

Rendering Discrete Participating Media with Geometrical Optics Approximation

Jie Guo, Bingyang Hu, Yanjun Chen, Yuanqi Li, Yanwen Guo and Ling-Qi Yan

Abstract—We consider the scattering of light in participating media composed of sparsely and randomly distributed discrete particles. The particle size is expected to range from the scale of the wavelength to the scale several orders of magnitude greater than the wavelength, and the appearance shows distinct graininess as opposed to the smooth appearance of continuous media. One fundamental issue in physically-based synthesizing this appearance is to determine necessary optical properties in every local region. Since these optical properties vary spatially, we resort to geometrical optics approximation (GOA), a highly efficient alternative to rigorous Lorenz-Mie theory, to quantitatively represent the scattering of a single particle. This enables us to quickly compute bulk optical properties according to any particle size distribution. Then, we propose a practical Monte Carlo rendering solution to solve the transfer of energy in discrete participating media. Results show that for the first time our proposed framework can simulate a wide range of discrete participating media with different levels of graininess and converges to continuous media as the particle concentration increases.

Index Terms—Light scattering, Geometrical optics approximation, Discrete participating media, Volume rendering.

1 INTRODUCTION

Rendering participating media is a long-standing problem in computer graphics, with much effort devoted to solving this problem plausibly and efficiently [1], [2]. From the physics point of view, radiative transfer is rather complicated and should be rigorously derived from Maxwell's electromagnetic theory. To make the simulation tractable, some compromises are made in physically-based rendering over the past decades. Two main assumptions are *independent scattering* [3] and *local continuity* (or *statistical homogeneity* [4]). The first assumption of independent scattering means that the particles forming the medium are far apart from each other and are mutually unaffected. Recently, this assumption has been relaxed in computer graphics via incorporating spatial correlations between scatterers into radiative transfer frameworks [5], [6], [7], [8], [9], leading to non-exponential attenuation of light.

The second assumption of local continuity implies that the medium is homogeneous and compact in each differential volume even if macroscopic heterogeneity exists. Under this circumstance, the light is not sensitive to the discrete spatial distribution of the scatterers, but only to their local average properties. Consequently, light scattering phenomena take place at any point of the medium, resulting in locally smooth renderings. However, many participating media are composed of separate particles distributed randomly within a given volume. The scattering will happen only at the particle positions [4], [10]. These facts indicate that the assumption of continuous media only holds when

the particle size is much smaller as compared with the resolution of the sensor (e.g., human eyes) and the quantity is sufficiently large [11]. Otherwise, individual grains can be observed when zoomed in. To break through this constraint, the graininess of the medium should be taken into consideration.

Several recent studies in computer graphics [12], [13], [14] have noticed the graininess in rendering granular materials. They typically rely on explicit geometries and precomputed transport functions to capture the appearance of discernible grains. As these approaches are designed for a certain amount of very large particles under geometric optics, the scattering behaviors of particles forming the media are rather limited. For instance, diffraction, which dominates light scattering from small particles, is often ignored. In this paper, we attempt to put forward a more general approach to model and render discrete participating media with a large amount of scatterers whose particle size distributions (PSDs) range widely. These media are omnipresent in natural and artificial environments, such as flying dusts, blowing snows, powder suspension, and air bubbles in liquid.

Unlike conventional continuous media, these discrete media have spatially-varying optical properties (e.g., the extinction coefficient) that cannot be determined in advance but should be evaluated on-the-fly. They are closely related to the scattering behavior of each individual particle and the fluctuation of PSDs. To derive necessary optical properties for any local region of a given discrete participating medium in a physically-based manner, one can resort to Lorenz-Mie theory [15], [16] which offers substantial realism in rendering participating media [17]. However, numerical evaluation of the Lorenz-Mie coefficients is known to be difficult and time-consuming when the particle size becomes large [18]. To ameliorate this issue, we introduce geometrical optics approximation (GOA) [18], [19] and use

- Jie Guo, Bingyang Hu, Yanjun Chen, Yuanqi Li and Yanwen Guo are with the State Key Lab for Novel Software Technology, Nanjing University, Nanjing, Jiangsu 210023, P.R. China.
E-mail: guojie@nju.edu.cn, fhymyang@gmail.com, cujooyer@gmail.com, dz1833015@smail.nju.edu.cn and ywguo@nju.edu.cn
- Ling-Qi Yan is with Department of Computer Science, UC Santa Barbara, United States.
E-mail: lingqi@cs.ucsb.edu

it to simplify the computation of light scattering when the particle is sufficiently large. We show how to perform a smooth transition between Lorenz-Mie theory and GOA in computing the optical properties, enabling both high accuracy and low computational cost.

Due to the variations in local PSDs, the optical properties exhibit multi-scale effects with respect to the scene configuration. We derive a novel *multi-scale volumetric rendering equation (VRE)* and propose a practical Monte Carlo rendering solution to solve it. Our solution only relies on the position and the radius of each particle distributed randomly according to some PSDs, avoiding cumbersome geometric modeling and lengthy precomputation. Experimental results verify that the proposed solution is able to capture the distinct grainy appearance of discrete participating media and guarantee temporal coherence in animation. We also show that it converges to continuous media in the limit of particle concentration.

In summary, the main contributions of this paper are:

- a general and physically-based framework for modeling and rendering discrete participating media, considering diffraction, polarization, a wide range of PSDs, etc,
- the use of GOA for efficient and accurate evaluation of the multi-scale bulk optical properties at any local region of a medium, and
- a new Monte Carlo rendering solution that captures both low-frequency haziness and high-frequency graininess in discrete participating media.

2 RELATED WORK

2.1 Participating Media Rendering

Rendering participating media is a challenging but important problem, which requires efficiently solving the VRE [1], [3] by means of Monte Carlo path integration [2], [20], [21], [22], [23], photon density estimation [24], [25], [26], [27], [28], [29], [30], or a combination of both [31]. In our current framework, we choose Monte Carlo path integration by virtue of its elegant simplicity, generality, and accuracy. This technique operates by stochastically constructing a large number of light paths between sensors and emitters to simulate the light transport in the scene. To facilitate the query of particles along paths, we augment each ray with a cylinder, in a way similar to the photon beam [25], [26], [31]. Multiple importance sampling [32], [33] algorithms are beneficial for reducing the large variance caused by Monte Carlo sampling.

2.2 Detailed Volumetric Modeling

Since the original VRE is only a rough approximation to the real radiative transport in participating media, the range of appearance that can be faithfully simulated is limited. Recent trend in computer graphics tries to capture more details in the volume by relaxing the assumptions in the original VRE, enriching the range of achievable appearances. For example, to account for angular anisotropy, the VRE is extended with local directional dependency based on the microflake model [34], [35], [36]. It is also possible to extend the VRE to simulate the effects of spatial correlations [5], [6],

[9], yielding non-exponential attenuation of light. Notably, these methods still assume the media to be statistically homogeneous within each differential volume [4], ignoring any sub-pixel details.

To handle participating media with complex 3D structures, volumetric representations of explicit geometries have been widely used. By capturing the geometric and optical properties of a fabric down to the fiber level, micro-appearance models, described using high-resolution volumes, offer state-of-the-art renderings for fabrics and textiles [37], [38], [39], [40], [41]. Unfortunately, these methods are highly data-intensive and plagued by heavy computation. To improve the performance while maintaining good accuracy, some downsampling strategies [42], [43] are developed. Current rendering solutions for granular materials are also based on explicit geometries and pre-captured optical properties of each individual grain [12], [13], [14]. They generally employ shell tracing to make large jumps inside media. Even so, the computational cost is still high. The goal of this work is to develop a general framework for handling participating media with graininess which also allows rapid computation and convenient usage.

2.3 Glittery Surface Simulation

Our work is also closely related to the simulation of glints on surfaces. Yan et al. [44], [45] suggested using explicit high-resolution normal maps to model sub-pixel surface details and successfully simulated spatially-varying glints with a patch-based normal distribution function. Subsequent work adopted a wave optics model to achieve more accurate results with noticeable color effects [46]. There are also other methods focusing specifically on capturing spatially-varying highlights from scratched surfaces, under either geometric optics [47], [48], [49] or wave optics [50]. To ease the burden of computation and storage, Kuznetsov et al. [51] proposed to learn high-frequency angular patterns from existing examples, using a generative adversarial network (GAN). Jakob et al. [52] addressed the problem of glittery surface simulation using a purely procedural approach which requires far less storage and supports on-the-fly point queries. This approach has been extended to incorporate iridescence [53] and allow fast global illumination [54].

