A survey on rendering homogeneous participating media

Wenshi Wu¹, Beibei Wang¹ (✉), and Ling-Qi Yan² (✉)

© The Author(s) 2021.

Abstract Participating media are frequent in real-world scenes, whether they contain milk, fruit juice, oil, or muddy water in a river or the ocean. Incoming light interacts with these participating media in complex ways: refraction at boundaries and scattering and absorption inside volumes. The radiative transfer equation is the key to solving this problem. There are several categories of rendering methods which are all based on this equation, but using different solutions. In this paper, we introduce these groups, which include volume density estimation based approaches, virtual point/ray/beam lights, point based approaches, Monte Carlo based approaches, acceleration techniques, accurate single scattering methods, neural network based methods, and spatially-correlated participating media related methods. As well as discussing these methods, we consider the challenges and open problems in this research area.

Keywords participating media; Monte Carlo methods; rendering; volume density estimation

1 Introduction

Participating media are frequent in real-world scenes; they include candles, olive oil, skin, and fog (see Fig. 1). They play an important role in realistic rendering for movie production, animation, and video games. Computing illumination simulation in scenes with participating media is a costly process, as light interacts with the participating media in complex ways.

Fig. 1 Different types of participating media. (a) From Ref. [1], reproduced with permission © ACM 2014. (b) From “*Finding Dory*”, © Disney Enterprises, Inc. and Pixar, Inc. (c) From “*Frozen*”, © Disney Enterprises, Inc.

Light is potentially refracted at the boundary if the indices of refraction differ inside and outside the media and may be absorbed or scattered as it travels inside the media. The refraction boundary can gather light in such media and cause high-frequency effects, called volumetric caustics, which are obvious in media with relatively large mean free paths. Scattering blurs incident light and causes low-frequency effects, which happen in media with small mean free paths. Directional phase functions and refraction at interfaces add to the computational complexity. The complex interplay between these different phenomena makes simulating light transport in participating media a difficult and ongoing research problem.

In recent years, several algorithms have been introduced for rendering participating media, such as the many lights based methods (virtual ray lights, VRL) [2], several extensions to photon mapping culminating with unified points, beams, and paths (UPBP) [1], Monte Carlo based methods [3], and point based methods [4]. All these methods greatly improve simulation of participating media. In this paper, we discuss each of these categories in Sections 3–6 and then cover some acceleration techniques in Section 7. Several approaches focus on single scattering only, or volume caustics, as shown in Section 8. More recently, deep learning has been exploited in participating media rendering, and we

¹ School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing, China. E-mail: W. Wu, wws19990401@njust.edu.cn; B. Wang, beibei.wang@njust.edu.cn (✉).

² University of California, Santa Barbara, USA. E-mail: lingqi@cs.ucsb.edu (✉).

Manuscript received: 2021-05-31; accepted: 2021-07-20

cover some related work in Section 9. In classical participating media, particles are assumed to have a white noise random distribution. However, particles in the media might be subject to various forces, leading to positive or negative random distributions. Spatially correlated media have been introduced and studied in the rendering domain in recent years, and we consider progress on this topic in Section 10. In Section 11, we compare some typical methods in the previous mentioned categories. We discuss challenges and future work in Section 12.

2 Background

In this section, we first introduce the properties of participating media, and then show the core of rendering participating media: the radiative transfer equation.

2.1 Participating media properties

In this paper, we focus on homogeneous participating media. Such participating media are described by an absorption coefficient σ_a , scattering coefficient σ_s , and phase function $p(\omega, \omega_t)$, which represent the absorption ratio, scattering ratio, and distribution of outgoing directions for scattering respectively. The sum of the absorption coefficient and scattering coefficient is called the attenuation coefficient, represented as $\sigma_t = \sigma_a + \sigma_s$. Another equivalent expression for a medium is scattering albedo α with $\alpha = \sigma_s / \sigma_t$ and the mean free path inside the material (mfp) l with $l = 1 / \sigma_t$. The mean free path denotes the average length before the first scattering event inside the medium.

According to these properties, we generally classify participating media as high-order scattering dominant or low-order scattering dominant, according to whether they are optically thick or thin, decided by the value of σ_t . Typical high-order scattering dominant media include wax, skin, marble, etc. Low-order scattering dominant media include olive oil, apple juice, etc. Furthermore, we distinguish single-, double-, and multiple-scattering effects, depending on the number of volume scattering events inside the translucent material. Single scattering corresponds to a light path with only one scattering event inside the material, double scattering to paths with two scattering events, and multiple scattering to paths with more than two scattering events. Single

scattering leads to high-frequency effects, while multiple scattering leads to low-frequency effects. Thus, separation in terms of these scattering events is reasonable. Both high-order scattering dominant and low-order scattering dominant media are challenging. For high-order scattering dominant media, we have to simulate a large number of scattering events before convergence. However, the overall appearance of these materials is often very smooth, meaning that much computational effort is used for an almost constant appearance. Low-order scattering dominant media usually introduce high-frequency volumetric caustic effects, due to the presence of double refraction; these are difficult to capture with either density estimation based approaches or Monte Carlo based approaches. Virtual point, ray, or beam methods cannot simulate single scattering. Thus, a specific group of methods are used to simulate single scattering.

2.2 Radiative transfer equation

Light transport within a participating medium is described by the radiative transfer equation [6], which defines the radiance that reaches a point x from direction ω as a sum of exit radiance from the nearest surface in this direction and in-scattered radiance from the medium among the whole length of the ray: see Fig. 2. This can be expressed as

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

where T_r is the transmittance, defined as

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

s is the distance through the medium to the nearest surface x_s , and x_t is a point with distance to surface x_s between 0 and s . $L(x_s, \omega)$ is governed by the rendering equation [7]. $L_i(x_t, \omega)$ is the in-scattering radiance at x_t from all directions ω_t over the sphere of directions $\Omega_{4\pi}$ using the phase function p , and is defined as

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

Fig. 2 The radiative transfer equation.

Like the rendering equation, the radiative transfer equation does not have an analytic solution. Computational solutions to the radiative transfer equation may be found by volumetric density estimation based, point based, many lights based, or Monte Carlo based methods. We discuss them in the following sections.

All of these methods have their advantages and disadvantages. Volumetric density estimation based, many lights based, and point based methods require a lighting pass to distribute points or other elements in the volume and then use the stored elements in the rendering pass. They use the elements in different ways: using density estimation, or gathering the contribution from these elements, or their combination. Thanks to the extra structure, these three types of methods are usually faster than Monte Carlo based methods, at the cost of extra storage. Regarding the scattering types, many lights based methods cannot handle single scattering, while the others are able to handle all types of scattering effects. Monte Carlo based methods and a subset of volumetric density estimation based methods [5] can produce unbiased results, while the others are biased. We discuss these methods in detail in the following sections.

3 Volumetric density estimation based methods

Among all the types of rendering algorithms for participating media, density estimation is the most commonly used. The basic idea of volumetric density estimation approaches is to distribute light photons, rays, or beams in the medium, and then estimate the density of these elements in a kernel size. In general, these methods include two passes: a lighting pass and a rendering pass. In the lighting pass, the rays are shot into the medium and are scattered until having negligible energy or reaching a maximum depth. The scattering events are stored to represent the light distribution, which can be represented in several manners: photons, beams, planes, etc. In the rendering pass, the camera rays are refracted in the media and gather the contribution from these representations with a certain density estimator, depending on the type of representation. The reuse of cached elements for all pixels makes it efficient.

