

A Survey on Homogeneous Participating Media Rendering

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

© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract Participating media are frequent in real-world scenes, whether it is milk, fruit juices, oil or muddy water in river or ocean scenes. Incoming light interacts with these participating media in complex ways: refraction at the boundaries and scattering and absorption inside the volumes. The radiative transfer equation is key point to solve this problem. There are several categories of the rendering methods which are all based on radiative transfer equation, but with different solutions. In this paper, we introduce these groups, more specifically, including: volume density estimation based approaches, virtual point / ray / beams 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. We discuss these methods, the challenges and open problems in this research direction.

Keywords Participating media; Monte Carlo based methods; Rendering; Volume density estimation.

1 Introduction

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




Fig. 1 Different types of participating media. (a) Figure reproduced from [42]. (b) Participating media from “Finding Dory” ©Disney Enterprises, Inc. and Pixar, Inc. (c) Participating media from “Frozen” ©Disney Enterprises, Inc.

The light is potentially refracted by the boundary if the indices of refractions differ between inside and outside the media and then absorbed or scattered as it travels inside the medium. The refracted boundary gathers the light in the media and causes high frequency effects, called volumetric caustics, which are obvious in media with relatively large mean free path. The scattering effects blur incident light and cause low frequency effects, which happen in media with small mean free path. Directional phase functions and refraction at the interface add to the computational complexity. This 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 many lights based method (Virtual Ray Lights, VRL) [57], several extensions to photon mapping culminating with Unified Points, Beams and Paths (UPBP) [42], Monte Carlo based methods [25] or point based methods [71]. All these methods greatly improve simulation of participating media. In this paper, we will discuss each of these categories from Section 3 to Section 6 and then will show some acceleration techniques in Section 6.4. Several approaches focus on single scattering only, or volume caustics, which will be shown in Section 8. More recently, deep learning has been exploited in participating media rendering, and

1 School of Computer Science and Engineering, Nanjing University of Science and Technology, China. E-mail: wws19990401@outlook.com, beibei.wang@njust.edu.cn.

2 University of California, Santa Barbara, US E-mail: lingqi@cs.ucsb.edu.

Manuscript received: 2014-12-31; accepted: 2015-01-30.

we will show some related work with deep learning in Section 9.

In the classical participating media, the particles are assumed to have a white noise random distribution. However, the particles in the media might have some kinds of forces, which leads to positive / negative random distribution. Spatial correlated media have been introduced and studied in the rendering domain in recent years, and we will show the progress on this topic in Section 10.

We will discuss challenges and the 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. This participating media is described by absorption coefficient σ_a , scattering coefficient σ_s and phase function $p(\omega, \omega_t)$, which represent the absorption ratio, scattering ratio and the distribution of outgoing directions for scattering respectively. The sum of the absorption coefficient and scattering coefficient is called attenuation coefficient, represented as $\sigma_t = \sigma_a + \sigma_s$. Another equivalent expression for a media 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 media.

According to these properties, we generally classify participating media into high-order scattering dominant and low-order scattering dominant media, considering whether they are optically thick or thin, according to 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 separate among 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 corresponds to paths with two scattering events, and multiple scattering corresponds to paths with more than two scattering events. The single scattering leads to high-frequency effects, and the multiple scattering leads to low-frequency effects.

Thus the separation between these scattering events is reasonable. Both high-order scattering dominant and low-order scattering dominant media are challenging. Regarding the 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 we used a lot of computational power for an almost constant appearance. The low-order scattering dominant media usually introduces high-frequency volumetric caustics effects, due to the presence of double refraction, which are difficult to capture with either density estimation based approaches or Monte Carlo based approaches. Virtual point, ray or beam methods can not simulate single scattering. Thus, a specific group of methods is proposed to simulate single scattering only.

2.2 Radiative Transfer Equation




Fig. 2 The radiative transfer equation.

Light transport within participating medium is described by the Radiative Transfer Equation [8], which defines the radiance that reaches a point x from direction ω as a sum of exited radiance from the nearest surface and this direction and in-scattered radiance from the medium among the whole length of the ray, as shown in 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

$$T_r(x \leftrightarrow x_s) = \exp^{-\sigma_t \|x - x_s\|}, \quad (2)$$

s is the distance through the medium to the nearest surface x_s , and x_t is a point and its distance to surface x_s is between 0 and s . $L(x_s, \omega)$ can be governed from the rendering equation [37]. $L_i(x_t, \omega)$ is the in-scattering radiance at x_t from all direction ω_t over the sphere of directions $\Omega_{4\pi}$ using the phase function p , defined as

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

Similar to rendering equation, the radiative transfer equation does not have any analytic solutions. The solutions to the radiative transfer equation include volumetric density estimation based, point based, many




Fig. 3 Different estimators and their results. Figure reproduced from [6].

lights based or Monte Carlo based methods. We will discuss them in the following sections.

All 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 usually are faster than Monte Carlo based methods, at the cost of the extra storage. Regarding the scattering types, many lights based methods can not 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 ([6]) based methods can produce unbiased results, while the others are biased. We will discuss these methods in detail in the following sections.

3 Volumetric Density Estimation

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 that distribute light photons, rays, or beams in the media, and then estimate the density of these elements in a kernel size. In general, the group of methods includes two passes: a lighting pass and a rendering pass. In the lighting pass, the rays are shot into the media and get scattered until having negligible energy or reaching the 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 types of the representations. The reusing of the cached elements for all pixels makes it efficient.

3.1 Different estimators

Jensen et al. [35] generalized photon-mapping algorithm from light travel among 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 sample with density estimation with a 3D kernel. The estimator is a point-to-point estimator. Later, Jarosz et al. [34] improved this algorithm, by gathering the contribution of photons map along with camera rays rather than camera samples, speeding up the convergence. The estimator is a point-to-beam estimator. Furthermore, Jarosz et al. [32] 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 [33] is proposed to eliminate memory limitation. The huge volumetric photon map is replaced by a lot of 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 it's discarded. This progressive strategy greatly improves the practicality of these work.

Křivánek et al. [42] found that different




Fig. 4 Contributions of the different estimators to the UPBP image of the Still life scene. Figure reproduced from [42].

representations of points, beams, paths, etc. are suitable for light transport of different participating media or surfaces. The photons are suitable for high-order scattering objects, and beams are suitable for low-order scattering objects, while path is suitable for volume effects evolved 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 the light transportation of any type of participating media, by automatic choosing the suitable representation, as shown in Fig. 4



Fig. 5 A comparison of the equal-time variance of different estimators in a scene containing participating media. This figure shows the full light transport in the scene (left), single scattering (middle/right, top half) and multiple scattering volumetric transport (middle/right, bottom half). Those estimators (middle) provide significant variance reduction compared to prior density estimators (right) at equal render time. Figure reproduced from [11].