2.4 Lorenz-Mie Theory and GOA

Lorenz-Mie theory [15], [16] develops a rigorous solution to the problem of light scattering by spherical particles. It was introduced to the graphics community by Rushmeier [55] to accurately simulate the physics of light transport in participating media. Later, Callet [56] used this theory to model pigmented materials consisting of pigmented particles in a transparent solvent. Atmospheric phenomena, such as halos and rainbows, are especially favored by this theory [57], [58], [59], [60]. Frisvad et al. [17] generalized the original Lorenz-Mie theory and used it to compute the appearance of materials with different mixed particle concentrations. Though accurate, this theory is computationally expensive. We show that GOA [18], [19], [61], [62], [63], [64], [65] is much more efficient than Lorenz-Mie theory in computing optical properties of individual particles in various media, especially when the particle size is large. Moreover, the

computation of GOA can be made in non-ideal situations such as absorbing particles [65] and non-spherical particles [61], [66], [67], [68]. Compared to Lorenz-Mie theory, GOA is less explored in computer graphics. We choose GOA in our framework, taking advantage of its high performance.

3 LIGHT SCATTERING BY A SINGLE PARTICLE

We first study light scattering by a single particle. We suppose that the particle is approximately spherical and has a set of physical properties including its radius r and the refractive index η_p . Currently, we assume that particles forming the medium have the same composition and only their sizes vary. In this case, the refractive index η_p is fixed. Supposing that the host medium has the refractive index η_m , we can define the relative refractive index of the particle as $\eta = \eta_p/\eta_m$. The size of a spherical particle may also be expressed in terms of the dimensionless size parameter $\alpha = kr = 2\pi\eta_m r/\lambda$, where k is the wave number defined by $k = 2\pi\eta_m/\lambda$ and λ is the wavelength of light in the medium.

To describe the scattering, we need two scattering amplitude functions: $S_1(\theta, \varphi, r)$ and $S_2(\theta, \varphi, r)$, where θ is the scattering angle and φ is the azimuth angle. The subscripts 1 and 2 denote perpendicular and parallel polarizations, respectively. For spherical particles, S_1 and S_2 are invariant with respect to φ , but they change depending on the radius r . For unpolarized light, these two functions define the phase function of a single particle as [4]

$$f_p(\theta, r) = \frac{|S_1(\theta, r)|^2 + |S_2(\theta, r)|^2}{2|k|^2 C_s(r)} \quad (1)$$

which is properly normalized by the scattering cross section C_s :

$$C_s(r) = \int_0^{2\pi} \int_0^\pi \frac{|S_1(\theta, r)|^2 + |S_2(\theta, r)|^2}{2|k|^2} \sin\theta d\theta d\phi. \quad (2)$$

Another important property of the particle is the extinction cross section C_t which is evaluated by

$$C_t(r) = 4\pi \text{Re} \left\{ \frac{S(0, r)}{|k|^2} \right\} \quad (3)$$

with $S(0, r) = S_1(0, r) = S_2(0, r)$. The notation Re takes the real part of a complex number. For particles with absorption, the absorption cross section is given by $C_a(r) = C_t(r) - C_s(r)$.

As seen, once the scattering amplitude functions $S_1(\theta, r)$ and $S_2(\theta, r)$ are available, we can easily find the scattering, extinction and absorption cross sections as well as the phase function of the particle. For light scattering of an electromagnetic wave from a homogeneous spherical particle, exact solutions of the two scattering amplitude functions are given by Lorenz-Mie theory [15], [16]. Its accuracy has been validated against real measurements in various literature [69], [70]. Please refer to Appendix A for more details.

As a rigorous and general electromagnetic treatment of light scattering by spherical particles, Lorenz-Mie theory can precisely handle a wide range of particle sizes. However, as the particle size increases, numerical calculations of the Lorenz-Mie coefficients become very tedious and time-consuming, due to the fact that the number of terms to be

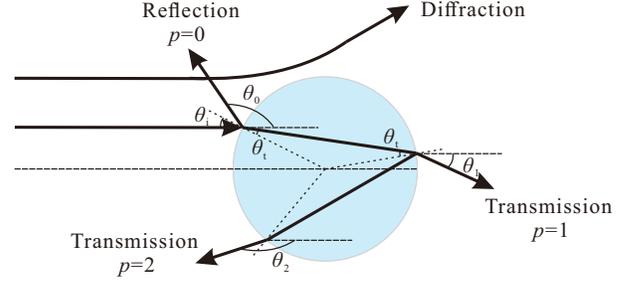


Fig. 1. Light scattering by a spherical particle in the GOA picture.

computed in the series for $S_1(\theta, r)$ and $S_2(\theta, r)$ is proportional to the size parameter α [17], [18]. For this reason, simpler approximate expressions should be developed to reduce the computational complexity. In the case that the particle size is large with respect to the wavelength of the illuminating light, geometrical optics approximation (GOA) [18], [19], [61], [62], [63], [64], [65] provides a simplified but also good solution.

3.1 Geometrical Optics Approximation

Within the framework of GOA, light scattering is calculated by a superposition of classical diffraction, geometrical reflection and transmission. The diffraction is independent of the particle's composition (i.e., the refractive index). Its amplitude functions for the forward direction are readily described by the Fraunhofer diffraction as [18]

$$S_{D,1}(\theta, r) = S_{D,2}(\theta, r) = \alpha^2 \frac{J_1(\alpha \sin \theta)}{\alpha \sin \theta} \quad (4)$$

where J_1 is the first-order Bessel function.

Leaving out diffraction, a light ray hitting a spherical particle at an incident angle θ_i is partially reflected and partially refracted depending on the properties of the interface, as sketched in Fig. 1. The refracted ray may undergo a number of internal reflections before leaving the particle. For each emerging ray, we use an integer p to denote the number of chords it makes inside the spherical particle. Obviously, the externally reflected ray has $p = 0$ while the other rays are transmitted with $p - 1$ internal reflections. The angle of deflection θ_p between the p th emerging ray and the direction of the incident ray is given by

$$\theta_p = 2p\theta_t - 2\theta_i - (p - 1)\pi \quad (5)$$

with $\sin \theta_i = \eta \sin \theta_t$ according to Snell's law. The scattering angle θ is further determined by the deflection angle θ_p as

$$\theta = q(\theta_p - 2\pi l) \quad (6)$$

where $q \in \{1, -1\}$ and l is an integer ensuring that the scattering angle θ is well defined in the range between 0 and π .

1. It is suggested that an appropriate number of terms to sum is $\lceil |\alpha| + 4.3|\alpha|^{1/3} + 1 \rceil$ [71].

2. $q = 1$ indicates that the incident ray hits the particle on the upper hemisphere and $q = -1$ for the lower.

Clearly, reflected and transmitted rays depend on the shape and composition of the particle. Their scattering amplitudes for each polarization are derived as [10], [19]

$$S_j^{(p)}(\theta, r) = \alpha \epsilon_j(\theta_i) \sqrt{\frac{\sin 2\theta_i}{2 \sin \theta |d\theta_p/d\theta_i|}} e^{i\phi} \quad j = 1, 2. \quad (7)$$

Here, the fraction $\epsilon_j(\theta_i)$, which is due to the reflection and/or refraction for an emergent ray of order p , is defined as

$$\epsilon_j(\theta_i) = \begin{cases} R_j(\theta_i) & p = 0 \\ (1 - R_j(\theta_i)^2)(-R_j(\theta_i))^{p-1} & p > 0 \end{cases} \quad (8)$$

with $R_j(\theta_i)$ being the Fresnel reflection coefficients. The phase difference $\phi = \phi_p + \phi_f$ includes ϕ_p due to the length of optical path:

$$\phi_p = 2\alpha(\cos \theta_i - p\eta \cos \theta_t) \quad (9)$$

and ϕ_f due to focal line:

$$\phi_f = \frac{\pi}{2} \left(1 + p - 2l - \frac{s}{2} - \frac{q}{2} \right) \quad (10)$$

with $s = \text{sgn}(d\theta_p/d\theta_i) = \text{sgn}(2p \tan \theta_t / \tan \theta_i - 2)$.

Putting together $S_{D,j}(\theta, r)$ and $S_j^{(p)}(\theta, r)$, we are able to get the total amplitude functions of GOA as

$$S_j(\theta, r) = \begin{cases} \sum_{p=0}^{\infty} S_j^{(p)}(\theta, r) + S_{D,j}(\theta, r) & \theta \in [0, \frac{\pi}{2}) \\ \sum_{p=0}^{\infty} S_j^{(p)}(\theta, r) & \theta \in [\frac{\pi}{2}, \pi] \end{cases} \quad (11)$$

with $j = 1, 2$. These expressions can be evaluated quite efficiently.