Jensen and Christensen [8] generalized the photon-mapping algorithm from light travel between object

surfaces to light transport in participating media. In the lighting pass, they traced the light in the volume and stored each volumetric interaction as a photon. In the rendering pass, they refracted the camera ray into the volume, sampled the refracted camera ray into camera samples, and then gathered the contribution of the stored photons to the camera samples with density estimation with a 3D kernel. The estimator is a point-to-point estimator. Later, Jarosz et al. [9] improved this algorithm, by gathering the contribution of the photon map along camera rays rather than camera samples, speeding up convergence. The estimator is a point-to-beam estimator. Furthermore, Jarosz et al. [10] replaced photons with photon beams in the lighting pass, which makes the representation more compact, resulting in a further faster convergence for some media, like fog. The estimator is beam-to-beam. All of these works rely on a huge volumetric photon map, which is memory costly. Thus, a progressive method with photon beams [11] was proposed to eliminate memory limitations. The huge volumetric photon map is replaced by many small photon maps which are generated iteratively. In each iteration, the small photon map is generated, used to update the contribution of the shading points, and then is discarded. This progressive strategy greatly improves the practicality of this approach.

Křivánek et al. [1] found that different representations of points, beams, paths, etc., are suitable for light transport of different participating media or surfaces. Photons are suitable for high-order scattering objects, and beams are suitable for low-order scattering objects, while paths are suitable for volume effects with specular surface rendering. Based on these observations, the estimators based on volumetric density estimation and estimators of Monte Carlo path tracing are combined through multiple importance sampling (MIS), resulting in a unified solution, called UPBP. This unified model can simulate light transport in any type of participating media, by automatically choosing a suitable representation, as shown in Fig. 3.

Note that all of these previous methods are biased, although they are consistent. Recent work by Bitterli and Jarosz [5] further improves the convergence, by extending to higher-dimensional expressions: photon plane and photon volume. They further improved convergence efficiency and achieved unbiased results by using a zero-order estimator. The different types

Fig. 3 Contributions of different estimators to the UPBP image of the Still Life scene. Reproduced with permission from Ref. [1], © ACM 2014.

of estimators are shown in Fig. 4. However, their estimators suffer from singularities, and require at least two bounces in the medium past any surface, leaving the remaining transport to other techniques. Deng et al. [12] solved these issues by generalizing photon planes to photon surfaces, resulting in different types of estimators, including the new “photon cone”, “photon cylinder”, “photon sphere”, and multiple new “photon plane” estimators. They are combined with multiple importance sampling to increase robustness for arbitrary types of participating media. Furthermore, the authors proposed a delta kernel to couple the light and camera subpath to avoid bias. Compared to the prior work [5], their method

reduces the variance significantly and is able to handle any scattering event, including single scattering (see Fig. 5). However, this method is unable to handle medium interactions immediately following scattering from a surface, e.g., volumetric caustics.

Qin et al. [13] introduced an unbiased photon gathering method for participating media. They combined each photon’s path with camera path into one path. They considered all factors in the Monte Carlo methods such as visibility and probability to remove the bias.

Density estimation based approaches are the most commonly used solutions for participating media rendering, since they support all types of scattering

Fig. 4 Different estimators and their results. Reproduced with permission from Ref. [5], © ACM 2017.

Fig. 5 Equal-time variance of different estimators in a scene containing participating media. Left: full light transport in the scene. Middle/right, above: single scattering. Middle/right, below: multiple scattering volumetric transport. The estimators in the middle provide significant variance reduction compared to prior density estimators on the right for equal rendering time. Reproduced with permission from Ref. [12], © ACM 2019.

events and are able to produce high-quality results. However, there are several limitations in this group of methods. First, most methods in this group are biased, except the zero-order kernel in Bitterli and Jarosz [5] and Deng et al. [12]. Second, this group of methods require an extra light pass and extra storage for the light distribution. Third, for high-order scattering dominant media, they take a long time to converge, and for single scattering, they require a large number of elements to simulate high-frequency volumetric caustics, as insufficient photons/rays/beams yield blurry caustics.

4 Virtual point/ray/beam light methods

In the previous section, we presented density estimation based approaches, which rely on caching the light distribution. Similarly, the virtual point, ray, and beam approaches have the same requirements, but use the cached light distribution in a different way, for gathering rather than density estimation. The count of photons/rays/beams is decreased significantly, compared to density estimation based approaches.

Keller [15] proposed instant radiosity, also called virtual point lights (VPLs); it has been used in a large number of surface rendering algorithms. It includes two passes: a lighting pass similar to photon mapping, and a rendering pass which is based on gathering rather than density estimation. Every photon is treated as a virtual point light, which emits light into the scene. The number of VPLs is much fewer than the number of photons, since visibility computation is required for VPLs. VPLs have been widely used for indirect illumination. Several improvements [16–18] have been made to accelerate the visibility computation, using hierarchical pruning. However, these methods suffer from singularity issues, which come from the potential tiny distance between the VPLs and the shading points, yielding spike artifacts in the rendered results. These artifacts may be removed by clamping or blurring [19], at the cost of bias or missing energy. Novák et al. [20] proposed a method that can approximately compensate the bias issue.

Later, VPLs were extended to participating media. Arbree et al. [21] combined lightcuts with diffusion dipoles [22] to approximate subsurface scattering. Multi-dimensional lightcuts [17] extend VPLs to

participating media and other effects (e.g., motion blur), thanks to the efficient hierarchy pruning. However, it still suffers from the singularity issue for multiple scattering events. Novák et al. [2] replaced virtual point lights with virtual ray lights (VRLs) to simulate the light transport in translucent materials. The light rays shot from the light source are scattered in participating media, and then the path segments in media are stored as virtual ray lights (see Fig. 6). The contribution from the stored VRLs to the camera rays is treated as a line-to-line double integral problem. This approach produces higher-quality results than VPLs, but it still suffers from the singularity issue. Later, virtual beam lights (VBLs) with finite thicknesses were proposed to replace VRLs to reduce the singularity issue, producing artifact-free images faster than VRLs (see Fig. 7). They also shoot the VBLs progressively, which significantly reduces the memory cost.

For high-order scattering media, many virtual ray/beam lights are required, and even a progressive solution takes a long time to converge. Thus, Wang et al. [23] proposed a precomputed solution to improve convergence in high-order scattering media, by precomputing the scattering events in an infinite participating medium (see Fig. 13). The precomputed scattering events are stored in two tables for both point light sources and ray light sources. Thanks to the rotational symmetry around the direction of propagation and the almost symmetric shape of

Fig. 6 Virtual point light methods (left) convert the vertices of a random-walk pre-process into a collection of virtual point lights (VPLs). Virtual ray lights (right) convert entire segments of the random-walk into virtual ray lights instead. This results in denser sampling, provably weaker singularities, and higher quality when estimating illumination from the collection of virtual lights (below). Reproduced with permission from Ref. [2], © ACM 2012.

Fig. 7 Indirect illumination on surfaces and in the medium. Progressive virtual ray beams (right) significantly reduce singularity artifacts compared to the virtual ray lights method (left). Progressive rendering effectively eliminates singularities in four indirect light transport paths (below) using novel importance sampling schemes and a new compatible lighting primitive. Reproduced with permission from Ref. [14], © The Author(s) 2012.

the lobes, the precomputed distribution for a single medium is reduced to three dimensions. In the lighting stage, only the rays after surface events are stored, while the volume events are ignored, since they are already included in the precomputed table. In the rendering stage, the contribution is gathered from the stored light rays to the camera ray, by querying the precomputed table, resulting in much faster convergence. As well as in VRL, this approach can also be used in other Monte Carlo rendering algorithms. Although this approach greatly improves the convergence, it has several limitations: it ignores visibility in the media, it is limited to homogeneous translucent media and only handles multiple scattering.

