.springer.com/journal/41095>. To submit a manuscript, please go to <https://www.editorialmanager.com/cvmj>.