In GOA, analytical expression of $C_t(r)$ can be derived as (see the derivation in Appendix B)

$$C_t(r) = 2\pi r^2 + \frac{2\pi r}{|k|} \sum_{p \in \mathcal{P}} \frac{\epsilon_j(0)}{|p/\eta - 1|} \cos(\phi_p + \phi_f) \quad (12)$$

where $\mathcal{P} = \{1, 3, 5, \dots\}$. This expression is fast to evaluate and well captures the ripple structures [72].

For absorbing particles, we certainly have $C_a(r) = C_t(r) - C_s(r) > 0$. Within GOA, the absorption cross section $C_a(r)$ is faithfully approximated by [4]

$$C_a(r) = \frac{16\pi^2 r^3 \eta_i}{3\lambda \eta_r} \left[\eta_r^3 - (\eta_r^2 - 1)^{\frac{3}{2}} \right] \quad (13)$$

in which η_r and η_i are the real and imaginary parts of η , respectively. The scattering amplitude functions $S_1(\theta, r)$ and $S_2(\theta, r)$ are also slightly different. The details are provided in Appendix C.

3.2 Discussions on p

There is also an infinite summation in computing S_1 and S_2 of GOA. However, unlike that in Lorenz-Mie theory, the number of terms needed is independent of the particle size, and a small p suffices in most cases. As shown in Fig. 2, when calculating $(|S_1| + |S_2|)/2$ with $p = 3$ in GOA, we get an almost identical curve with that of $p = 100$, irrespective of the particle size. This is because higher-order reflections ($p > 3$) carry much less energy and have very little impact on the scattered light intensities. Regarding this, p is safely set to 3 in what follows. Please see more discussions in Appendix E.

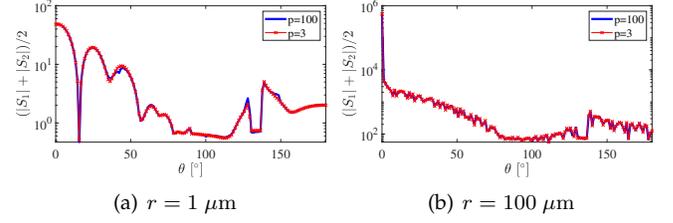


Fig. 2. Impact of p over $(|S_1| + |S_2|)/2$ in GOA. Here, $\eta = 1.33$ and $\lambda = 0.6 \mu\text{m}$.

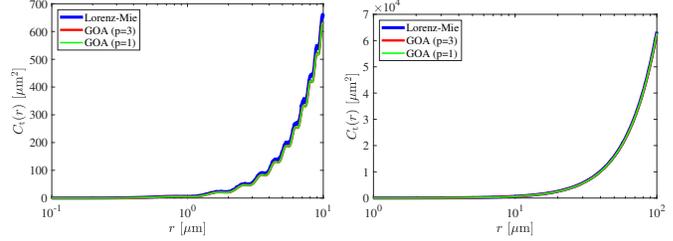


Fig. 3. Variation of the extinction cross section C_t as a function of r . Here, we compare our GOA (red: $p = 3$, green: $p = 1$) calculation against that of Lorenz-Mie theory (blue) with $\eta = 1.33$ and $\lambda = 0.6 \mu\text{m}$. The ripple structures [72] that are well captured by GOA are clearly shown in the left diagram.

Moreover, we can further simplify the extinction cross section to

$$C_t(r) = 2\pi r^2 + \frac{4r\lambda\eta^2}{(\eta + 1)^2|\eta - 1|} \sin\left(\frac{4\pi r}{\lambda}(1 - \eta)\right) \quad (14)$$

by setting $p = 1$, since C_t only relies on the value of S_1 (or S_2) evaluated at $\theta = 0$, and the light rays with $p > 1$ contribute little to the forward scattering. This is evidenced in Fig. 3 where the C_t curves of $p = 1$ (green) and $p = 3$ (red) are virtually indistinguishable for a very wide range of r . In Appendix E, we show that the Relative Mean Squared Error (RelMSE)³ is less than 10^{-5} for $\eta = 1.33$ and $\lambda = 0.6 \mu\text{m}$ using $p = 1$ and $p = 3$, respectively.

3.3 Comparisons between GOA and Lorenz-Mie Theory

To investigate the range of validity of GOA in simulating the scattering patterns of spherical particles, we compare the results with the rigorous Lorenz-Mie results on a wide range of particle radii in Fig. 4. The wavelength of the incident rays is set to $0.6 \mu\text{m}$ and the relative index of refraction is $\eta = 1.33$ in all the calculations. Fig. 4 reveals that the scattering amplitude distributions by GOA align well with those obtained with Lorenz-Mie theory for large particles with $r > 1 \mu\text{m}$. The agreement of these two methods is especially good in almost all directions when the radius is large enough (e.g., $r = 100 \mu\text{m}$). However, when $r \approx 1 \mu\text{m}$, some discrepancies between the two methods appear. These discrepancies become large as the radius of particle decreases further.

To show the influence of these discrepancies on the perception of the translucent appearance, we render a smooth medium comprising monodisperse particles of radius r . We determine the phase function using either Lorenz-Mie

3. The RelMSE for C_t is calculated by $(C_t^{p=3} - C_t^{p=1})^2 / (C_t^{p=3})^2$

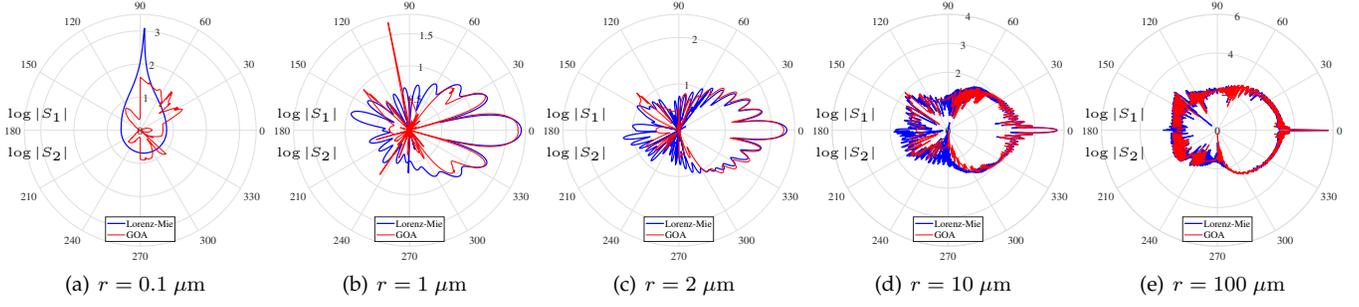


Fig. 4. Visual comparisons of $\log |S_1|$ (upper half) and $\log |S_2|$ (lower half) by Lorenz-Mie calculations (blue curves) with those by GOA (red curves) for $\eta = 1.33$ and $\lambda = 0.6 \mu\text{m}$. The particle radius is set to $r = 0.1, 1, 2, 10$ and $100 \mu\text{m}$, respectively.

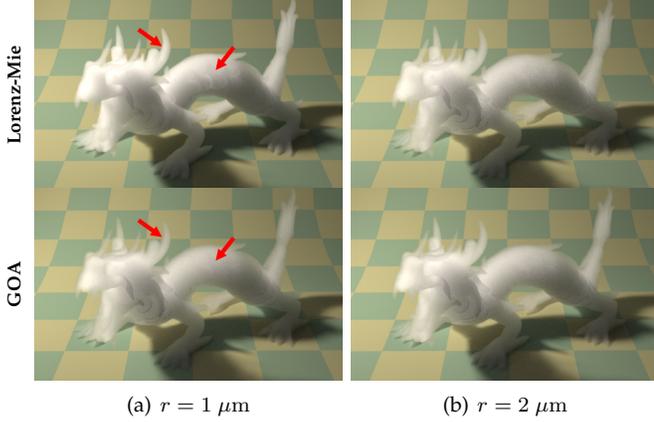


Fig. 5. Rendering a smooth medium with phase functions derived from Lorenz-Mie theory (top row) and GOA (bottom row), respectively. Here, $\eta = 1.33$. The red arrows highlight the differences.