Differently, Georgiev et al. [35] proposed a joint importance sampling method for low-order scattering. They devised joint importance sampling of path vertices in participating media to construct paths that explicitly account for the product of all scattering and geometry terms along a sequence of vertices instead of just locally at a single vertex. Many rendering algorithms could benefit from this approach, including VRLs, to significantly reduce noise and increase performance in renderings with both isotropic and highly anisotropic, low-order scattering.

Similar to using lightcuts to organize virtual point lights, virtual ray lights can also be organized into a hierarchical structure for pruning VRLs. Frederickx et al. [24] clustered the VRLs into a series of ray

slices in the precomputation pass and estimated the required number of clusters by analysis of variance in the rendering pass, resulting in faster convergence. Yuksel and Yuksel [25] treated self-illumination in explosion rendering as a VPL lighting problem, and organized these VPLs into a hierarchy. In addition, multiple scattering is precomputed in the hierarchy, resulting in efficient explosion rendering.

To solve the expensive visibility computation for a large number of virtual light sources in participating media, Huo et al. [26] proposed a sparse sampling and reconstruction method, based on the smooth characteristics of multiple scattering in participating media. The virtual light sources are organized into a small number of representative points using clustering. Accurate visibility computation is performed for these representative points and the visibility of all virtual light sources is obtained through matrix completion, which greatly improves efficiency.

VRL and VBL based participating media rendering methods are faster than density estimation based methods, since much fewer rays or photons are required in this type of methods. However, the visibility computation is still expensive, while density estimation based methods do not have this issue. Also, they are unable to handle single scattering, and rely on density estimation based methods. Furthermore, they have singularity issues, especially in high-order scattering media, although this issue has been reduced by many techniques. All VRL based methods are

biased, while some density estimation based methods are unbiased.

5 Point based methods

The point-based global illumination (PBGI) algorithm [27] was first proposed to compute diffuse light transport for surface rendering and is widely used in movie production, since it is noise-free and much faster than Monte Carlo based methods. Like the previous two groups of methods, PBGI also includes two passes: a lighting pass and a rendering pass. However, the lighting pass is quite different, since the cached point cloud includes geometric information, which serves to approximate the geometry and is used for visibility computation with rasterization in the rendering pass.

Wang et al. [28] expanded PBGI to participating media rendering. As well as surface points, points are also placed inside the medium. The volume samples include the geometric information, using bounding boxes. The volume and surface samples are organized into hierarchies. During rendering, single, double, and multiple scattering are computed separately (Fig. 8). Single scattering is computed directly, finding light samples closer to the camera ray. To compute double scattering, they traverse the spatial hierarchy to obtain the best tree cut and gather the contributions from these nodes. For multiple scattering, they use a precomputed table to store the resulting contribution. They then add the contributions from single, double,

and multiple scattering. For indirect lighting and multiple scattering after several bounces on the refractive surface, they use the surface samples. This method has the advantage of fast convergence, with little noise during simulation, and it performs well for a large range of materials, from low to high albedo, and from isotropic to highly anisotropic. However, this method also has certain limitations, such as the large amount of multiple scattering data, and does not support scenes containing complex materials (such as those with glossy reflections). Wang and Holzschuch [4] further improved it, with faster single scattering due to tighter bounding boxes, and a GPU implementation to support interactive media rendering and editing. Rendered results for a wax object are shown in Fig. 9.

Recently, Liang et al. [29] introduced frequency analysis theory to single scattering computation in PBGI. Since single scattering is usually high-frequency, and a large number of volume samples are required to produce sharp volumetric caustics, they proposed to use covariance tracing of the volume samples and then adjust the kernel size for single scattering computation, resulting in higher quality with fewer volume samples.

Point based methods for volumetric rendering scale well with scene complexity and provide noise-free results in much shorter time. They provide a natural compromise between computation time and quality, by varying the number of samples. Unlike density estimation based methods, the points in

Fig. 8 PBGI framework for participating media rendering. The algorithm begins by computing incoming light for volume and surface samples. Then it computes single-, double-, and multiple-scattering effects for each camera ray using these volume and surface samples. Reproduced with permission from Ref. [4], © The Author(s) 2012.

Fig. 9 Comparison between Ref. [4] and other volumetric methods for a wax object. For this material, with a large albedo α and a small mean free path ℓ , multiple scattering effects dominate. Reproduced with permission from Ref. [4], © The Author(s) 2012.

PBGI are carefully distributed and include geometric information, so are able to produce high-quality single scattering results with much fewer points and in a shorter time than density estimation based methods. Regarding multiple scattering, the precomputed multiple scattering table avoids tracing scattering events and helps to decrease convergence time significantly. Compared to many lights based methods, more points are used in PBGI to capture high-frequency single scattering, while many lights based methods cannot simulate single scattering. The main limitation for this group of algorithms is that homogeneous materials are assumed. Extension to heterogeneous materials, with spatially varying scattering properties, still requires further work. They also do not consider visibility when summing the contributions from volume and surface samples to a camera sample, resulting in over-estimation when there is occlusion. Furthermore, point based methods are biased. In summary, point based methods are suitable for rendering applications needing limited rendering cost and interactive rendering applications, e.g., interactive media editing and lighting editing.

6 Monte Carlo based methods

Monte Carlo based methods have been widely used in participating media, since they are unbiased and simple. We focus on efficient sampling for Monte Carlo solutions to transport problems. Interested readers can refer to broader surveys of volumetric media rendering in research [30] and production [31].

Monte Carlo based rendering was first proposed for forward path tracing integration [32], which relies on two sampling operations: distance sampling and

phase function sampling. The distance sampling usually considers transmittance attenuation within the medium, varying exponentially with distance. For high-order scattering media, millions of scattering events happen before leaving objects, due to the small steps in distance sampling which depend on the mean free path. Highly anisotropic media require sampling of a high-frequency phase function, with noisy results, even with importance sampling. When the medium is enclosed in a refractive boundary, path sampling is even more difficult, with unidirectional path tracing. Lafortune and Willems [33] later expanded the ideas to bidirectional path tracing, which improves the convergence rate for light transport with refractive boundaries. Pauly et al. [34] proposed the Metropolis light transport (MLT) approach for participating media.

The above Monte Carlo methods take a long time to converge when simulating high-order scattering or highly anisotropic participating media. Thus, further advanced approaches build on them to improve the convergence rate, such as next event estimation, zero-variance random walks, or path guiding.

6.1 Advanced sampling

Joint importance sampling [35] constructs single and double scattering sub-paths, while accounting for the product of phase functions and geometry terms along the sub paths. For isotropic scattering, a fully analytic formula for sampling phase functions and geometry terms for double scattering at once are derived using marginalization. Using tabulation, a generalized method is provided that can handle anisotropic scattering as well. Joint importance sampling samples distances based on geometry terms,

and the transmittance is not importance sampled, so it is unsuitable for high-order scattering media.

6.2 Next event estimation

Jakob and Marschner [36] proposed a new mutation strategy for Metropolis light transport based on manifold exploration, improving sampling for paths involving specular and highly glossy surfaces. It can also be used for participating media, especially for media surrounded by refractive boundaries. The manifold exploration idea further benefits next event estimation (NEE).

NEE is usually used in Monte Carlo (MC) rendering to reduce variance, by estimating direct illumination by sampling a point on the emitter and testing its visibility by casting a shadow ray. NEE is not suitable for participating media with refractive boundaries. Hanika et al. [37] proposed manifold next event estimation (MNEE), which leverages manifold exploration [36] to connect to the light source through multiple refracting surfaces. The main drawback of these techniques is that the search for the boundary vertex may not succeed. Since it relies on geometric derivatives, it does not handle detailed or displaced surfaces well. It was improved by Koerner et al. [38], by searching for a point in the volume which satisfies Fermat's principle instead of walking over the surface. This method is suitable for highly directional, near-delta distributions, since the computational overhead is too expensive. Both of these methods can be integrated into a unidirectional path tracer using MIS.