Note that, all of these previous methods are biased, although they are consistent. A recent work by Bitterli et al. [6] further has improved the convergence, by extending to higher dimensional expressions: photon plane and photon volume. They further improved the convergence efficiency and achieved unbiased results by using a zero order estimator. The different types of estimators are shown in Fig. 3. 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. [11] solved these issues by generalizing photon

planes to photon surfaces, results in different types of estimators, including new “photon cone”, “photon cylinder”, “photon sphere”, and multiple new “photon plane” estimators. The estimators are combined with multiple importance sampling to increase robustness for arbitrary types of participating media. Furthermore, they proposed a delta kernel to couple the light and camera subpath to avoid bias. Compared to the prior work ([6]), their method significantly reduces variance and is able to handle any scattering event, including single scattering (see Fig. 5). However, this method is not able to handle medium interactions immediately following scattering from a surface, e.g. volumetric caustics.

Qin et al. [61] introduced an unbiased photon gathering for participating media. They combined each photon’s path with camera path into one path. They considered all factors in Monte Carlo method such as visibility and probability to ensure unbiasedness.

Discussion. Density estimation based approaches are the most commonly used solutions for participating media rendering, since it supports all types of scattering events and is able to produce high quality results. However, there are several limitations of this group of methods. First, most methods in this group are biased, except the zero order kernel in Bitterli et al. [6] and Deng et al. [11]. Second, this group of methods requires an extra light pass and extra storage for the light distribution. Third, for high-order scattering dominant media, they require long time to converge, and for single scattering, density estimation based methods requires a large number of elements to simulate high-frequency volumetric caustics, as insufficient photons / rays / beams yields blurry caustics.

3.2 Combination with Other Techniques




Fig. 6 Virtual point light methods (left) convert the vertices of a random-walk pre-process into a collection of virtual point lights (VPLs). VRLs (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 (bottom). Figure reproduced from [57].



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 render effectively eliminates singularities in four indirect light transport paths (bottom) using novel importance schemes and a new lighting primitive compatible. Figure reproduced from [56].

4 Virtual Point / Ray / Beam Lights Methods

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

Keller et al. [40] proposed instant radiosity, which is also called Virtual Point Lights (VPLs) and 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 less than number of photons, since visibility computation is required for VPLs. VPLs have been widely used for indirect illumination. Several improvements [59, 66, 67] have been made to accelerate the visibility computation, using hierarchical pruning. However, these methods suffer from singularity issue, which comes from the potential tiny distance between the VPLs and the shading points, yielding spike artifacts in the rendering results. These artifacts could be removed by clamping or blurring [23], at the cost of bias or energy missing. The bias issue has been compensated approximately [55].

Later, VPLs have been extended for participating media. Arbree et al. [2] combine lightcuts with diffusion dipole [36] to approximate subsurface scattering. Multidimensional lightcuts [66] extends VPLs to participating media and other effects (e.g. motion blur), thanks to its efficient hierarchy pruning. However, it still suffers from the singularity issue for multiple scattering events. Novák et al. [57] replaced virtual point lights with virtual ray lights (VRLs) to simulate 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 treat as a line-to-line double integral problem. This approach produces higher quality than VPLs. However, it still suffers from singularity issue. Later, virtual beam lights (VBLs) with finite thicknesses are 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, a lot of virtual ray/beam lights are required, even with progressive solution, long time is required to converge. Thus, Wang et al. [3] proposed a precomputed solution to improve the convergence in high-order scattering media, by precomputing the scattering events in an infinite participating media (see Fig. 13). The precomputed scattering events are stored in two tables for both point light source and ray light source. Thanks to the symmetry of revolution around the direction of propagation and the almost symmetric shape of the lobes, the dimension of precomputed distribution for a single media is reduced to three dimensions. In the lighting stage, only the rays after the 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. Besides 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 the visibility in the media, limits to homogeneous translucent media and only handles multiple scattering and relegates single scattering to other rendering methods.

On the contrary, Georgiev *et al.* [18] 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 the lightcuts to organize the virtual point lights, the virtual ray lights can also be organized into a certain hierarchical structure for pruning VRLs. Frederickx *et al.* [14] clustered the VRLs into a series of ray slices in the precomputation pass and estimate the required cluster number using sufficient analysis of variance in the rendering pass, resulting in faster convergence. Yuksel and Yuksel [74] treated the self-illumination in explosion rendering as the lighting problem in VPLs, and organized these VPLs into a hierarchy. In addition, multiple scattering was precomputed in the hierarchy, resulting in efficient explosion rendering.

To solve the expensive visibility computation of a large number of virtual light source in participating media, Huo *et al.* [28] proposed a sparse sampling and reconstruction method, based on the smooth characteristics of the multiple scattering in the participating media. The virtual light sources are organized into a small number of representative points using clustering. The 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.

Discussion. The VRL (VBL) based participating media rendering methods are faster than density estimation based methods, since much less rays or photons are required in this types of method. However, the visibility computation is still expensive, while

the density estimation based methods do not have this issue. Also, they are not able to handle single scattering, and rely on density estimation based methods to compute it. Furthermore, they have singularity issue, especially in high-order scattering media, although this issue has been reduced by many techniques. All the VRL based methods are biased, while some of the density estimation based methods are unbiased.

5 Point based Methods

Point-based global illumination (PBGI) algorithm [9] was first proposed to compute the diffuse light transport for surface rendering and widely used in movie production, since it's noise-free and much faster than Monte Carlo based methods. Similar to 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 the geometric information, which serves as an approximation of the geometry representation and will be used for visibility computation with rasterization in the rendering pass.

Wang *et al.* [69] expanded PBGI to participating media rendering. Besides the points on the surface, they are also placed inside the medium. The volume samples including geometric information, using bounding boxes. The volume and surface samples are organized into hierarchies respectively. During rendering, single-, double- and multiple- scattering are computed separately (Fig. 8). Single scattering is computed directly, finding light samples are closer to the camera ray. To compute double scattering, they traversed the spatial hierarchy to obtain the best tree cut and gathered the contributions from these nodes. For multiple scattering, they used a precomputed table to store the resulting contribution. They then added the contributions from single, double and multiple scattering. For indirect lighting and multiple scattering after several bounces on the refractive surface, they used 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 albedo to high albedo, and for 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 glossy reflections). Wang and Holzschuch [71] further improved it, with a faster single scattering tighter bounding boxes, and a GPU implementation to support