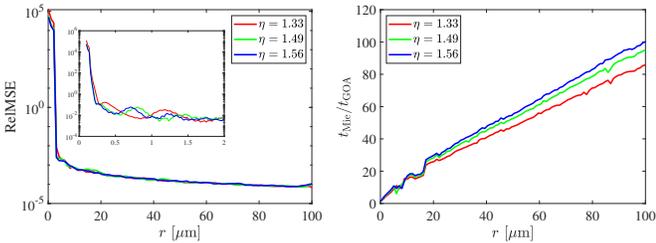


Fig. 6. Comparisons between GOA and Lorenz-Mie theory in RelMSE and runtime complexity. Left: Variation of RelMSE as a function of r . Right: Variation of runtime ratio $t_{\text{Mie}}/t_{\text{GOA}}$ as a function of r .

theory or GOA according to r^4 . The extinction coefficient is set to a constant 0.6 for a fair comparison. The synthesized images are presented in Fig. 5 with r setting to $1 \mu\text{m}$ and $2 \mu\text{m}$, respectively. Clearly, the differences of phase functions between Lorenz-Mie and GOA in the case of $r = 1 \mu\text{m}$ result in inconsistency of appearance. However, this inconsistency almost disappears completely when r goes up to $2 \mu\text{m}$, although there are some mismatches on the backward peaks of $S_1(\theta)$ (see Fig. 4(c), same for $S_2(\theta)$). Fig. 4 and Fig. 5 together have verified the accuracy of choosing GOA to compute $S_1(\theta)$ and $S_2(\theta)$ when $r \geq 2 \mu\text{m}$. Please see more comparisons and discussions in Appendix E.

To further show the similarity between Lorenz-Mie theory and GOA in computing $S_1(\theta)$ and $S_2(\theta)$, we report their

4. These phase functions are precomputed and stored in tables.

Relative Mean Squared Error (RelMSE) in Fig. 6 left. The RelMSE is computed on $(|S_1(\theta)| + |S_2(\theta)|)/2$ and different relative refractive indexes are tested. The results confirm that subtle errors exist when r is large: despite some fluctuations, these calculations of GOA exhibit errors of less than 0.01 as compared with exact Lorenz-Mie calculations when $r \geq 2 \mu\text{m}$.

In the calculations of $(|S_1(\theta)| + |S_2(\theta)|)/2$, Lorenz-Mie theory generally consumes much more time than GOA as evidenced in Fig. 6 right. As r increases, the runtime ratio of Lorenz-Mie theory and GOA $t_{\text{Mie}}/t_{\text{GOA}}$ grows linearly with respect to r , and can easily reach two orders of magnitude difference in performance. This is explained by the fact that the number of terms in Lorenz-Mie theory is linearly proportional to the size parameter α , as we mentioned previously. In comparison, the runtime for GOA is independent of the radius r .

Considering the trade-off between accuracy and time complexity, we choose GOA when $r \geq 2 \mu\text{m}$ and switch to Lorenz-Mie theory otherwise. This makes the runtime of computing $S_1(\theta)$ (or $S_2(\theta)$) almost constant with respect to r while retaining the accuracy as much as possible.

4 BULK OPTICAL PROPERTIES WITH GRAININESS

Now, we consider light scattering by a cloud of spherical particles of the same composition but of different sizes. The particles are assumed to be in each other's far-field regimes and their sizes are likely to range from wavelength-scale to the scale much larger than the wavelength. In this section, we first discuss the particle size distribution that may vary spatially and then study the bulk optical properties of the discrete participating medium considering graininess. Thanks to the high efficiency of GOA, we are able to evaluate the bulk optical properties on-the-fly.

4.1 Particle Size Distribution

We use the particle size distribution (PSD) $N(r)$ to describe the population of particles in a discrete participating medium. Thus $N(r)dr$ is the total concentration (particle number per unit volume) of particles with sizes in the domain $[r, r + dr]$. The total particle number concentration within some limited interval $[r_{\min}, r_{\max}]$ of sizes is obtained by $N_0 = \int_{r_{\min}}^{r_{\max}} N(r)dr$. To use $N(r)$ as a probability density distribution (PDF), we have to normalize $N(r)$ via $N(r)/N_0$.

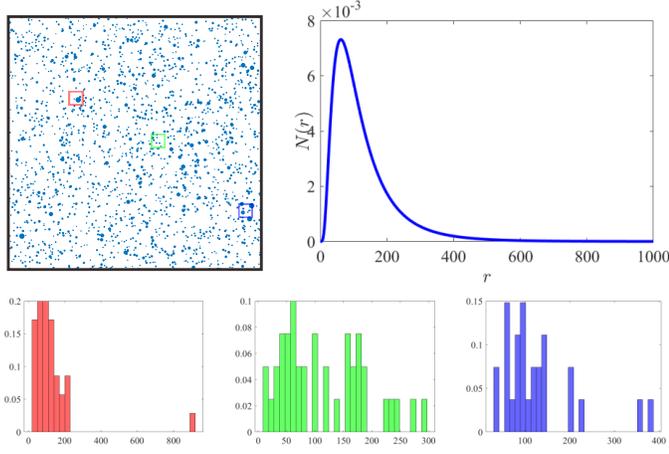


Fig. 7. Illustration of spatially-varying PSD. First row: A realization of 2000 polydisperse particles and its global PSD. Second row: Actual PSDs for particles in three small patches.

It is generally reported that particle sizes follow close to a log-normal distribution [17], [68]:

$$N(r) = \frac{N_0}{\sqrt{2\pi} r \ln \sigma_g} \exp \left[-\frac{1}{2} \left(\frac{\ln r - \ln \bar{r}_g}{\ln \sigma_g} \right)^2 \right] \quad (15)$$

in which σ_g is the geometric standard deviation and \bar{r}_g is the geometric mean radius. Obviously, this statistical tendency stems from the observation of a large number of particles. A small number of particles will give rise to a size distribution deviating from the log-normal distribution. To demonstrate this in 2D, we generate 2000 particles in a box with uniformly distributed positions and log-normal distributed radii. The visualization of these particles and its PSD are shown in the first row of Fig. 7. We then extract three small patches from the box and estimate their actual PSDs by binning. As expected, the PSDs plotted in the second row of Fig. 7 vary spatially and contain quite different features. In what follows, we use the notation $N(r, \mathbf{x})$ to emphasize that the PSD varies spatially. Nevertheless, the ensemble average over these spatially-varying PSDs converges to the log-normal distribution shown in the first row of Fig. 7. Rendering with this global PSD yields a smooth appearance similar to that from a traditional continuous medium.

4.2 Bulk Optical Properties

With the spatially-varying PSD $N(r, \mathbf{x})$, we are able to obtain the bulk optical properties of a local area in which many independent particles are immersed. Supposing that $\mathcal{V}_{\mathbf{x}}$ is a small volume centered around \mathbf{x} , the bulk extinction coefficient σ_t ⁵ of this volume is evaluated by

$$\sigma_t(\mathcal{V}_{\mathbf{x}}) = \frac{1}{\mu(\mathcal{V}_{\mathbf{x}})} \int_{r_{\min}(\mathcal{V}_{\mathbf{x}})}^{r_{\max}(\mathcal{V}_{\mathbf{x}})} C_t(r) \int_{\mathbf{x} \in \mathcal{V}_{\mathbf{x}}} N(r, \mathbf{x}) d\mathbf{x} dr \quad (16)$$

in which $\mu(\mathcal{V}_{\mathbf{x}})$ is the measurement of $\mathcal{V}_{\mathbf{x}}$. $r_{\min}(\mathcal{V}_{\mathbf{x}})$ and $r_{\max}(\mathcal{V}_{\mathbf{x}})$ return the minimum and maximum particle radii inside $\mathcal{V}_{\mathbf{x}}$, respectively. For brevity, we simplify both $r_{\min}(\mathcal{V}_{\mathbf{x}})$ and $r_{\max}(\mathcal{V}_{\mathbf{x}})$ by dropping $\mathcal{V}_{\mathbf{x}}$ henceforth. For the

5. In rendering literature, the symbol σ refers to the cross section sometimes, while using μ for the coefficient.

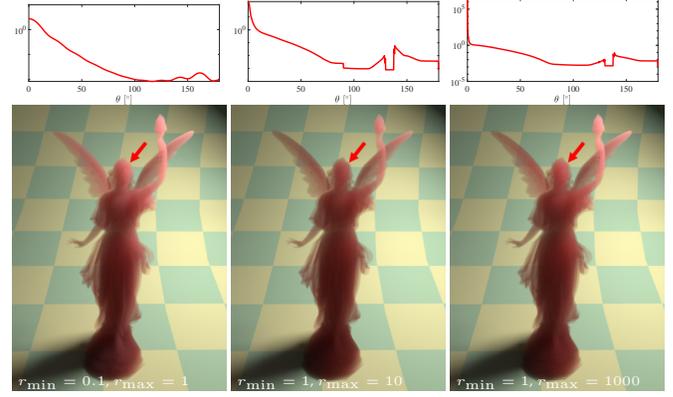


Fig. 8. Top row: Visualizations of the phase functions estimated by our method. Bottom row: The corresponding renderings of a smooth homogeneous medium with these phase functions, while keeping other properties fixed. The red arrows highlight the differences.

scattering coefficient and the absorption coefficient, they can be defined in a similar way by replacing $C_t(r)$ with $C_s(r)$ and $C_a(r)$, respectively⁶. Generally, these properties exhibit multi-scale effects with respect to the size of $\mathcal{V}_{\mathbf{x}}$.