Recently, Weber et al. [39] proposed multiple vertex next event estimation (MVNEE), which connects the point to the light source with sub-paths generated by perturbation, instead of one segment connections. The sub-paths are generated by perturbing seed paths generated by a path tracer. This approach is intended for multiple scattering in a high-order scattering homogeneous participating medium, and it could be combined with a path tracer via multiple importance sampling. This method significantly reduces noise and increases performance of multiple scattering rendering in highly anisotropic, optically dense media, but it is unsuitable for participating media with boundaries.

6.3 Zero-variance random walk

A zero-variance random walk path means creating a random walk without variance. At every scattering

point, an outgoing direction as well as a distance to the next scattering event has to be perfectly importance sampled by all terms of the measurement equation: the product of phase function, transmittance, and incoming importance (or radiance, depending on the direction of the random walk). In practice, it is impossible to sample such a path, since it is hard to obtain the incoming radiance distribution. Even with path guiding, the incoming radiance distribution is an approximation, which cannot guarantee zero variance. Thus, the present approach works for some simple situations. For high scattering and isotropic participating media, Dwivedi sampling [40] biases the sampling probability distributions to exit the medium as quickly as possible, based on the idea that paths close to the boundary have a higher contribution than paths deep within the volume. Dwivedi sampling approximates the geometry by locally fitting a slab to compute an analytic approximation of the light field. This method leads to lower variance, but is unsuitable for thin geometries (such as ears) with a strong backlight.

Later, Meng et al. [41] solved this issue by taking the geometric characteristics of the object into account in the sampling process, to guide the scattering away from the object as quickly as possible. They proposed two biasing sampling methods: closest points and incident illumination sampling. The first searches for the closest point to the boundary at every scattering vertex to increase the chances to escape thin geometry. The second chooses light vertices proportional to their emission using importance sampling, specifically improving the variance of the random walk for backlit cases.

The main limitation of these works is that they are only suitable for high-order scattering and isotropic participating media.

6.4 Path guiding

Path guiding was first introduced for surface rendering [43–46]. The common goal is to find an “optimal” distribution that can approximate the actual path integral and make convergence faster. In these works, the incoming radiance distribution of some samples is learned and further used by combining it with the bidirectional reflectance distribution function using multiple importance sampling [44] or product importance sampling [45]. Besides path guiding methods in path space, several works have

focused on the primary sample space [47–49]. Deep learning has also been used in path guiding. Müller et al. [50] proposed a deep neural network model for the probability density function of samples. The learned model is leveraged for sampling the ray's outgoing direction. The advantage of this work is the independence from scenes and local point parameters (e.g., including the textured BRDF). However, this method is biased and sampling is very expensive.

Later, path guiding was introduced for volume rendering. Deng et al. [42] extended Müller et al. [44] to translucent materials. They also used a spatial-directional tree to represent the incident radiance distribution and then sample the outgoing direction after scattering considering both the learned lighting distribution and the phase function. More specifically, they introduced resample importance sampling (RIS) to jointly sample the lighting distribution and the phase function, as they observed the low sampling quality of MIS for the product of two high-frequency functions. RIS is combined with MIS, depending on the sharpness of the phase function, to save the computational cost in RIS. Figure 10 shows results of sampling 2D functions by both RIS and MIS. RIS can also improve surface rendering results, when importance sampling the lighting distribution and the BRDF. The proposed method significantly improves the performance of light transport simulation in participating media, especially for small lights and

media with refractive boundaries. This method can handle any homogeneous participating media, from high scattering to low scattering, from high absorption to low absorption, and from isotropic media to highly anisotropic media. Unfortunately, this approach uses naive distance sampling, which has a significant impact on convergence speed.

Herholz et al. [3] proposed a path guiding approach using zero-variance path sampling theory, which is able to guide all sampling decisions (Fig. 11), including distance sampling, direction sampling, Russian roulette, and splitting. Zero sampling is guided by a cached estimate of the adjoint transport solution, which is represented as a k -d-tree in the spatial domain and directional distributions using a parametric mixture model based on the vMF distribution. For distance sampling, they considered the product of transmittance and the adjoint transport solution (e.g., in-scattered radiance), using an incremental sampling strategy to avoid expensive cumulative density function (CDF) sampling. For direction sampling, they computed the product of the vMF and the phase function also represented by a vMF approximation, and then sampled this product (see Fig. 12). Compared to the method of Deng et al. [42], sampling the product of two functions should have higher quality than resample importance sampling of the two functions, although the product operation has extra computational cost. The vMF representation in the spherical domain makes the product operation easier than the quad tree used in Deng et al. [42]. With all the sampling decisions guided by the cached estimate of the adjoint transport solution, the method of Herholz et al. [3] has significantly faster convergence than that of Deng et al. [42] and an unguided path tracer.

Both of the path guiding approaches are unbiased, and are able to handle all types of scattering events, including single scattering and multiple scattering.

Monte Carlo based methods have been used in movie production, since they are unbiased, robust, and have simple parameters. They are able to handle all types of scattering, including single and multiple scattering. However, these methods usually are slow to converge, for both volumetric caustics and smooth multiple scattering effects. Advanced sampling approaches (e.g., path guiding) improve their convergence. However, compared to the previous

Fig. 10 Comparison between multiple importance sampling and RIS of the product of two functions f_1 and f_2 . RIS produces results much closer to the target product than MIS. Reproduced with permission from Ref. [42], © The Author(s) 2020.

Fig. 11 Above: example path containing volume (blue) and surface (yellow) vertices. Note that, in the path tracing algorithm, the path carries visual importance, and this is generated in the opposite direction to the flow of light. Below: the four zero-variance sampling decisions considered. Optimal decisions are whether to scatter within or outside the medium, how far to travel until the next scattering event, how to choose the scattering direction, and when to terminate the path. Reproduced with permission from Ref. [3], © The Author(s) 2019.

Fig. 12 Comparison between standard sampling and path guiding from Herholz et al. [3]. Reproduced with permission from Ref. [3], © The Author(s) 2019.

groups of methods, Monte Carlo based methods still have slow convergence. Thus, Monte Carlo based methods target high-quality rendering applications with a high time budget.

7 Acceleration techniques

In the previous sections, we presented and discussed four groups of participating media rendering methods. The performance of these methods can be further improved by combining them with acceleration techniques, like gradient domain rendering, frequency analysis, radiance caching, and precomputation.

7.1 Gradient-based rendering

The gradient-based method [51] was introduced for homogeneous participating media by Gruson et al. [52], taking advantage of the smoothness characteristics of rendered participating media. In their paper, four different gradient-based estimation methods were introduced, including point to point,

point to beam, beam to beam, beam to light plane. With gradient domain rendering, smoother results are obtained than the original method. More recently, a deep learning based reconstruction approach for participating media has been proposed [53], yielding a better reconstruction solution than Poisson reconstruction. We will describe this approach in Section 9.

7.2 Frequency analysis

Frequency analysis of light transport was introduced for prefiltering and adaptive sampling [54, 55] in surface rendering. Later, Belcour et al. [56] introduced it to participating media, based on the observation that multiple scattering leads to a smooth appearance. They proposed analysis of absorption and scattering of local light fields in the Fourier domain, and derived the corresponding set of operators on the covariance matrices of the power spectrum of the light field. Using these covariance matrices, they proposed several

Fig. 13 The original path is replaced with a special path. Reproduced with permission from Ref. [23], © IEEE 2018.

improvements: adaptive sampling in image space and during contribution gathering for the camera and light samples, etc., resulting in faster convergence of the photon beam approach [11].