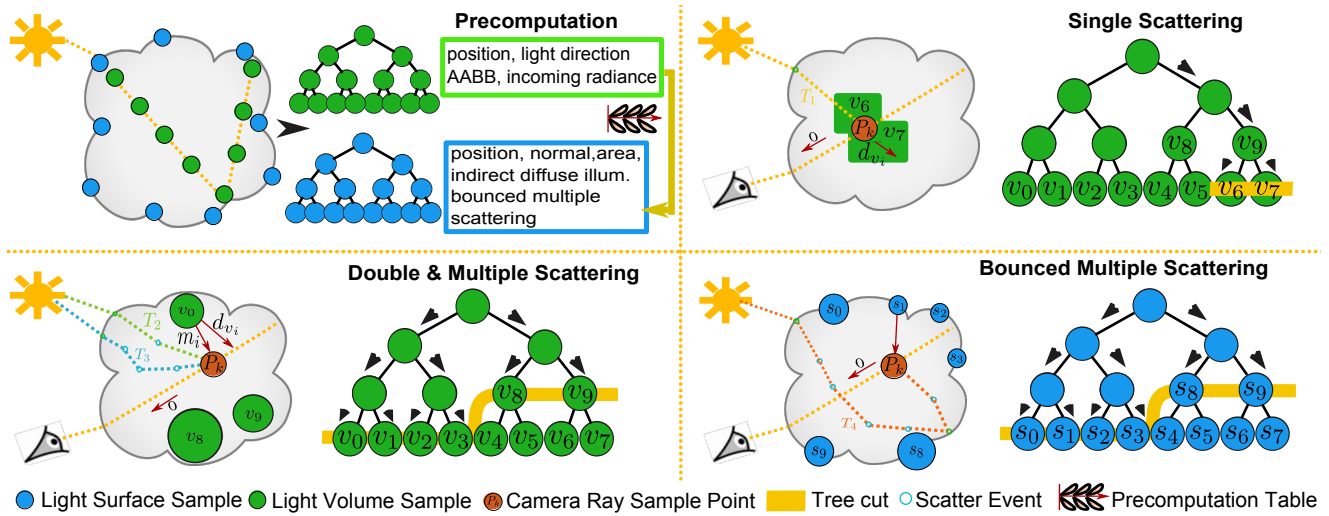


Fig. 8 The framework of PBGI in participating media rendering. The algorithm: it begins by computing incoming light at volume and surface samples. Then it computes Single-, Double- and Multiple scattering effects for each camera ray using these volume and surface samples. Figure reproduced from [71].

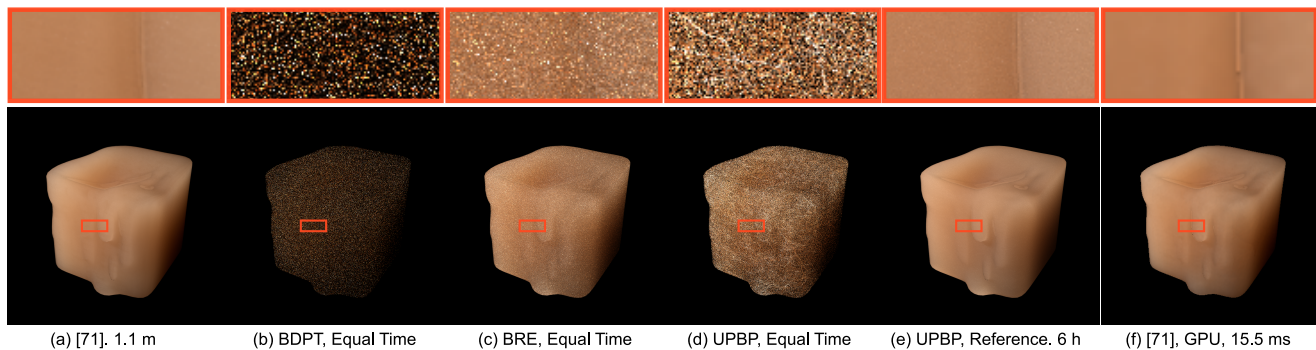


Fig. 9 Comparison between [71] and other volumetric methods on a wax media. For this material, with a large albedo α and a small mean free path ℓ , multiple scattering effects dominate. Figure reproduced from [71].

media interactive rendering and editing. The rendered results for wax media is shown in Fig. 9.

Recently, Liang et al. [46] introduced the frequency analysis theory to single scattering computation in PBGI. Since the single scattering is usually high frequency, and a large number of volume samples are required to produce the sharpness of the 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 less volume samples.

Discussion. The 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 acting on the number of samples. Differ from the density estimation-based methods, the points in PBGI are carefully distributed

and include geometric information, thus they are able to produce high-quality single scattering result with much less points and shorter time than density estimation-based methods. Regarding the multiple scattering, the precomputed multiple scattering table avoids the tracing of scattering events and helps to decrease the converge time significantly. Compared to many-light based methods, there are more points used in PBGI to capture the high-frequency in the single scattering, while many-light based methods cannot simulate single scattering. The main limitation for this group of algorithms is that homogeneous materials is assumed. Extension to heterogeneous materials, with spatially varying scattering properties, will require future 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 are occlusions. Furthermore, the

point based methods are also biased. In summary, point based-methods are suitable for rendering applications with 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 [58] and production [13].

Monte Carlo based rendering was first proposed for forward path tracing integration [38], which relies on two sampling operations: distance sampling and phase function sampling. The distance sampling usually considers the transmittance attenuation along the media, thus is performed exponentially to the distance. For high-order scattering media, millions of scattering events happen before leaving the objects, due to the small step with the distance sampling which is relative to the mean free path. A highly anisotropic medium requires the sampling of a high frequency phase function, resulting in noisy result, even with importance sampling. When the medium is enclosed in a refractive boundary, the path sampling is even more difficult with unidirectional path tracing. Lafortune and Willems [44] later expanded it to bidirectional path tracing, which improves the convergence rate of light transport with refractive boundaries. Pauly et al. [60] proposed the Metropolis light transport (MLT) approach for participating media.

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

6.1 Advanced Sampling

Joint importance sampling [17] 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 is 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, thus it's not suitable for high-order scattering media.

6.2 Next Event Estimation

Jakob and Marschner [29] proposed a new mutation strategy for metropolis light transport based on manifold exploration, improving the sampling for paths involving specular and highly glossy surfaces. It can also be used for participating media, especially for media surround by refractive boundaries. The manifold exploration idea further benefits the next event estimation (NEE).

NEE is usually used in Monte Carlo (MC) rendering to reduce variance, via 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. [22] proposed Manifold next event estimation (MNEE), which leverages manifold exploration [29] 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 the geometric derivative, it does not handle detailed or displaced surfaces well. It was improved by Koerner et al. [41], via searching 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. [72] proposed Multiple Vertex Next Event Estimation (MVNEE), which connects the point to the light source with sub-path generated by perturbation, instead of one segment connections. The sub-paths are generated by perturbing seed paths generated with path tracer. This approach is proposed for the multiple scattering in a high-order scattering homogeneous participating media, and it could be combined with path tracer via multiple importance sampling. This method significantly reduces noise and increases performance of multiple scattering renderings in highly anisotropic, optically dense media, but it is not suitable for participating media with boundaries.

6.3 Zero-variance Random Walk

Zero-variance random walk path means creating random walk without variance. This means that 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's impossible to sample such a path, since it's hard to obtain the incoming radiance distribution. Even with path guiding, the incoming radiance distribution is an approximation, which can not guarantee zero variance. Thus the present work is only for some simple situation. For high scattering and isotropic participating media, Dwivedi sampling [43] biases the sampling probability distributions to exit the medium as quickly as possible, based on the idea that the path close to boundary has higher contribution than path deep into the body. They approximate the geometry via locally fitting a slab to compute an analytic approximation of the light field for Dwivedi sampling. This method leads to less variance, however, this method is not suitable for thin geometries (such as ears) with a strong backlight.