The ensemble phase function is derived as

$$\begin{aligned} f(\mathcal{V}_{\mathbf{x}}, \theta) &= \frac{\int_{r_{\min}}^{r_{\max}} C_s(r) f_p(\theta, r) \int_{\mathbf{x} \in \mathcal{V}_{\mathbf{x}}} N(r, \mathbf{x}) d\mathbf{x} dr}{\mu(\mathcal{V}_{\mathbf{x}}) \sigma_s(\mathcal{V}_{\mathbf{x}})} \\ &= \frac{\int_{r_{\min}}^{r_{\max}} (|S_1(\theta, r)|^2 + |S_2(\theta, r)|^2) \int_{\mathbf{x} \in \mathcal{V}_{\mathbf{x}}} N(r, \mathbf{x}) d\mathbf{x} dr}{2|k|^2 \mu(\mathcal{V}_{\mathbf{x}}) \sigma_s(\mathcal{V}_{\mathbf{x}})}. \end{aligned} \quad (17)$$

in which $\sigma_s(\mathcal{V}_{\mathbf{x}})$ serves as the normalization factor for $f(\mathcal{V}_{\mathbf{x}}, \theta)$. Fig. 8 visualizes the phase functions generated by the above formula. Here, we use uniform sampled radii between r_{\min} and r_{\max} ⁷. By fixing other properties, we show the influence of these phase functions on the final appearance of a smooth homogeneous medium in Fig. 8. Clearly, this formula is only valid for $\sigma_s(\mathcal{V}_{\mathbf{x}}) \neq 0$. When $\sigma_s(\mathcal{V}_{\mathbf{x}}) = 0$, i.e., the volume $\mathcal{V}_{\mathbf{x}}$ is free of particles, $f(\mathcal{V}_{\mathbf{x}}, \theta)$ degenerates into a delta function: $\sigma_s(\mathcal{V}_{\mathbf{x}}) f(\mathcal{V}_{\mathbf{x}}, \theta) = \delta(\theta)$.

Similarly, we can derive the transmittance along a light beam of length s as

$$T_A(s) = \exp \left\{ -\frac{1}{\mu(\mathcal{A})} \int_{r_{\min}(\mathcal{C})}^{r_{\max}(\mathcal{C})} C_t(r) \int_{\mathbf{x} \in \mathcal{C}} N(r, \mathbf{x}) d\mathbf{x} dr \right\} \quad (18)$$

in which $\mathcal{C} = \mathcal{A} \times s$ represents a small cylinder around the light beam and \mathcal{A} is its cross section. The derivation is provided in Appendix D.

For particles with a monodisperse distribution, the transmittance is simplified to

$$T_A(s) = \exp \left\{ -\frac{C_t}{\mu(\mathcal{A})} \int_{\mathbf{x} \in \mathcal{C}} N(\mathbf{x}) d\mathbf{x} \right\} \quad (19)$$

in which C_t is a constant and $N(\mathbf{x})$ is a function of \mathbf{x} only. The integral in the above formula simply counts the number of particles located in the query region $\mathcal{A} \times s$.

6. These properties can be viewed as the properties at position \mathbf{x} when $\mathcal{V}_{\mathbf{x}}$ is infinitely small, i.e., $\sigma_t(\mathbf{x}) = \sigma_t(\mathcal{V}_{\mathbf{x}})$, $\sigma_s(\mathbf{x}) = \sigma_s(\mathcal{V}_{\mathbf{x}})$ and $\sigma_a(\mathbf{x}) = \sigma_a(\mathcal{V}_{\mathbf{x}})$.

7. This is similar to the log-normal distribution with a very large σ_g .

5 RENDERING SOLUTION

With these bulk optical properties, we are able to derive a multi-scale *volumetric rendering equation* (VRE) describing radiative transfer in discrete random media. Then, we develop a Monte Carlo sampling based solution to solve the VRE. This solution only requires the position and the size of each particle, avoiding the explicit tessellation of its shape. In a preprocessing stage, we generate and store N_{tot} particles with random positions and log-normal distributed radii for a discrete participating medium. During rendering, the stored particles are queried to determine the optical properties for each traced ray. A uniform grid is developed for acceleration.

5.1 Multi-scale Volumetric Rendering Equation

Conventionally, the VRE describing macroscopic light scattering in participating media is written as

$$L(\mathbf{x}, \boldsymbol{\omega}) = T(\mathbf{x}, \mathbf{z})L(\mathbf{z}, \boldsymbol{\omega}) + \int_0^z T(\mathbf{x}, \mathbf{y}) \int_{\mathbb{S}^2} \sigma_s(\mathbf{y}) f(\mathbf{y}, \boldsymbol{\omega}', \boldsymbol{\omega}) L_i(\mathbf{y}, \boldsymbol{\omega}') d\boldsymbol{\omega}' dy \quad (20)$$

in which $L(\mathbf{x}, \boldsymbol{\omega})$ is the radiance arriving at $\mathbf{x} \in \mathbb{R}^3$ along a direction $\boldsymbol{\omega} \in \mathbb{S}^2$, f represents the scattering phase function characterizing the probability of radiation incident from $\boldsymbol{\omega}'$ being scattered into direction $\boldsymbol{\omega}$, z is the distance through the medium to the nearest boundary at $\mathbf{z} = \mathbf{x} - z\boldsymbol{\omega}$ and $\mathbf{y} = \mathbf{x} - y\boldsymbol{\omega}$ is a point at distance $y \in (0, z)$. The conventional transmittance between \mathbf{x} and \mathbf{y} is computed as $T(\mathbf{x}, \mathbf{y}) = \exp\{-\int_0^y \sigma_t(\mathbf{x} - s\boldsymbol{\omega}) ds\}$.

By substituting the multi-scale properties into the above equation, we arrive at a multi-scale version of the VRE:

$$L(\mathbf{x}, \boldsymbol{\omega}) = T_{\mathcal{A}}(\mathbf{x}, \mathbf{z})L(\mathbf{z}, \boldsymbol{\omega}) + \int_0^z T_{\mathcal{A}}(\mathbf{x}, \mathbf{y}) \int_{\mathbb{S}^2} L_i(\mathbf{y}, \boldsymbol{\omega}') Q(\theta) d\boldsymbol{\omega}' dy \quad (21)$$

with

$$Q(\theta) = \begin{cases} \sigma_s(\mathcal{V}_{\mathbf{y}}) f(\mathcal{V}_{\mathbf{y}}, \theta) & N(\mathcal{V}_{\mathbf{y}}) \neq 0 \\ \delta(\theta) & N(\mathcal{V}_{\mathbf{y}}) = 0. \end{cases} \quad (22)$$

Since this multi-scale VRE is a general extension to the conventional one, it naturally supports multiple scattering.

5.2 Query Cylinder

In our multi-scale VRE, every optical property depends on a PSD while the PSD is defined on a differential volume. To evaluate the transmittance $T_{\mathcal{A}}(\mathbf{x}, \mathbf{y})$ between any two positions \mathbf{x} and \mathbf{y} in the medium, we need a differential volume around the ray $\mathbf{x} \rightarrow \mathbf{y}$. This volume is used to query particles which contribute to the transmittance $T_{\mathcal{A}}(\mathbf{x}, \mathbf{y})$. In our implementation, we design it to be a thin cylinder centered around $\mathbf{x} \rightarrow \mathbf{y}$, as illustrated in Fig. 9. We name such a cylinder as a query cylinder. In this sense, we view each ray as a ‘‘fat ray’’ which gathers small particles along its trajectory. This is quite different from the implementation of rendering continuous media in which the optical properties are determined globally, without explicitly querying particles in a local area.

In theory, the cross section \mathcal{A} should be infinitely small. However, a too small one may have numerical issues and

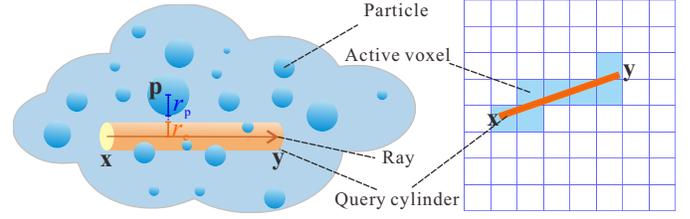


Fig. 9. Illustration of a query cylinder and the active voxels in our ray traversal algorithm. Only the particles contained in the active voxels will be queried.

cause large variance. Conversely, bias will be introduced in $T_{\mathcal{A}}(\mathbf{x}, \mathbf{y})$ when \mathcal{A} is very large. In practice, we select \mathcal{A} as follows and keep it unchanged as the ray traverses the medium. Supposing that z_{near} and z_{med} respectively denote the depth of the near plane and the smallest depth of the medium in the view frustum, the size of \mathcal{A} is selected according to

$$S_{\mathcal{A}} = k S_{\text{pix}} \frac{z_{\text{med}}}{z_{\text{near}}}. \quad (23)$$

Here, S_{pix} is the pixel’s size and k can be viewed as the percentage of the pixel’s footprint at the distance z_{med} . Typically, satisfactory results are obtained when k is in the range $[0.5, 1]$. The influence of k on the visual effect is discussed in the next section.