7.3 Radiance caching

The radiance caching method [57] uses spherical harmonic functions to represent and store radiance distributions of certain sparse points and related gradient information. The illumination at a new position is interpolated from the existing distribution and gradient, thereby reducing the amount of calculation and providing a smoother result. Since the stored sample points are sparse, the points to be calculated may be differently occluded compared to the stored points, resulting in different illumination; the authors did not consider visibility during interpolation. This problem was solved by Marco et al. [58], by using the second-order Hessian error metric to determine whether the interpolation error is acceptable. Since the visibility is considered, it produces results with higher quality than those of Jarosz et al. [57].

7.4 Precomputation

High-order scattering media can be rendered by density based methods (e.g., UPBP), many lights based methods (e.g., VRL), and Monte Carlo based methods (e.g., MEMLT). However, all of these methods have very slow convergence. This issue was solved by Wang et al. [23], by precomputing multiple scattering in the infinite participating media. The precomputed multiple scattering data is stored in a table of two-dimensional position and one-dimensional directions, utilizing the symmetry of lobes. The precomputed multiple scattering is used during mutation. For a seed path generated with path tracing, the sub-paths within the media are replaced by a special virtual path with a converged contribution computed from the precomputation

table. During mutation, only surface events and the endpoints of the special path may be mutated, so the path lengths are greatly decreased using this special path, and many fewer mutations are required, since all vertices within the special path remain untouched. This approach results in much faster convergence than the original MEMLT, especially for high-order scattering. Besides MEMLT, this method is also suitable for other bidirectional tracing methods, like VRL and UPBP. Recently, Ge et al. [59] used neural networks instead of tables to represent multiple scattering, which reduces the memory usage from hundreds of MB to kB without significant reduction in quality. Although the precomputation greatly improves the convergence rate, it introduces bias to the rendering results. The contribution of the special path is biased, as it is essentially based on density estimation in the precomputation pass.

8 Accurate single scattering approaches

Single scattering corresponds to light entering the material, being refracted at the first interface, scattering once inside the material, and leaving the material, being refracted a second time at the interface before reaching the camera. The presence of two refractions makes it difficult to compute single scattering using standard methods. Single scattering effects can produce volume caustics, with complicated shapes, under sharp lights, e.g., point lights. With a point light, single scattering can be solved in a deterministic way. In this section, we will show two methods which focus on accurate single scattering solutions.

Walter et al. [60] introduced a method for accurate computation of single scattering effects in participating media under a point light source. This configuration is impossible to render with volumetric path tracing or MEMLT. They treated the single

scattering computation as a search problem: find all points on the surface that connect the point light and the camera sample via Fermat's principle. The presence of a shading normal at surface points makes the problem more complex (see Fig. 14). They computed these entry points on the surface using Newton–Raphson optimization. Using an interval Newton method, all solutions can be found. However, to capture the sharpness of the caustics, a large number of camera samples are required along the camera ray. Wang et al. [62] can also be used for single scattering computation. It produces similar results to those of Walter et al. [60].

Holzschuch [61] improved it by computing the extent of the influence of each triangle on the camera ray, based on the observation that the radiance caused by an individual triangle on the surface varies smoothly and discontinuities correspond to triangle

Fig. 14 (a) Multiple paths (black lines) may connect light L to scattering point V ; the non-refractive approximation (purple line) cannot account for these paths. (b) Geometry of problem at a point P with geometric normal N_g and shading normal N_s . Reproduced with permission from Ref. [60], © ACM 2009.

edges. This method is significantly faster than Walter et al. [60] while providing higher-quality results (see Fig. 15).

Sun et al. [63] proposed an analytical solution for single scattering by participating media without boundaries under a point light, which can achieve a real-time frame rate.

All of the above methods can compute accurate single scattering without noise. However, the computation time highly depends on scene complexity. Moreover, these methods are limited to participating media under point lights.

9 Participating media rendering based on neural networks

Neural networks have also been used for participating media rendering, including atmospheric cloud rendering, multiple scattering representations, BSSRDF models, and reconstruction of gradient-domain volumetric rendering.

9.1 Atmospheric cloud rendering

Atmospheric clouds are high-order scattering media; a very long time is required to simulate the multiple scattering events. Kallweit et al. [64] proposed to render atmospheric clouds using a radiance-predicting neural networks (RPNNs) for multiple scattering. An RPNN represents the radiance for each shading configuration, which includes location, direction, the light source, and the density structure of the entire cloud. They proposed a hierarchy of point stencils to represent varying scales of cloud density, as shown

Fig. 15 Comparison among Holzschuch [61], Walter et al. [60], photon mapping, and path tracing. Reproduced with permission from Ref. [61], © The Author(s) 2014.

in Fig. 16. The network structure is based on a multilayer perceptron (MLP), shown in Fig. 16. The hierarchical descriptor is fed progressively into the network. During rendering, Monte Carlo rendering is used for direct lighting and single scattering, and

the neural network is queried for multiple scattering based on the shading configuration. The rendering results of this method are almost identical to those of path tracing, but it is a thousand times faster: see Fig. 17.

Fig. 16 Stencil grid (left) and progressive training network (right). Reproduced with permission from Ref. [64], © The Author(s) 2017.

Fig. 17 Comparison of path-traced references and RPNN approach. Reproduced with permission from Ref. [64], © The Author(s) 2017.

9.2 Multiple scattering representation

Wang et al. [23] proposed to use a precomputed table to represent multiple scattering and used it to accelerate several rendering algorithms. Each medium requires a 3D representation, which makes it impossible to represent the entire media space, further considering the albedo and phase function. To solve this issue, Ge et al. [65] proposed to use a neural network to represent the multiple scattering distribution for arbitrary infinite media, which maps a seven-dimensional function to a radiance, via a four-layer neural network structure (as shown in Fig. 18). The input layer includes the anisotropy value g of the participating medium, the scattering albedo α , the location of the sampling point $r(\rho; z)$, and direction $(\theta; \varphi)$, and the output layer is the radiance. They use this neural network to represent double scattering and multiple scattering, and leave single scattering to other methods. The storage for both double and multiple scattering is reduced from 50 GB to 23.6 kB. During rendering, they reproduced the precomputed table for a specific medium from the neural network, and then used the precomputed table for multiple scattering computation, following Wang et al. [23]. They integrated it into virtual ray light (VRL) method

Fig. 18 Network structure of Ge et al. [65]. Reproduced with permission from Ref. [65], © IEEE 2019.

with an efficient GPU implementation, resulting in an interactive frame rate: see Fig. 19.

Vicini et al. [66] introduced a shape-adaptive BSSRDF model using a conditional variational autoencoder, which learns to sample from a reference distribution produced by a brute-force volumetric path tracer. This model relies on the combination of three neural networks, which together constitute the probability generation model of the BSSRDF sampling surface. The first feature network extracts features from the input parameters. These input parameters include material properties (reflectivity, anisotropy, and refractive index) and a set of geometric features. In order to describe the local geometry, this method uses low-order ternary polynomials to encode approximate distance functions, adapting the model to geometric details, including curvature, thickness, angle, etc. The second scatter network learns to sample from the reference distribution generated by a volume path tracer; the third absorption network performs regression fitting on the scale factor of the distribution of multiple scattering. Figure 20 shows the structure of the three networks. This method supports arbitrary homogeneous media parameters, so maintaining the efficiency of the classic BSSRDF model without the assumption of diffusion theory, and greatly improves overall accuracy. However, it can only handle high-order scattering effects and does not work well on shapes with sharp features.

9.3 Reconstruction for gradient domain rendering

Xu et al. [53] proposed an unsupervised neural network for image reconstruction of gradient-domain volumetric photon density estimation, more specifically for volumetric photon mapping, using a variant of GradNet [67] with an encoded shift

Fig. 19 Results from Ge et al. [65] and UPBP. Reproduced with permission from Ref. [65], © IEEE 2019.