Later, Meng et al. [48] solved this issue by taking the geometric characteristics of the object into account in the sampling process, so as 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 one searches for the closest point to the boundary at every scattering vertex to increase the chances to escape thin geometry. The second one 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 suitable for high-order scattering and isotropic participating media.

6.4 Path Guiding

Path guiding was first introduced for surface rendering [24, 52, 62, 65]. 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 with Bidirectional Reflection Distribution Function with multiple importance sampling [52] or product importance sampling [24]. Besides path guiding methods in path space, several works have focused on the primary sample space [19, 53, 75]. Deep learning has also been used in path guiding. Müller et al. [54] proposed a deep neural network model to present the probability density function of samples. The learned model is leveraged for sampling the ray




Fig. 10 Comparison between multiple importance sampling and RIS of two function f_1 and f_2 product. The RIS produces result much closer to the target function (product of two functions) than MIS. Figure reproduced from [10].

outgoing direction. The advantage of this work is independence from scenes and local point parameters (e.g. including textured BRDF). However, this method is biased and the sampling is very expensive. Path guiding makes convergence faster compared to the original path tracing.

Later, path guiding has been introduced for volume rendering. Deng et al. [10] extended [52] to translucent materials. They also used a spatial-directional tree to represent the incident radiance distribution and then sample the outgoing direction after a scattering considering both this learned lighting distribution and the phase function. More specifically, they introduced the resample importance sampling (RIS) to joint 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. The RIS is combined with MIS, depending on the sharpness of the phase function to save the computational cost in RIS. Fig. 10 shows the sampling results of both RIS and MIS of 2D functions. As demonstrated in the paper, the RIS is also able to improve the results in the surface rendering, 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, from isotropic media to highly anisotropic media. Unfortunately, this approach uses




Fig. 11 Top: An example path containing volume (blue) and surface (yellow) vertices. Note that, in the path tracing algorithm, the path carries visual importance, and thus is generated in the opposite direction of the flow of light. Bottom: The four considered zero-variance sampling decisions. 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. Figure reproduced from [25]




Fig. 12 Comparison between standard sampling and path guiding from Herholz et al. [25]. Figure reproduced from [25].

the naive distance sampling, which has a significant impact on the convergence speed.

Herholz et al. [25] proposed a path guiding approach using zero-variance path sampling theory, which is able to guide all the sampling decisions (Fig. 11), including distance sampling, direction sampling, Russian roulette and splitting. The zero sampling is guided by a cached estimate of the adjoint transport solution, which is represented as a kd-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 with vMF approximation, and then sampled this product (see Fig. 12). Compared to Deng et al. [10], sampling the product of two functions should have higher quality than resample importance sampling of the two functions, although the product operation requires extra computational cost. The vMF representation in the spherical domain makes the product operation easier than the quad tree used in Deng et al. [10]. With all the sampling decisions

guided by the cached estimate of the adjoint transport solution, Herholz et al. [25] led to significantly faster convergence compared to Deng et al. [10] and 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.

Discussion. 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 the types of scattering, including single and multiple scattering. However, these methods usually take long time 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 groups of methods, Monte Carlo based methods still have a slow convergence. Thus, the targeting application of Monte Carlo based methods are high quality rendering applications with enough 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 could be further improved by combining with acceleration techniques, like gradient domain rendering, frequency analysis, radiance caching and precomputation.

7.1 Gradient-based Rendering

The gradient-based method [27] was introduced for the homogenous participating media by Gruson et al. [1], taking advantage of the smoothness characteristics of the participating media's rendering




Fig. 13 The original path is replaced with a special path. Figure reproduced from [3].

results. 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 the 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 [73], yielding a better reconstruction solution than Poisson reconstruction. We will describe this approach in Section 9.

7.2 Frequency Analysis

Frequency analysis on light transport has been introduced for prefiltering and adaptive sampling ([12],[5]) in surface rendering. Later, Belcour et al. [4] introduced it to participating media, based on the observation that multiple scattering leads to smooth appearance. They proposed the 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 matrix, they proposed several improvements: adaptive sampling in the image space and during the contribution gathering among the camera and the light samples, etc., resulting in faster convergence of photon beams approach [33].

7.3 Radiance Caching

Radiance caching method [31] uses a spherical harmonic function to represent and store radiance distribution of some sparse points and related gradient information. The illumination at a new position is calculated through the existing distribution and gradient to complete the interpolation, thereby reduces the amount of calculation and getting smoother effect. Since the stored sampling points are sparse, the points to be calculated may be different from the stored points due to different occlusions, resulting in different illumination, and the previous method did not consider the visibility problem during interpolation. This problem was solved by Marco et al. [47], by using the second-order Hessian error metric to determine whether

the interpolation error is acceptable. Thanks to the visibility, it produces results with higher quality than Jarosz et al. [31].

7.4 Precomputation

High-order scattering media can be rendered by density based methods (e.g. UPBP), many-light based methods (e.g. VRL) and Monte Carlo based methods (e.g. MEMLT). However, all of these methods have very low convergence rate. This issue was solved by Wang et al. [3], via 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 the mutation. For a seed path generated with path tracing, the sub-paths within the media are replaced with a special path, which is a virtual path with a converged contribution computed from the precomputation table. During mutation, only the surface events and the endpoints of the special path could be mutated, thus the path lengths are greatly decreased using this special path, and much less mutation count is required, since all vertices within the special path will not be touched. This approach results in much faster convergence than the original MEMLT, especially for high-order scattering approach. Besides MEMLT, this method is also suitable for other bidirectional tracing methods, like VRL and UPBP. Recently, Ge et al. [15] 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.




Fig. 14 (a) Multiple paths (black lines) may connect light L to scatter point V, while the non-refractive approximation (purple line) cannot account for these paths. (b) Geometry of problem at a point P with geometric normal \hat{N}_g and shading normal \hat{N}_s . Figure reproduced from [68].

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, the 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. [68] introduced a method for accurate computation of single scattering effects in participating media under point light source. Under this configuration, it's impossible to render with volumetric path tracing or MEMLT. They treated single scattering computation problem as a searching problem: finding all the points on the surface that connect the point light and the camera sample which satisfy Fermat's principle. The presence of shading normal of the surface points makes the problem more complex (see Fig. 14). They computed these entry points on the surface using Newton-Raphson optimization. With the interval Newton, all the solutions can be found. However, to capture the sharpness of the caustics, a large number of camera samples are required along the camera ray. Similar to this method, Wang et al. [70] extended it to multiple-bounce pure specular light transport computation, however, it can also be used for single scattering computation. It produces similar results as Walter et al. [68].

Holzschuch [26] improved it by computing the extent of the influence of each triangle over 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 edges. This method is significantly faster than Walter et al. [68] while providing higher quality results (see Fig. 15).

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

Both of the above methods can compute accurate single scattering without any noise. However, computation time for this method depends strongly on scene complexity. Moreover, these methods are limited to participating media under point lights.

9 Participating Media Rendering Based on Neural Networks

Neural network has also been leveraged in participating media rendering, including atmospheric clouds rendering, multiple scattering representation, BSSRDF models and reconstruction of gradient-domain volumetric rendering.