5.3 Gathering Particles

Gathering particles within a query cylinder (central ray: $\mathbf{x} \rightarrow \mathbf{y}$, radius: r_c and cross section: \mathcal{A}) requires conducting sphere-cylinder intersection test for every particle in the medium. Given a particle with the position \mathbf{p} and radius r_p , it is supposed to be inside the query cylinder if the distance from \mathbf{p} to the central ray $\mathbf{x} \rightarrow \mathbf{y}$ is smaller than $r_p + r_c$, as illustrated in Fig. 9 left.

Testing all particles of the medium is notoriously time-consuming. To boost the performance, we accelerate the process of ray traversing the medium using the 3D digital differential analyzer (3D-DDA) [73], [74]. Specifically, we construct a uniform grid for the medium and adopt a ray traversal algorithm, similar to that in [73], to find the active voxels intersected by the query cylinder. Fig. 9 right illustrates all the active voxels corresponding to the orange query cylinder. Only those particles inside the active voxels will be tested against the query cylinder. We determine the active voxels simply by the ray $\mathbf{x} \rightarrow \mathbf{y}$. This introduces negligible bias, because the radius of the cross section is more than two orders of magnitude smaller than the side length of the voxel. This is significantly different to the thick beams used in beam radiance estimation [25]. After collecting all the particles inside the query cylinder, we accumulate their contributions to the transmittance $T_{\mathcal{A}}(\mathbf{x}, \mathbf{y})$ according to Eq. (18).

To construct a uniform grid, its resolution should be carefully determined. We have observed by experiments that high performance is achieved when roughly one particle resides in each voxel after space subdivision.

5.4 Computing $Q(\theta)$

To solve the multi-scale VRE, we also have to compute $Q(\theta)$ at any sample position \mathbf{y} . $Q(\theta)$ describes the angular distribution of scattering at \mathbf{y} . To quickly compute it, a small query region around \mathbf{y} should be defined. We set this query region to a small sphere with radius r_c . If this query region contains $N_q > 0$ particles, we evaluate $Q(\theta)$ with the following formula:

$$Q(\theta) = \frac{1}{2|k|^2\mu(\mathcal{V}_{\mathbf{y}})} \sum_{i=1}^{N_q} (|S_1(\theta, r_i)|^2 + |S_2(\theta, r_i)|^2). \quad (24)$$

5.5 Importance Sampling

Similar to the simplifications used in rendering surface glints [52], [53], importance sampling is performed according to the global optical properties of the medium, assuming it to be continuous. Specifically, we use the global extinction coefficient for free-flight sampling and use the tabulated global phase function for angular sampling. These global optical properties only need to be determined once in the preprocessing stage, assuming the entire bounding box of the medium to be $\mathcal{V}_{\mathbf{x}}$ in Eq. (16) and Eq. (17).

6 RESULTS

We have implemented the rendering solution on top of the Mitsuba renderer [75], with spectral rendering enabled. We use 8 spectral samples in the range of the visible spectrum at equally-spaced locations [46]. After rendering, we convert the spectral values to the sRGB color space. All synthesized images are created on a PC with an Intel 16-core i7-6900K CPU and 16G RAM.

To compute the bulk optical properties of a discrete participating medium, we need to specify the complex refractive index (η) for the particles involved and the global PSD (\bar{r}_g and σ_g). The physical unit for the particle radius is μm . We also provide an upper bound to the particle size ($r_{\max} = 2000$) to avoid unreasonably large particles which are unusual and are no longer suitable to be treated as participating media. As mentioned previously, we use GOA in the calculation of the scattering amplitude functions when $r \geq 2$ and switch to Lorenz-Mie theory otherwise. Except the STAIRCASE scene, the refractive index is set according to the data of ice selected from [17], and the host medium is set to be air with $\eta_m = 1$.

6.1 Comparisons Against Explicit Path Tracing

We first compare our method with the traditional path tracing. Previous methods simulating the grainy appearance of discrete participating media mostly rely on explicit path tracing (EPT), with potential approximations to simplify the computation of high-order scattering [12], [13], [14]. However, EPT and other approximations are restricted to geometric optics. This means that only surface reflection and refraction are properly handled. In principle, the mesh of every particle should be explicitly generated and costly ray-object intersections are required.

Compared with EPT, our method offers at least two benefits. First, particle scattering is considered which includes the Fraunhofer diffraction and phase differences, etc.

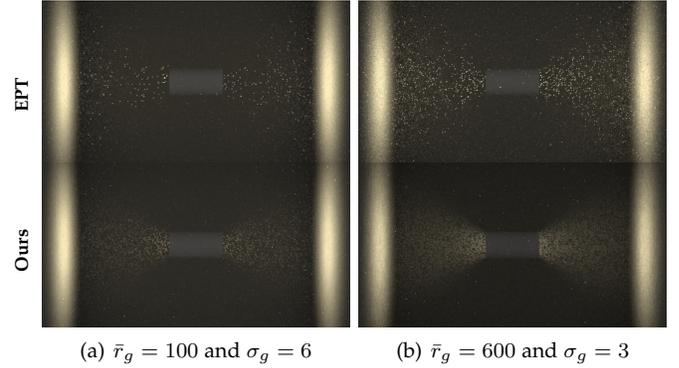


Fig. 10. Comparisons between EPT and our method in simulating the grainy appearance of discrete participating media with different PSDs ($N_{\text{tot}} = 5 \cdot 10^5$).

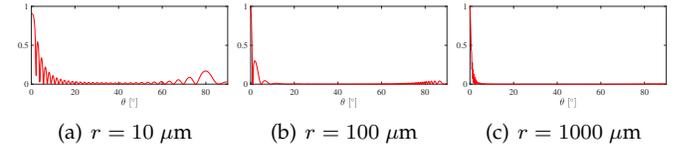


Fig. 11. Angular percentages of energy contributed by the Fraunhofer diffraction for different sized particles.

This significantly expands the range of particles that can be handled. Notably, these optical phenomena are quite important in correctly simulating light scattering, especially for very small particles. To verify this, we produce and render $5 \cdot 10^5$ spherical particles with different PSDs in Fig. 10. These particles have random positions and log-normal sampled radii. For EPT, small transparent balls are instantiated in the scene. Since only surface reflection and refraction are computed for EPT (the top row), the energy inherently belonging to the Fraunhofer diffraction is not correctly captured by EPT, resulting in overly sparse and specular volumetric glints in this LAMP scene. The importance of the Fraunhofer diffraction is visualized in Fig. 11. This figure plots the percentages of energy contributed by the Fraunhofer diffraction at different scattering angles and for different sized particles. Obviously, the Fraunhofer diffraction cannot be ignored especially for small particles. Moreover, EPT easily misses many small particles that are hard to be gathered along an ordinary light path, leading to a slow convergence rate. Also, even substantially increasing the sampling rate or using “fat ray” tracing similar to ours, the appearance is still quite sparse. On the contrary, our method (in the bottom row) preserves this energy from the Fraunhofer diffraction and produces smoother appearance that is closer to real scenarios, thanks to the Airy’s pattern⁸ caused by the Fraunhofer diffraction and the efficient rendering solution tailored for sparse media.

When only very large particles exist in the medium, our method and EPT tend to produce the similar grainy appearance, as shown in Fig. 12. Here, we render a discrete medium with 10^5 particles of the same size ($1000 \mu\text{m}$). Since the radius is sufficiently large, pure geometric optics

⁸ The Fraunhofer diffraction will distort light passing through the particle, making it visually large. The smaller the particle is, the larger the light’s distribution is.

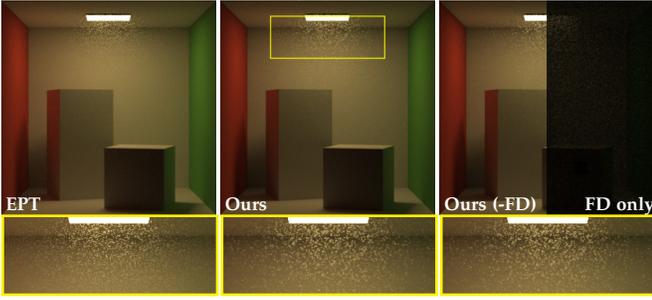


Fig. 12. Comparisons between EPT and our method with or without the Fraunhofer diffraction (-FD) in rendering 10^5 large particles of the same size ($1000 \mu\text{m}$). The impact of the Fraunhofer diffraction is highlighted in the difference image (FD only).