Fig. 20 Network structure of the shape-adaptive BSSRF model. Reproduced with permission from Ref. [66], © ACM 2019.

connection and a separated auxiliary feature branch; it includes volume based auxiliary features such as transmittance and photon density. This network smooths images at a global scale and preserves high-frequency details on a small scale. Their network produces a higher-quality result than previous works (L_1 , L_2 , or GradNet). Although they only considered volumetric photon mapping, it is straightforward to extend this method to other tasks, like beam radiance estimation. However, this method cannot reconstruct images with strong noise, and tends to over-blur some features.

10 Spatially-correlated participating media

All the rendering algorithms mentioned so far assume that the particles in the media have white-noise random distribution. However, media in the real-world may not follow this kind of distribution, but may be distributed according to certain rules.

This change in internal properties means that transmittance may no longer obey an exponential law. The earliest research on spatially-correlated participating media mainly concerned discrete media, such as sand (see Fig. 21). Recently, continuous spatially-correlated participating media have received attention, opening a new research direction. In this section, we first consider previous work on discrete media, and then discuss related work on continuous spatially-correlated participating media.

10.1 Discrete participating media

Moon et al. [69] proposed an importance sampling approach for path tracing in discrete media, via precomputation. In the precomputation stage, a shell function is used to represent the probability density function on a sphere around the center point. In the rendering process, importance sampling is guided by the precomputed shells to sample at a large distance in the discrete media, rather than via many small steps. This method reduces the

Fig. 21 Participating media in discrete space. Reproduced with permission from Ref. [68], © ACM 2015.

number of sampling steps and improves speed of convergence. It has been further improved by Meng et al. [68] to support multi-scale rendering: using path tracing with Ref. [69] for accurate rendering at small scales, and using diffusion theory to approximate at large scales. Later, Müller et al. [70] treated mixed media in multiple discrete participating media as equivalent continuous participating media to support mixed media rendering. In addition, they proposed a different multi-scale solution: a particle scattering distribution function for small scale rendering and the method of Moon et al. [69] for large scale rendering.

10.2 Continuous spatially-correlated participating media

In spatially-correlated participating media, particles no longer have white-noise random distribution, but follow certain laws, such as mutual attraction or repulsion caused by forces between particles. In a randomly distributed medium, light attenuates exponentially; in a mutually attracting medium, because there are relatively large gaps, the attenuation is slower than exponential; in a mutually repulsive medium, the opposite is true and attenuation is faster than exponential.

The generalized Boltzmann equation (GBE) from the field of neutron transfer [72] solves neutron transfer in non-exponential media based on several assumptions: an isotropic medium without boundaries, the phase function and scattering albedo of the medium do not depend on distance, and so on. The theory of light transport in continuous spatially-correlated participating media involved in computer graphics [71] eliminates several limitations on the basis of GBE, such as the no-boundary limit, so that the theory can be applied to rendering (see

Fig. 22). This method can simulate light transport in homogeneous spatially-correlated participating media. Although the proposed theory can support heterogeneous participating media, it does not provide a feasible solution, and does not explain the reciprocity and the path integral form, so it cannot be applied to any light transport method.

In a new form of path integration [73], non-exponential medium rendering is converted into the average over a series of participating media with different indexes, based on two hypotheses [72]. These are that the phase function of the medium and the scattering albedo do not depend on distance and the free-flights of photons, but only depend on the last scattering event. They lead to an approximate solution using averaging formulae for different participating media. In addition, this method guarantees reciprocity and supports the rendering of non-uniform participating media.

Guo et al. [74] proposed a general, physically-based framework for modeling and rendering such correlated media with non-exponential decay of transmittance (see Fig. 23). They described spatial correlations by introducing the fractional Gaussian field (FGF), a powerful mathematical tool that has proven useful in many areas but remains under-explored in graphics. With the FGF, they studied the effects of correlations in a unified manner, by modeling both high-frequency, noise-like fluctuations and k -th order fractional Brownian motion (fBm) with a stochastic continuity property. As a result, a wide variety of appearances stemming from different types of spatial correlations can be reproduced. Compared to the previous work, this method is the first that addresses both short-range and long-range correlations using physically-based fluctuation models. This method can simulate

Fig. 22 Results for randomly distributed, mutually attracted, and mutually repelling particles. Reproduced with permission from Ref. [71], © ACM 2018.

Fig. 23 Left: a complex scene containing several spatially-correlated media: the fractional Gaussian field (FGF) is able to reproduce a wide range of appearances stemming from short-range to long-range correlations and supports macroscopic heterogeneity. Right: reference generated by classical transport theory for comparison. Reproduced with permission from Ref. [74], © ACM 2019.

different extents of randomness in spatially-correlated media, resulting in a smooth transition in a range of appearances from exponential falloff to complete transparency.

11 Comparison and discussion

In Table 1, we compare several typical methods from several groups, in terms of the supporting scattering events, bias, storage cost, and rendering efficiency.

- *Scattering events.* Many lights based methods cannot handle single scattering, while others can handle both single and multiple scattering.
- *Bias.* Monte Carlo based methods produce unbiased results, while the others are mostly biased, except for the method of Bitterli and Jarosz [5], which is unbiased for multiple scattering when using a zero-order estimator, and Deng et al. [12] which is also unbiased when using a delta kernel.

- *Storage.* Density based approaches and point based approaches use more memory than others, where Bitterli and Jarosz [5] and Deng et al. [12] introduced higher-dimensional elements, e.g., photon planes and volumes, reducing memory cost compared to UPBP [1]. Many lights based methods require less memory than density based methods, since they are based on gathering rather than density estimation. Monte Carlo based methods usually do not need extra structures, resulting in low memory cost. When Monte Carlo based methods are accelerated with path guiding, more memory is required to store the learned lighting distribution.
- *Rendering speed.* Among these methods, point based methods are the fastest, and can achieve interactive frame rates. However, their rendering results have the lowest quality of all these methods. Monte Carlo based methods are the

Table 1 Comparison of approaches, considering supporting scattering events, performance, storage cost, etc. Unbiased (multiple) means only unbiased for multiple scattering and biased for single scattering

Method	Type	Scattering event	Biased/unbiased	Storage cost	Rendering efficiency
UPBP [1]	Density-based	Single + Multiple	Biased	●●●●●	●●○○○
Bitterli and Jarosz [5]	Density-based	Single + Multiple	Unbiased (multiple)	●●●●○	●●●○○
Deng et al. [12]	Density-based	Single + Multiple	Unbiased (delta kernel)	●●●○○	●●●○○
Wang and Holzschuch [4]	Point-based	Single + Multiple	Biased	●●●●○	●●●●●
VRL [2]	Many lights-based	Multiple only	Biased	●●●○○	●●●○○
MEMLT [36]	Monte Carlo-based	Single + Multiple	Unbiased	●○○○○	●○○○○
Wang et al. [23]	Monte Carlo-based	Multiple only	Biased	●●○○○	●●○○○
Herholz et al. [3]	Monte Carlo-based	Single + Multiple	Unbiased	●●○○○	●●○○○
Deng et al. [42]	Monte Carlo-based	Single + Multiple	Unbiased	●●○○○	●●○○○

slowest, since they have slow convergence, but they produce the highest quality. Path guiding based methods improve the convergence rate. Many lights based methods are less efficient than PBGI, but are still faster than density based methods and Monte Carlo based methods. Density based methods have different rendering cost. The methods of Bitterli and Jarosz [5] and Deng et al. [12] converge faster than UPBP, thanks to their higher-dimensional elements.