9.1 Atmospheric Cloud Rendering

Atmospheric cloud is high-order scattering media, and it requires very long time to simulate the multiple scattering event. Kallweit et al. [39] proposed to render atmospheric clouds using a radiance-predicting neural networks (RPNN) for multiple scattering. 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 the cloud density, as shown 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 path tracing, but with thousand times of speedup, as shown in Fig. 17.

9.2 Multiple Scattering Representation

Wang et al. [3] proposed to use 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

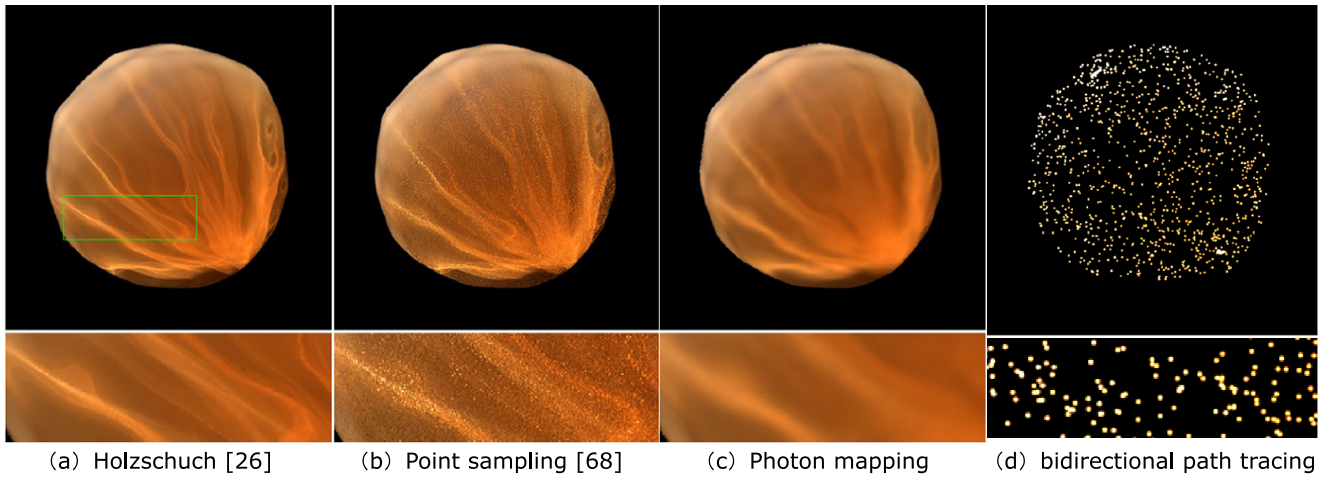


Fig. 15 Comparison among, Holzschuch [26], Walter et al. [68], photon mapping and path tracing. Figure reproduced from [26].




Fig. 16 Stencil grid (left) and progressively training network (right). Figure reproduced from [39].

represent the entire media space, further considering the albedo and phase function. To solve this issue, Ge et al. [16] proposed to use 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 used this neural network to represent double scattering and multiple scattering, and left single scattering with other methods. The storage for both double and multiple scattering is reduced from 50 GB to 23.6KB. During rendering, they reproduced the precomputed table for a specific media from the neural network, and then used the precomputed table for multiple scattering computation, similar to Wang et al. [3]. They integrated it into virtual ray light (VRL) with an efficient GPU implementation, resulting in interactive frame rate, as shown in Fig. 19.




Fig. 17 Comparison of path-traced references and RPNN approach. Figure reproduced from [39].

Vicini et al. [64] introduced a shape-adaptive BSSRDF model via 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 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 proposes to use 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




Fig. 18 Network structure of Ge *et al.* [16]. Figure reproduced from [16].

performs regression fitting on the scale factor of the distribution of multiple scattering. Fig. 20 shows the structure of the three networks. This method supports arbitrary homogeneous media parameters, which maintains 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.* [73] 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 [21] with an encoded shift connection and a separated auxiliary feature branch, which includes volume based auxiliary features such as transmittance and photon density. This network smooths images on global scale and preserves the high frequency details on a small scale. Their network produces a higher quality result, compared to previous works (L_1 or L_2 or GradNet). Although they only considered volumetric photon mapping, it's straightforward to extend this method for other forms, like beam radiance estimation. However, this method can not reconstruct images with a lot of noise, and tends to over-blur some features.

10 Spatially-correlated Participating Media

All the rendering algorithms mentioned in the prior sections assume that the particles in the media have white-noise random distribution. However, the media in the real-world might not follow this kind of distribution, but are distributed according to certain rules. This change in internal properties determines that the transmittance no longer obeys the exponential

law. The earliest research on spatially-correlated participating media was mainly for discrete media, such as sand (see Fig. 21). Recently, the study on continuous spatially-correlated participating media has got attention, opening a new research direction. In this section, we first show previous work on the discrete media, and then show the related work on continuous spatially-correlated participating media.

10.1 Discrete Participating Media

Moon *et al.* [51] proposed an importance sampling approach for path tracing in discrete media, via precomputation. In the precomputed stage, a shell function is used to represent the probability density function on the sphere around the center point. In the rendering process, importance sampling is guided by the precomputed shells to sample a large distance in the discrete media, rather than a lot of small steps. This method reduces the number of sampling steps and improves the convergence efficiency. It's further improved [49] to support multi-scale rendering: using path tracing with [51] for accurate rendering at small scales, and using the diffusion theory to approximate at large scales. Later, Muller *et al.* [50] treated mixed media of 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 Moon *et al.* [51] for large scale rendering.

10.2 Continuous Spatially-correlated Participating Media

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

The Generalized Boltzmann Equation (GBE) in the field of neutron transfer [45] solves the neutron transfer in non-exponential media based on several assumptions. These assumptions include: 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



Fig. 19 Comparison of Ge et al. [16] and UPBP. Figure reproduced from [16].

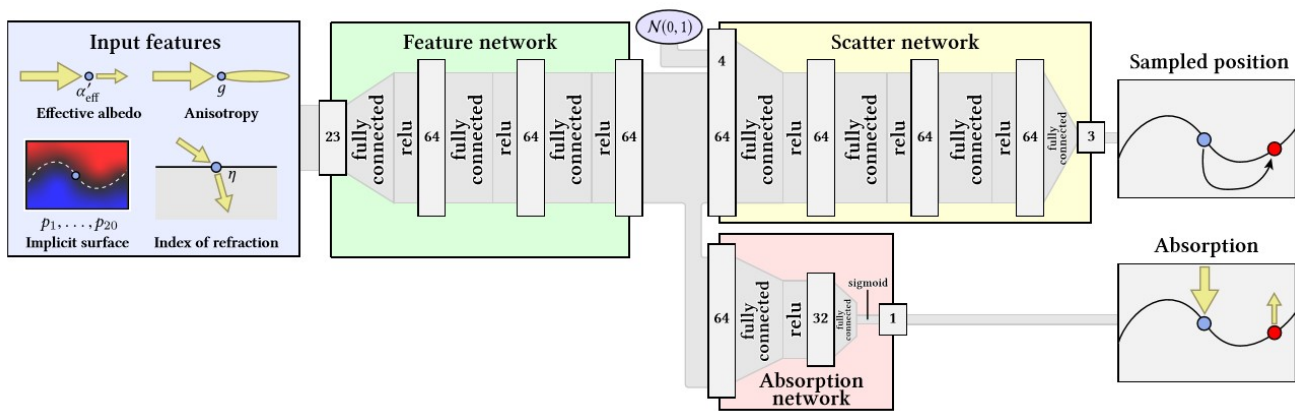


Fig. 20 The network structure of shape-adaptive BSSRDF model. Figure reproduced from [64]




Fig. 21 Participating media in discrete space. Figure reproduced from [49].