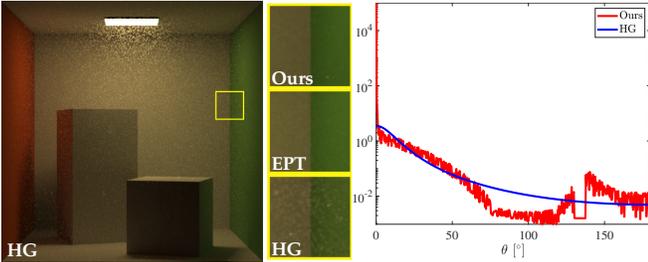


Fig. 13. Comparisons between our directional scattering model and the best-fit Henyey and Greenstein (HG) model [76]. Due to obvious discrepancies between our estimated phase function (red curve in the right diagram) and the HG phase function (blue curve in the right diagram), the renderings are slightly different (see the closeups).

becomes applicable and serves as a valid approximation to the particle scattering. This is further verified in the third column of Fig. 11. We see that the contribution of the Fraunhofer diffraction concentrates in a very narrow angle in this case, making it hard to be observed. Consequently, if we remove the Fraunhofer diffraction in our model and leave only reflection and refraction (the third column of Fig. 12), we will achieve the grainy appearance similar to that of our full model. However, there are still subtle differences (highlighted in the difference image) contributed by the Fraunhofer diffraction.

The second strength of our rendering solution lies in the runtime performance. For rendering these sparse media in the LAMP scene and the CORNELL BOX scene in Figs. 10 and 12, our method achieves a $1.2\times$ speedup over EPT in the rendering time.

6.2 Comparisons Against the Henyey and Greenstein model

In our framework, we derive the phase function from Lorenz-Mie theory and GOA. In computer graphics, it is more common to adopt an empirical model, e.g., the Henyey and Greenstein (HG) model [76], because of its simplicity and well-defined behavior. However, the HG model is not very accurate, as pointed out by various literature [69], [77], [78]. In Fig. 13 we compare our estimated phase function with the best-fit HG phase function in rendering the same scene as in Fig. 12. Since the HG phase function cannot faithfully encode the scattering pattern from these relatively large particles ($r = 1000 \mu\text{m}$), the rendering with the best-fit HG phase function is slightly different from ours, as

compared in the insets. Recall that our rendering is close to that generated using EPT in this specific scene with large particles.

6.3 Comparisons Against Continuous Media

We also compare the grainy appearance of discrete media simulated by our method with the smooth appearance of continuous media. In Fig. 14, we assume that the particles in the DRAGON scene possess the same radius r and are randomly distributed in a cube of the volume V . Under this configuration, it is easy to derive the extinction coefficient and the scattering coefficient of the continuous media as $C_t(r)N_{\text{tot}}/V$ and $C_s(r)N_{\text{tot}}/V$, respectively. The phase function can be computed in a similar way and stored in a table. With these global properties, traditional volumetric path tracing is applicable to render these continuous media. Generally, each discrete medium and its paired continuous medium have the same overall brightness. For the discrete media, when the number of particles N_{tot} is small, we can clearly observe individual particles lit by the lamps. As N_{tot} increases, the appearance tends to become hazy, and the rendering result gets closer to the corresponding continuous medium. When N_{tot} is sufficiently large (e.g., $N_{\text{tot}} = 10^7$), the discrete medium and its continuous counterpart will achieve quite similar appearance. In Fig. 15, a much denser medium ($N_{\text{tot}} = 10^8$) is rendered by our method, which again produces smooth appearance matching that from a continuous medium. These pair-wise comparisons demonstrate that our rendering solution converges to the traditional volumetric rendering of continuous media in the limit of particle concentration.

6.4 Choice of Query Cylinder's Cross Section

We determine the query cylinder's radius r_c according to Eq. (24) in which the parameter k plays an important role. We suggest to choose its value in the range $[0.5, 1]$ which yields reasonable grainy appearance as shown in the second row of Fig. 16. Values in this range allow us to faithfully capture almost all grains in the scenes with little bias. Generally, a too large k will produce uncomfortable aliasing as shown in the last two rows of Fig. 16. In these two cases, although the overall brightness is similar to that of $k = 0.5$, the volumetric glints are overly blurred. On the other hand, a too small k will miss many particles during query and is therefore inefficient, especially when the medium is very sparse, e.g., $N_{\text{tot}} = 10^4$. However, k with a value smaller than 0.5 is also acceptable sometimes. For instance, the first image in the bottom row of Eq. 16 is rendered with $k = 0.2$ which achieves the similar effect with that of $k = 0.5$ for this relatively dense medium ($N_{\text{tot}} = 10^5$).

6.5 Choice of Grid Resolution

We employ a uniform grid to accelerate the ray traversal process, considering that particles are uniformly distributed in the scene. As shown in the right diagram of Fig. 17, the grid resolution (Res) will influence the performance. Although no analytical analysis can be referred to select the best resolution, we empirically observe that a grid resolution yielding roughly one particle per voxel achieves

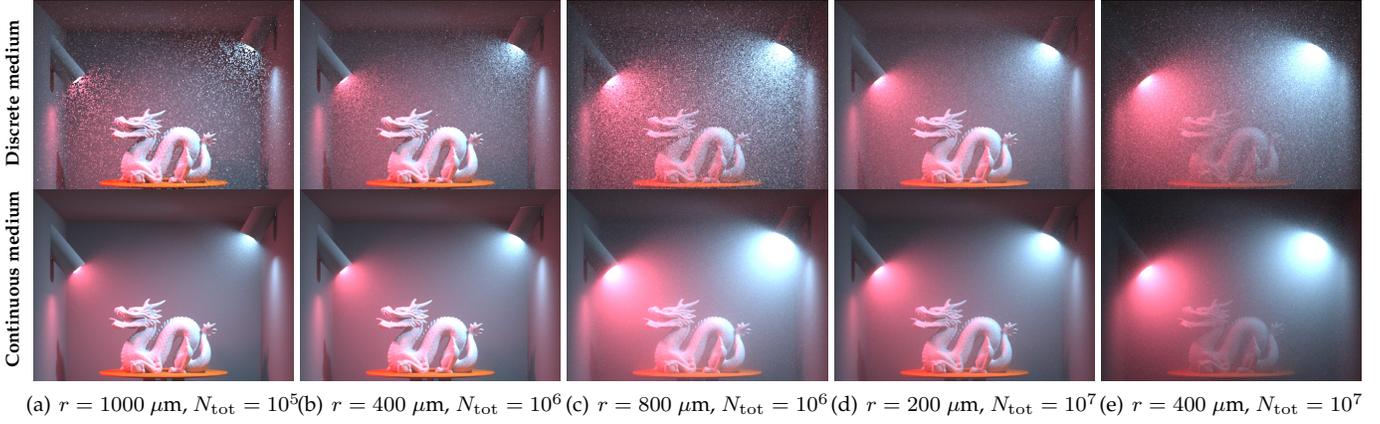


Fig. 14. Visual comparisons against continuous media with varying optical properties. As the particle number N_{tot} increases, the appearance of the discrete medium exhibits obvious low-frequency haziness and converges to that of a continuous medium.

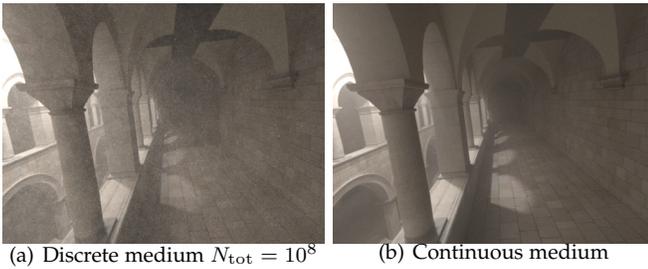


Fig. 15. Comparison between our rendering solution and its continuous counterpart on a very dense medium ($r = 400 \mu\text{m}$ and $N_{\text{tot}} = 10^8$).