12 Challenges and future work

This article summarizes several recently proposed efficient rendering methods for participating media, including volume density estimation based approaches, point based approaches, Monte Carlo based approaches, etc. These methods improve the rendering efficiency of participating media in different ways.

At present, media rendering and surface rendering are relatively independent, but the two are closely coupled in the actual scene. Although there are some attempts to link the two and study their relationship, there is still no complete theory. Therefore, building a unified model of surface rendering and participating media has important theoretical and practical significance.

The participating media model introduced in this article is limited to homogeneous participating media, while participating media in the real world may be more complex, such as spatially-correlated participating media. In these participating media, the distribution of particles is no longer random, but is affected by a certain gravitational or repulsive force, which makes the previous theory no longer applicable. Some methods have been proposed to solve this problem, but these methods still have low rendering efficiency, thus, efficiently rendering spatially-correlated participating media has important practical significance.

At present, noise reduction based on deep learning has been tried extensively in surface rendering, and it is also effective. However, the performance of these methods in participating media is not satisfying. One important reason is that the characteristics of the participating media are very different from the characteristics of surface materials. How to design and make full use of the characteristics

of participating media to support noise reduction requires further research. In addition, some participating media have a strong global effect due to the scattering events, such as fog or smoke, and the current network architecture used for surface rendering will cause aliases on the participating media. Therefore, noise reduction for participating media requires a different network architecture.

References

- [1] Křvák, 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.
- [2] Novák, J.; Nowrouzezahrai, D.; Dachsbacher, C.; Jarosz, W. Virtual ray lights for rendering scenes with participating media. *ACM Transactions on Graphics* Vol. 31, No. 4, Article No. 60, 2012.
- [3] Herholz, S.; Zhao, Y. Y.; Elek, O.; Nowrouzezahrai, D.; Lensch, H. P. A.; Křvák, J. Volume path guiding based on zero-variance random walk theory. *ACM Transactions on Graphics* Vol. 38, No. 3, Article No. 25, 2019.
- [4] Wang, B. B.; Holzschuch, N. Point-based rendering for homogeneous participating media with refractive boundaries. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 10, 2743–2757, 2018.
- [5] 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.
- [6] Chandrasekhar, S. *Radiative Transfer*. New York: Dover Publications, 1960.
- [7] Kajiya, J. T. The rendering equation. *ACM SIG-GRAPH Computer Graphics* Vol. 20, No. 4, 143–150, 1986.
- [8] 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.
- [9] 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.
- [10] 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.

- [11] Jarosz, W.; Nowrouzezahrai, D.; Thomas, R.; Sloan, P. P.; Zwicker, M. Progressive photon beams. *ACM Transactions on Graphics* Vol. 30, No. 6, 1–12, 2011.
- [12] Deng, X.; Jiao, S. J.; Bitterli, B.; Jarosz, W. Photon surfaces for robust, unbiased volumetric density estimation. *ACM Transactions on Graphics* Vol. 38, No. 4, Article No. 46, 2019.
- [13] Qin, H.; Sun, X.; Hou, Q. M.; Guo, B. N.; Zhou, K. Unbiased photon gathering for light transport simulation. *ACM Transactions on Graphics* Vol. 34, No. 6, Article No. 208, 2015.
- [14] Novák, J.; Nowrouzezahrai, D.; Dachsbacher, C.; Jarosz, W. Progressive virtual beam lights. *Computer Graphics Forum* Vol. 31, No. 4, 1407–1413, 2012.
- [15] Keller, A. Instant radiosity. In: Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques, 49–56, 1997.
- [16] Walter, B.; Fernandez, S.; Arbree, A.; Bala, K.; Donikian, M.; Greenberg, D. P. Lightcuts: A scalable approach to illumination. *ACM Transactions on Graphics* Vol. 24, No. 3, 1098–1107, 2005.
- [17] Walter, B.; Arbree, A.; Bala, K.; Greenberg, D. P. Multidimensional lightcuts. *ACM Transactions on Graphics* Vol. 25, No. 3, 1081–1088, 2006.
- [18] Ou, J.; Pellacini, F. LightSlice: Matrix slice sampling for the many-lights problem. *ACM Transactions on Graphics* Vol. 30, No. 6, 1–8, 2011.
- [19] Hašan, M.; Křivánek, J.; Walter, B.; Bala, K. Virtual spherical lights for many-light rendering of glossy scenes. *ACM Transactions on Graphics* Vol. 28, No. 5, 1–6, 2009.
- [20] Novák, J.; Engelhardt, T.; Dachsbacher, C. Screen-space bias compensation for interactive high-quality global illumination with virtual point lights. In: Proceedings of the Symposium on Interactive 3D Graphics and Games, 119–124, 2011.
- [21] Arbree, A.; Walter, B.; Bala, K. Single-pass scalable subsurface rendering with lightcuts. *Computer Graphics Forum* Vol. 27, No. 2, 507–516, 2008.
- [22] Jensen, H. W.; Marschner, S. R.; Levoy, M.; Hanrahan, P. A practical model for subsurface light transport. In: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, 511–518, 2001.
- [23] Wang, B. B.; Ge, L. S.; Holzschuch, N. Precomputed multiple scattering for rapid light simulation in participating media. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 7, 2456–2470, 2020.
- [24] Frederickx, R.; Bartels, P.; Dutré, P. Adaptive LightSlice for virtual ray lights. In: Proceedings of the EG 2015 – Short Papers, 2015.
- [25] Yuksel, C.; Yuksel, C. Lighting grid hierarchy for self-illuminating explosions. *ACM Transactions on Graphics* Vol. 36, No. 4, Article No. 110, 2017.
- [26] Huo, Y. C.; Wang, R.; Hu, T. L.; Hua, W.; Bao, H. J. Adaptive matrix column sampling and completion for rendering participating media. *ACM Transactions on Graphics* Vol. 35, No. 6, Article No. 167, 2016.
- [27] Christensen, P. Point-based approximate color bleeding. Technical Report 08-01. Pixar Technical Notes. 2008. Available at <https://graphics.pixar.com/library/PointBasedColorBleeding/>.
- [28] Wang, B. B.; Gascuel, J.-D.; Holzschuch, N. Point-based light transport for participating media with refractive boundaries. In: Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations, 109–119, 2016.
- [29] Liang, Y. L.; Wang, B. B.; Wang, L.; Holzschuch, N. Fast computation of single scattering in participating media with refractive boundaries using frequency analysis. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 10, 2961–2969, 2020.
- [30] Novák, J.; Georgiev, I.; Hanika, J.; Jarosz, W. Monte Carlo methods for volumetric light transport simulation. *Computer Graphics Forum* Vol. 37, No. 2, 551–576, 2018.
- [31] Fong, J.; Wrenninge, M.; Kulla, C.; Habel, R. Production volume rendering: SIGGRAPH 2017 course. In: Proceedings of the SIGGRAPH '17: ACM SIGGRAPH 2017 Courses, Article No. 2, 2017.
- [32] Kajiya, J. T.; Herzen, B. P. V. Ray tracing volume densities. *ACM SIGGRAPH Computer Graphics* Vol. 18, No. 3, 165–174, 1984.
- [33] 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.
- [34] 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.
- [35] 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.
- [36] Jakob, W.; Marschner, S. Manifold exploration: A Markov Chain Monte Carlo technique for rendering scenes with difficult specular transport. *ACM Transactions on Graphics* Vol. 31, No. 4, Article No. 58, 2012.
- [37] Hanika, J.; Droske, M.; Fascione, L. Manifold next event estimation. *Computer Graphics Forum* Vol. 34, No. 4, 87–97, 2015.