Fig. 22 The result of randomly distributed, mutually attracted and mutually repelling particles. Figure reproduced from [30].

involved in computer graphics [30] 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 the light transport of homogeneous spatially-correlated participating media. Although the proposed theory can support non-uniform participating media, it does not

provide a feasible solution, and does not explain the reversibility and the path integral form, so it cannot be applied to any light transport method.

A new form of path integration [7]: the non-exponential medium rendering is converted into the average of a series of different index participating media, in which two hypotheses [45] are introduced. These hypotheses 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 the approximate solution of a series of average formulas for different participating media. In addition, this method guarantees reversibility and supports the rendering of non-uniform participating media.

Guo et al. [20] 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




Fig. 23 A complex scene containing several spatially-correlated media, demonstrating that the Fractional Gaussian Field (FGF) is able to reproduce a wide range of appearances stemming from short-range to long-range correlations and support macroscopic heterogeneity (left). A reference generated by the classical transport theory is provided for a comparison (right). [20].

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 are able to be reproduced. Compared to 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 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, on the supporting scattering events, biased or not, storage cost and rendering efficiency.

Scattering event. Many-light based methods can not handle single scattering, while the others can handle both single and multiple scattering.

Unbiasedness. Monte Carlo-based methods produce unbiased results, while the others are mostly biased, except for Bitterli et al. [6], which is unbiased for multiple scattering when using zero-order estimator, and Deng et al. [11] which is also unbiased when using delta kernel.

Storage cost. Density based approaches and point based approaches have more memory cost than others, where Bitterli et al. [6] and Deng et al. [11] introduce higher-dimension elements, e.g. photon planes and volumes, thus less memory cost is required compared to UPBP [42]. Many-lights based method has less memory cost compared to the 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 less memory cost. When Monte Carlo based methods are accelerated with path guiding, more memory cost required to store the learned lighting distribution.

Rendering efficiency. Among all these methods, point based methods are the fastest, since they are able to achieve interactive frame rate. However, their rendering results have the lowest quality among all these methods. Monte Carlo based methods are the slowest, since they requires long time to converge, 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. Methods of Bitterli et al. [6] and Deng et al. [11] converge faster than UPBP, thanks to their higher-dimension elements.

12 Challenges and Future Work

This article summarizes several recently proposed efficient rendering methods for participating media, including point-based rendering methods, precomputed methods, and path guiding methods. These three methods improve the rendering efficiency of participating media from different perspectives.

At present, the 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 uniform participating media, while the participating media in the real world are 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 the problem of low rendering efficiency, so how to efficiently render spatially-correlated participating media has important practical significance.

At present, the noise reduction method based

Tab. 1 Comparison among several approaches, considering supporting scattering events, performance, storage cost, etc. Unbiased (multiple) means only unbiased for multiple scattering and biased for singles scattering.

Method	Type	Scattering Event	biased / unbiased	Storage Cost	Rendering Efficiency
UPBP [42]	Density-based	Single + Multiple	biased	●●●●●	●●○○○
Bitterli et al. [6]	Density-based	Single + Multiple	unbiased (multiple)	●●●●○	●●●○○
Deng et al. [11]	Density-based	Single + Multiple	unbiased (delta kernel)	●●●○○	●●●●○
Wang et al. [71]	Point-based	Single + Multiple	biased	●●●●○	●●●●●
VRL [57]	Many light-based	Multiple only	biased	●●●○○	●●●○○
MEMLT [29]	Monte Carlo-based	Single + Multiple	Unbiased	●○○○○	●○○○○
Wang et al. [3]	Monte Carlo-based	Multiple only	biased	●●○○○	●●●○○
Herholz et al. [25]	Monte Carlo-based	Single + Multiple	Unbiased	●●●○○	●●○○○
Deng et al. [10]	Monte Carlo-based	Single + Multiple	Unbiased	●●●○○	●●○○○

on deep learning has been tried a lot in surface rendering, and it is also effective. However, the performance of these methods in participating media is not satisfactory. One important reason is that the characteristics of the participating media are very different from the characteristics of the surface material. So how to design and make full use of the characteristics of the participating media to support the noise reduction requires future research. In addition, some participating media have a strong global effect due to scattering, such as the effect of fog or smoke, and the current network architecture used for surface rendering will cause aliases on the participating media. Therefore, the noise reduction of participating media requires a different network architecture than before.

References

- [1] G. Adrien, H. Binh-Son, V. Nicolas, N. Derek, and H. Toshiya. Gradient-domain volumetric photon density estimation. *Acm Transactions on Graphics*, 37(4):1–13, 2018.
- [2] A. Arbree, B. Walter, and K. Bala. Single-pass scalable subsurface rendering with lightcuts. *Computer Graphics Forum (Proc. Eurographics 2008)*, 27(2):507–516, 2008.
- [3] W. Beibei, G. Liangsheng, and H. Nicolas. Precomputed multiple scattering for light simulation in participating medium. *IEEE Transactions on Visualization and Computer Graphics*, 26(7), 2020.
- [4] L. Belcour, K. Bala, and C. Soler. A local frequency analysis of light scattering and absorption. *ACM Trans. Graph.*, 33(5), Sept. 2014.
- [5] L. Belcour, C. Soler, K. Subr, N. Holzschuch, and F. Durand. 5d covariance tracing for efficient defocus and motion blur. *ACM Trans. Graph.*, 32(3), July 2013.
- [6] B. Bitterli and W. Jarosz. Beyond points and beams: Higher-dimensional photon samples for volumetric light transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 36(4), July 2017.
- [7] B. Bitterli, S. Ravichandran, T. Müller, M. Wrenninge, J. Novák, S. Marschner, and W. Jarosz. A radiative transfer framework for non-exponential media. *Acm Transactions on Graphics*, 37(6):1–17, 2018.
- [8] S. Chandrasekhar. *Radiative transfer*. Dover publications, New York, 1960.
- [9] P. Christensen. Point-based approximate color bleeding. Technical Report 08-01, Pixar Technical Notes, 2008.
- [10] H. Deng, B. Wang, R. Wang, and N. Holzschuch. A practical path guiding method for participating media. *Computational Visual Media*, 6, 02 2020.
- [11] X. Deng, S. Jiao, B. Bitterli, and W. Jarosz. Photon surfaces for robust, unbiased volumetric density estimation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 38(4), 2019.
- [12] F. Durand, N. Holzschuch, C. Soler, E. Chan, and F. X. Sillion. A frequency analysis of light transport. 24(3):1115–1126, July 2005.
- [13] J. Fong, M. Wrenninge, C. Kulla, and R. Habel. Production volume rendering: Siggraph 2017 course. In *ACM SIGGRAPH 2017 Courses*, SIGGRAPH '17, 2017.
- [14] R. Frederickx, P. Bartels, and P. Dutré. Adaptive LightSlice for Virtual Ray Lights. In B. Bickel and T. Ritschel, editors, *EG 2015 - Short Papers*, 2015.
- [15] L. Ge, B. Wang, L. Wang, and N. Holzschuch. A compact representation for multiple scattering in participating media using neural networks. In *ACM SIGGRAPH 2018 Talks*, 2018.
- [16] L. Ge, B. Wang, L. Wang, X. Meng, and N. Holzschuch. Interactive simulation of scattering effects in participating media using a neural network model. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–1, 2019.
- [17] I. Georgiev, J. Krivanek, T. Hachisuka, D. Nowrouzezahrai, and W. Jarosz. Joint importance sampling of low-order volumetric sampling. *ACM Transactions on Graphics*, 32(6):164:1–164:14, 2013.
- [18] I. Georgiev, J. Krivanek, T. Hachisuka, D. Nowrouzezahrai, and W. Jarosz. Joint importance sampling of low-order volumetric scattering. *ACM*