- [38] Koerner, D.; Novák, J.; Kutz, P.; Habel, R.; Jarosz, W. Subdivision next-event estimation for path-traced subsurface scattering. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations*, 91–96, 2016.
- [39] Weber, P.; Hanika, J.; Dachsbacher, C. Multiple vertex next event estimation for lighting in dense, forward-scattering media. *Computer Graphics Forum* Vol. 36, No. 2, 21–30, 2017.
- [40] 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.
- [41] Meng, J.; Hanika, J.; Dachsbacher, C. Improving the dwivedi sampling scheme. *Computer Graphics Forum* Vol. 35, No. 4, 37–44, 2016.
- [42] Deng, H.; Wang, B. B.; Wang, R.; Holzschuch, N. A practical path guiding method for participating media. *Computational Visual Media* Vol. 6, No. 1, 37–51, 2020.
- [43] 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.
- [44] Müller, T.; Gross, M.; Novák, J. Practical path guiding for efficient light-transport simulation. *Computer Graphics Forum* Vol. 36, No. 4, 91–100, 2017.
- [45] 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.
- [46] 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.
- [47] Guo, J. J.; Bauszat, P.; Bikker, J.; Eisemann, E. Primary sample space path guiding. In: *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations*, 73–82, 2018.
- [48] Zheng, Q.; Zwicker, M. Learning to importance sample in primary sample space. *Computer Graphics Forum* Vol. 38, No. 2, 169–179, 2019.
- [49] Müller, T.; McWilliams, B.; Rousselle, F.; Gross, M.; Novák, J. Neural importance sampling. *arXiv preprint arXiv:1808.03856*, 2018.
- [50] Müller, T.; McWilliams, B.; Rousselle, F.; Gross, M.; Novák, J. Neural importance sampling. *ACM Transactions on Graphics* Vol. 38, No. 5, Article No. 145, 2019.
- [51] Hua, B. S.; Gruson, A.; Petitjean, V.; Zwicker, M.; Nowrouzezahrai, D.; Eisemann, E.; Hachisuka, T. A survey on gradient-domain rendering. *Computer Graphics Forum* Vol. 38, No. 2, 455–472, 2019.
- [52] Gruson, A.; Hua, B. S.; Vibert, N.; Nowrouzezahrai, D.; Hachisuka, T. Gradient-domain volumetric photon density estimation. *ACM Transactions on Graphics* Vol. 37, No. 4, Article No. 82, 2018.
- [53] Xu, Z. L.; Sun, Q.; Wang, L.; Xu, Y. N.; Wang, B. B. Unsupervised image reconstruction for gradient-domain volumetric rendering. *Computer Graphics Forum* Vol. 39, No. 7, 193–203, 2020.
- [54] Durand, F.; Holzschuch, N.; Soler, C.; Chan, E.; Sillion, F. X. A frequency analysis of light transport. *ACM Transactions on Graphics* Vol. 24, No. 3, 1115–1126, 2005.
- [55] Belcour, L.; Soler, C.; Subr, K.; Holzschuch, N.; Durand, F. 5D covariance tracing for efficient defocus and motion blur. *ACM Transactions on Graphics* Vol. 32, No. 3, Article No. 31, 2013.
- [56] Belcour, L.; Bala, K.; Soler, C. A local frequency analysis of light scattering and absorption. *ACM Transactions on Graphics* Vol. 33, No. 5, Article No. 163, 2014.
- [57] 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.
- [58] 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.
- [59] Ge, L. S.; Wang, B. B.; Wang, L.; Holzschuch, N. A compact representation for multiple scattering in participating media using neural networks. In: *Proceedings of the ACM SIGGRAPH 2018 Talks*, Article No. 16, 2018.
- [60] Walter, B.; Zhao, S.; Holzschuch, N.; Bala, K. Single scattering in refractive media with triangle mesh boundaries. *ACM Transactions on Graphics* Vol. 28, No. 3, Article No. 92, 2009.
- [61] Holzschuch, N. Accurate computation of single scattering in participating media with refractive boundaries. *Computer Graphics Forum* Vol. 34, No. 6, 48–59, 2015.
- [62] Wang, B. B.; Hašan, M.; Yan, L.-Q. Path cuts: Efficient rendering of pure specular light transport. *ACM Transactions on Graphics* Vol. 39, No. 6, Article No. 238, 2020.
- [63] Sun, B.; Ramamoorthi, R.; Narasimhan, S. G.; Nayar, S. K. A practical analytic single scattering model for real time rendering. *ACM Transactions on Graphics* Vol. 24, No. 3, 1040–1049, 2005.

- [64] Kallweit, S.; Müller, T.; McWilliams, B.; Gross, M.; Novák, J. Deep scattering: Rendering atmospheric clouds with radiance-predicting neural networks. *ACM Transactions on Graphics* Vol. 36, No. 6, Article No. 231, 2017.
- [65] Ge, L. S.; Wang, B. B.; Wang, L.; Meng, X. X.; Holzschuch, N. Interactive simulation of scattering effects in participating media using a neural network model. *IEEE Transactions on Visualization and Computer Graphics* Vol. 27, No. 7, 3123–3134, 2021.
- [66] Vicini, D.; Koltun, V.; Jakob, W. A learned shape-adaptive subsurface scattering model. *ACM Transactions on Graphics* Vol. 38, No. 4, Article No. 127, 2019.
- [67] Guo, J.; Li, M. T.; Li, Q. W.; Qiang, Y. T.; Yan, L. Q. GradNet: Unsupervised deep screened Poisson reconstruction for gradient-domain rendering. *ACM Transactions on Graphics* Vol. 38, No. 6, Article No. 223, 2019.
- [68] Meng, J.; Papas, M.; Habel, R.; Dachsbacher, C.; Marschner, S.; Gross, M.; Jarosz, W. Multi-scale modeling and rendering of granular materials. *ACM Transactions on Graphics* Vol. 34, No. 4, Article No. 49, 2015.
- [69] Moon, J. T.; Walter, B.; Marschner, S. R. Rendering discrete random media using precomputed scattering solutions. In: *Proceedings of the Eurographics Symposium on Rendering*, 2007.
- [70] Müller, T.; Papas, M.; Gross, M.; Jarosz, W.; Novák, J. Efficient rendering of heterogeneous polydisperse granular media. *ACM Transactions on Graphics* Vol. 35, No. 6, Article No. 168, 2016.
- [71] 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.
- [72] Larsen, E. W.; Vasques, R. A generalized linear Boltzmann equation for non-classical particle transport. *Journal of Quantitative Spectroscopy and Radiative Transfer* Vol. 112, No. 4, 619–631, 2011.
- [73] Bitterli, B.; Ravichandran, S.; Müller, T.; Wrenninge, M.; Novák, J.; Marschner, S.; Jarosz, W. A radiative transfer framework for non-exponential media. *ACM Transactions on Graphics* Vol. 37, No. 6, Article No. 225, 2018.
- [74] Guo, J.; Chen, Y. J.; Hu, B. Y.; Yan, L. Q.; Guo, Y. W.; Liu, Y. T. Fractional Gaussian fields for modeling and rendering of spatially-correlated media. *ACM Transactions on Graphics* Vol. 38, No. 4, Article No. 45, 2019.

Wenshi Wu is a master student at Nanjing University of Science and Technology. She received her bachelor degree from the Department of Applied Mathematics at the same university in 2021. Her research interests are mainly in rendering participating media.

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

Ling-Qi Yan is an assistant professor of computer science at UC Santa Barbara, co-director of the MIRAGE Lab, and affiliated faculty in the Four Eyes Lab. Before that, he received his doctoral degree from the Department of Electrical Engineering and Computer Sciences at UC Berkeley and obtained his bachelor degree in computer science from Tsinghua University. His research interests include physically-based rendering, real-time ray tracing, and realistic appearance modeling.

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>.