- Transactions on Graphics (Proceedings of SIGGRAPH Asia 2013)*, 32(6), nov 2013.
- [19] J. Guo, P. Bauszat, J. Bikker, and E. Eisemann. Primary sample space path guiding. In W. Jakob and T. Hachisuka, editors, *Eurographics Symposium on Rendering - EI & I*, pages 73–82, July 2018.
- [20] J. Guo, Y. Chen, B. Hu, L.-Q. Yan, Y. Guo, and Y. Liu. Fractional gaussian fields for modeling and rendering of spatially-correlated media. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2019)*, 38(4), 2019.
- [21] J. Guo, M. Li, Q. Li, Y. Qiang, B. Hu, Y. Guo, and L.-Q. Yan. Gradnet: Unsupervised deep screened poisson reconstruction for gradient-domain rendering. *ACM Trans. Graph.*, 38(6), Nov. 2019.
- [22] J. Hanika, M. Droske, and L. Fascione. Manifold next event estimation. *Computer Graphics Forum*, 34(2):87–97, 2015.
- [23] M. Hašan, J. Krivánek, B. Walter, and K. Bala. Virtual spherical lights for many-light rendering of glossy scenes. *ACM Trans. Graph. (proc. Siggraph Asia)*, 28(5):143:1–143:6, Dec. 2009.
- [24] S. Herholz, O. Elek, J. Vorba, H. Lensch, and J. Krivanek. Product importance sampling for light transport path guiding. In *Eurographics Symposium on Rendering*, 2016.
- [25] S. Herholz, Y. Zhao, O. Elek, D. Nowrouzezahrai, H. P. A. Lensch, and J. Krivanek. Volume Path Guiding Based on Zero-Variance Random Walk Theory. *ACM Transactions on Graphics*, Mar. 2019.
- [26] N. Holzschuch. Accurate computation of single scattering in participating media with refractive boundaries. *Computer Graphics Forum*, 34(6):48–59, 2015.
- [27] B.-S. Hua, A. Gruson, V. Petitjean, M. Zwicker, D. Nowrouzezahrai, E. Eisemann, and T. Hachisuka. A survey on gradient-domain rendering. *Computer Graphics Forum*, 38(2):455–472, 2019.
- [28] Y. Huo, R. Wang, T. Hu, W. Hua, and H. Bao. Adaptive matrix column sampling and completion for rendering participating media. *Acm Transactions on Graphics*, 35(6cd):167, 2016.
- [29] W. Jakob and S. Marschner. Manifold exploration: A markov chain monte carlo technique for rendering scenes with difficult specular transport. *ACM Trans. Graph.*, 31(4):58:1–58:13, July 2012.
- [30] A. Jarabo, C. Aliaga, and D. Gutierrez. A radiative transfer framework for spatially-correlated materials. *ACM Transactions on Graphics*, 37(4CD):83:1–83:13, 2018.
- [31] W. Jarosz, C. Donner, M. Zwicker, and H. W. Jensen. Radiance caching for participating media. *ACM Transactions on Graphics*, 27(1):7:1–7:11, March 2008.
- [32] W. Jarosz, D. Nowrouzezahrai, I. Sadeghi, and H. W. Jensen. A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM Trans. Graph.*, 30(1):5:1–5:19, Jan. 2011.
- [33] W. Jarosz, D. Nowrouzezahrai, R. Thomas, P.-P. Sloan, and M. Zwicker. Progressive photon beams. *ACM Trans. Graph. (proc. SIGGRAPH Asia)*, 30(6), 2011.
- [34] W. Jarosz, M. Zwicker, and H. W. Jensen. The beam radiance estimate for volumetric photon mapping. *Comput. Graph. Forum (proc. Eurographics)*, 27(2):557–566, Apr. 2008.
- [35] H. W. Jensen and P. H. Christensen. Efficient simulation of light transport in scenes with participating media using photon maps. In *SIGGRAPH*, pages 311–320, 1998.
- [36] H. W. Jensen, S. R. Marschner, M. Levoy, and P. Hanrahan. A practical model for subsurface light transport. In *SIGGRAPH*, pages 511–518. ACM, 2001.
- [37] J. T. Kajiya. The rendering equation. *SIGGRAPH Comput. Graph.*, 20(4):143–150, 1986.
- [38] J. T. Kajiya and B. P. V. Herzen. Ray tracing volume densities. In *Computer Graphics (ACM SIGGRAPH '84 Proceedings)*, volume 18, pages 165–174, July 1984.
- [39] S. Kallweit, T. Müller, B. McWilliams, M. Gross, and J. Novák. Deep scattering: Rendering atmospheric clouds with radiance-predicting neural networks. *ACM Trans. Graph.*, 36(6), Nov. 2017.
- [40] A. Keller. Instant radiosity. In *SIGGRAPH*, pages 49–56, 1997.
- [41] D. Koerner, J. Novák, P. Kutz, R. Habel, and W. Jarosz. Subdivision next-event estimation for path-traced subsurface scattering. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations*, EGSR '16, page 91–96, 2016.
- [42] J. Krivánek, I. Georgiev, T. Hachisuka, P. Vévoda, M. Šik, D. Nowrouzezahrai, and W. Jarosz. Unifying points, beams, and paths in volumetric light transport simulation. *ACM Trans. Graph. (proc. SIGGRAPH)*, 33(4):1–13, Aug. 2014.
- [43] J. Krivánek and E. d'Eon. A zero-variance-based sampling scheme for monte carlo subsurface scattering. In *ACM SIGGRAPH 2014 Talks*, SIGGRAPH '14, 2014.
- [44] E. P. Lafortune and Y. D. Willems. Rendering participating media with bidirectional path tracing. In *Rendering Techniques '96 (Proceedings of the Seventh Eurographics Workshop on Rendering)*, pages 91–100. Springer-Verlag/Wien, 1996.
- [45] E. W. Larsen and R. Vasques. A generalized linear boltzmann equation for non-classical particle transport. *Journal of Quantitative Spectroscopy & Radiative Transfer*, 112(4):619–631, 2011.
- [46] Y. Liang, B. Wang, L. Wang, and N. Holzschuch. Fast computation of single scattering in participating media with refractive boundaries using frequency analysis. *IEEE Transactions on Visualization and Computer Graphics*, 26(10):2961–2969, 2020.
- [47] J. Marco, A. Jarabo, W. Jarosz, and D. Gutierrez. Second-order occlusion-aware volumetric radiance caching. *ACM Transactions on Graphics*, 37(2), 2018.
- [48] J. Meng, J. Hanika, and C. Dachsbacher. Improving the

- dwivedi sampling scheme. *Computer Graphics Forum*, 35(4):37–44, 2016.
- [49] J. Meng, M. Papas, R. Habel, C. Dachsbacher, S. Marschner, M. Gross, and W. Jarosz. Multi-scale modeling and rendering of granular materials. *ACM Trans. Graph.*, 34(4):49:1–49:13, July 2015.
- [50] T. Müller, M. Papas, M. Gross, W. Jarosz, and J. Novák. Efficient rendering of heterogeneous polydisperse granular media. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 35(6):168:1–168:14, December 2016.
- [51] J. T. Moon, B. Walter, and S. R. Marschner. Rendering Discrete Random Media Using Precomputed Scattering Solutions. In J. Kautz and S. Pattanaik, editors, *Rendering Techniques*. The Eurographics Association, 2007.
- [52] T. Müller, M. Gross, and J. Novák. Practical path guiding for efficient light-transport simulation. *Computer Graphics Forum (Proceedings of EGSR)*, 36(4):91–100, June 2017.
- [53] T. Müller, B. McWilliams, F. Rousselle, M. Gross, and J. Novák. Neural importance sampling. *arXiv preprint arXiv:1808.03856*, 2018.
- [54] T. Müller, B. McWilliams, F. Rousselle, M. Gross, and J. Novák. Neural importance sampling. *ACM Trans. Graph.*, 38(5), Oct. 2019.
- [55] J. Novak, T. Engelhardt, and C. Dachsbacher. Screen-space bias compensation for interactive high-quality global illumination with virtual point lights. In *Proc. Symposium on Interactive 3D Graphics and Games (I3D '11)*, pages 119–124, February 2011.
- [56] J. Novák, D. Nowrouzezahrai, C. Dachsbacher, and W. Jarosz. Progressive virtual beam lights. *Comput. Graph. Forum (Proc. EGSR)*, 31(4), 2012.
- [57] J. Novák, D. Nowrouzezahrai, C. Dachsbacher, and W. Jarosz. Virtual ray lights for rendering scenes with participating media. *ACM Trans. Graph. (proc. SIGGRAPH)*, 31(4), July 2012.
- [58] J. Novák, I. Georgiev, J. Hanika, and W. Jarosz. Monte Carlo methods for volumetric light transport simulation. *Computer Graphics Forum (Proceedings of Eurographics - State of the Art Reports)*, 37(2), May 2018.
- [59] J. Ou and F. Pellacini. Lightslice: Matrix slice sampling for the many-lights problem. *ACM Transactions on Graphics*, 30(6):179:1–179:8, December 2011.
- [60] M. Pauly, T. Kollig, and A. Keller. Metropolis light transport for participating media. In B. Peroche and H. Rushmeier, editors, *Rendering Techniques 2000 (Proceedings of the Eleventh Eurographics Workshop on Rendering)*, pages 11–22. Springer Wien, 2000.
- [61] H. Qin, X. Sun, Q. Hou, B. Guo, and K. Zhou. Unbiased photon gathering for light transport simulation. *ACM Transactions on Graphics*, 34(6), 2015.
- [62] F. Simon, A. Jung, J. Hanika, and C. Dachsbacher. Selective guided sampling with complete light transport paths. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 37(6), 2018.
- [63] B. Sun, R. Ramamoorthi, S. G. Narasimhan, and S. K. Nayar. A practical analytic single scattering model for real time rendering. *ACM Trans. Graph.*, 24(3):1040–1049, July 2005.
- [64] D. Vicini, V. Koltun, and W. Jakob. A learned shape-adaptive subsurface scattering model. *ACM Trans. Graph.*, 38(4), 2019.
- [65] J. Vorba, O. Karlík, M. Šik, T. Ritschel, and J. Krivánek. On-line learning of parametric mixture models for light transport simulation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2014)*, 33(4), aug 2014.
- [66] B. Walter, A. Arbree, K. Bala, and D. P. Greenberg. Multidimensional lightcuts. *ACM Trans. Graph. (proc. Siggraph)*, 25(3):1081–1088, 2006.
- [67] B. Walter, S. Fernandez, A. Arbree, K. Bala, M. Donikian, and D. P. Greenberg. Lightcuts: A scalable approach to illumination. *ACM Trans. Graph. (proc. Siggraph)*, 24(3):1098–1107, 2005.
- [68] B. Walter, S. Zhao, N. Holzschuch, and K. Bala. Single scattering in refractive media with triangle mesh boundaries. *ACM Trans. Graph.*, 28(3), 2009.
- [69] B. Wang, J.-D. Gascuel, and N. Holzschuch. Point-Based Light Transport for Participating Media with Refractive Boundaries. In *Eurographics Symposium on Rendering 2016 (EI&I)*, Dublin, Ireland, June 2016.
- [70] B. Wang, M. Hašan, and L.-Q. Yan. Path cuts: Efficient rendering of pure specular light transport. *ACM Trans. Graph.*, 39(6), Nov. 2020.
- [71] B. Wang and N. Holzschuch. Point-based rendering for homogeneous participating media with refractive boundaries. *IEEE Transactions on Visualization and Computer Graphics*, 24(10):2743–2757, 2018.
- [72] P. Weber, J. Hanika, and C. Dachsbacher. Multiple vertex next event estimation for lighting in dense, forward-scattering media. *Computer Graphics Forum*, 36(2):21–30, 2017.
- [73] Z. Xu, Q. Sun, L. Wang, Y. Xu, and B. Wang. Unsupervised image reconstruction for gradient-domain volumetric rendering. *Computer Graphics Forum*, 39(7):193–203, 2020.
- [74] C. Yuksel and C. Yuksel. Lighting grid hierarchy for self-illuminating explosions. *ACM Transactions on Graphics*, 36(4CD):1–10, 2017.
- [75] Q. Zheng and M. Zwicker. Learning to importance sample in primary sample space. *Computer Graphics Forum (Proc. Eurographics)*, 38(2), 2019.




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



Beibei Wang received the PhD degree from Shandong University, in 2014 and visited Telecom ParisTech from 2012 to 2014. She is an associate professor with the Nanjing University of Science and Technology. She worked as a postdoc with 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 doctor 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.