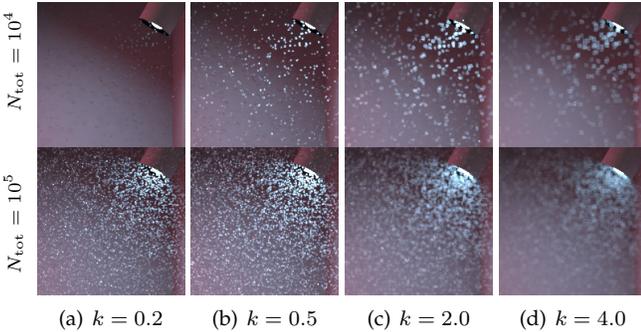


Fig. 16. Impact of the query cylinder's cross section on the simulated grainy appearance. Here, k is used in Eq. (24) which determines the percentage of the pixel's footprint at the nearest boundary of the medium.

the optimum solution for most scenes. The BENCH scene in the left panel of Fig. 17 contains $2 \cdot 10^5$ particles and achieves the best performance at a resolution of $60 \times 60 \times 60$. As the resolution increases, the runtime grows steadily due to the additional cost introduced by the grid. However, a resolution much lower than $60 \times 60 \times 60$ will have a very poor performance since too many particles reside in each voxel. For other scenes, a similar conclusion can be drawn.

6.6 Impact of the Refractive Index

The proposed rendering solution can be easily generalized to support absorbing particles. The absorption cross section of each particle is computed by Eq. (13) which varies linearly

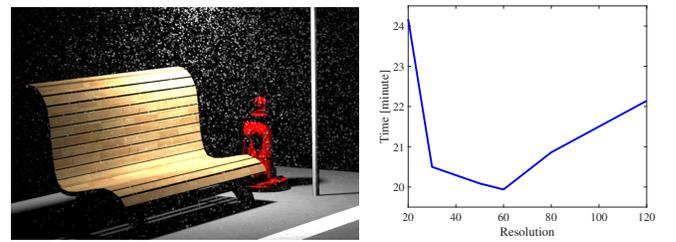


Fig. 17. Evolution of the rendering time (at 2048 spp) over the grid resolution. This BENCH scene contains $N_{\text{tot}} = 2 \cdot 10^5$ particles ($\bar{r}_g = 800$ and $\sigma_g = 4$) and achieves the lowest rendering time at a resolution of $60 \times 60 \times 60$ which is roughly one particle per voxel.

with the imaginary part of the complex refractive index, i.e., η_i . The extinction cross section is slightly modified according to the formulas in Appendix C. The results of changing η_i are shown in Fig. 18. This STAIRCASE scene aims to simulate flying dusts in a dirty room lit by a local area light through the window. As expected, an increase of η_i will increasingly dim the intensity of scattering.

6.7 Impact of the Global PSD

Fig. 19 analyzes the impact of the global PSD on the appearance of discrete participating media. Here, we generate $N_{\text{tot}} = 4 \cdot 10^6$ particles with different global PSDs in the HOUSE scene. This scene is designed to simulate the appearance of blowing snows in the sky. The first row shows the impact of the geometric mean radius \bar{r}_g . The general trend is that the scattering effects become increasingly prominent as \bar{r}_g grows, since the scattering coefficient is positively correlated with the particle radius. Concerning the geometric standard deviation σ_g , it is responsible for the level of graininess as shown in the second row of Fig. 19. A small σ_g tends to generate smoother appearance than a large one. This is to be expected since a large σ_g means local PSDs changing widely, leading to stronger graininess.

In fact, the PSD is not limited to the log-normal distribution. Other distributions also work. For instance, in the third row of Fig. 19, we show a cloud of blowing snows following a bimodal log-normal distribution. We generate $1.99 \cdot 10^7$ small particle with $\bar{r}_{g,1} = 100$ (or $\bar{r}_{g,1} = 200$)

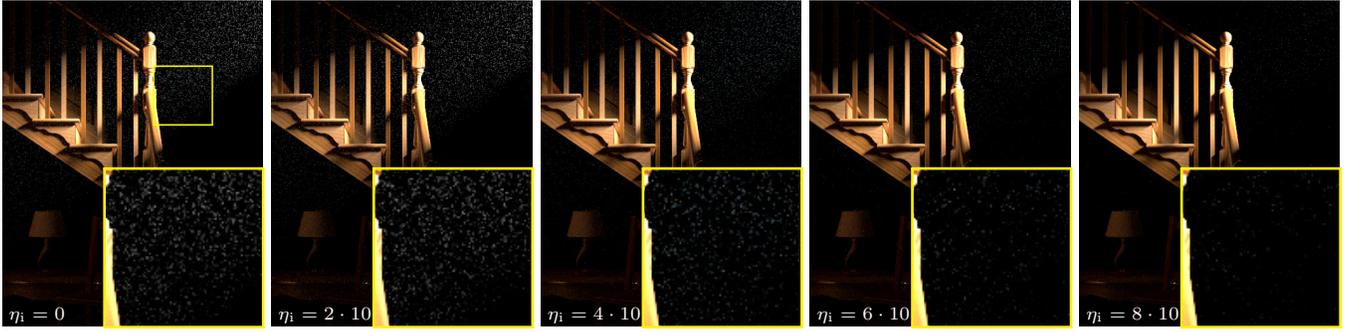


Fig. 18. Impact of the imaginary part of the refractive index (η_i) on the STAIRCASE scene. The real part is set to 1.5. Other scene parameters are $N_{\text{tot}} = 3 \cdot 10^5$, $\bar{r}_g = 800$ and $\sigma_g = 4$.

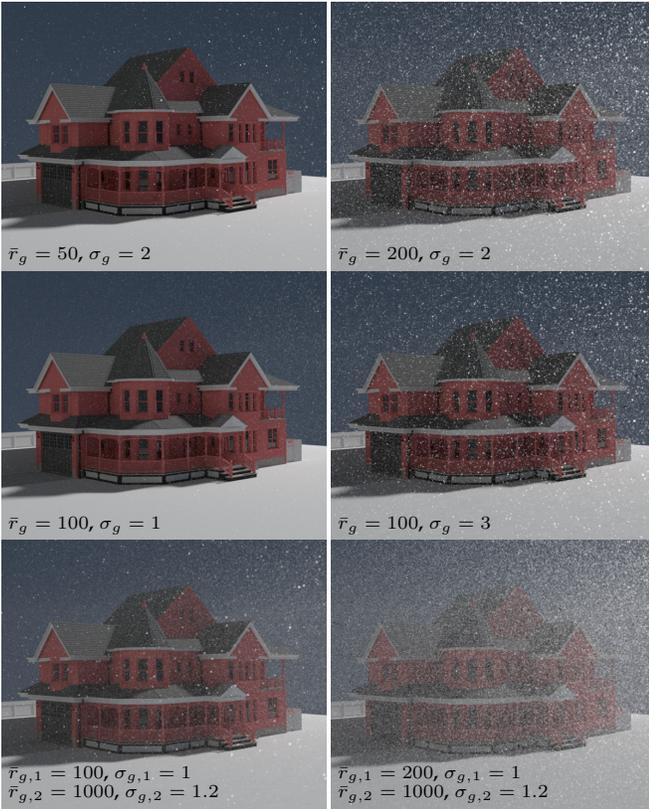


Fig. 19. Impact of the global PSD. First row: Changing \bar{r}_g while fixing σ_g to 2 and N_{tot} to $4 \cdot 10^6$. Second row: Changing σ_g while fixing \bar{r}_g to 100 and N_{tot} to $4 \cdot 10^6$. Third row: Changing $\bar{r}_{g,1}$ in the first mode of the bimodal log-normal distribution while fixing others ($N_{\text{tot}} = 2 \cdot 10^7$).

and $\sigma_{g,1} = 1$, and also generate 10^5 large particles with $\bar{r}_{g,2} = 1000$, $\sigma_{g,2} = 1.2$, leading to $2 \cdot 10^7$ particles in total. Following this complex distribution, we can observe both haziness from massive small particles and graininess from a handful of large particles.

6.8 Performance Analysis

The runtime performance of some test scenes can be found in Table 1. As we mentioned previously, our rendering solution achieves a roughly $1.2\times$ speed improvement over EPT for the LAMP scene and the CORNELL BOX scene. As the particle number N_{tot} increases, the improvement will be more evident. As shown in Fig. 20, $2\times$ speed improvement

TABLE 1

Rendering time performance (in minutes) and memory consumption for six typical scenes shown in this paper. Absorption is not included in the computation.

Scene	N_{tot}	Spp	Memory	Rendering time	
				EPT	Ours
LAMP	$5 \cdot 10^5$	1024	49M	44	36
CORNELL BOX	10^5	1024	60M	16	13
DRAGON	10^5	2048	397M	—	16
BENCH	$2 \cdot 10^5$	2048	52M	—	20
STAIRCASE	$3 \cdot 10^6$	4096	138M	—	53
HOUSE	$4 \cdot 10^6$	1024	545M	—	38

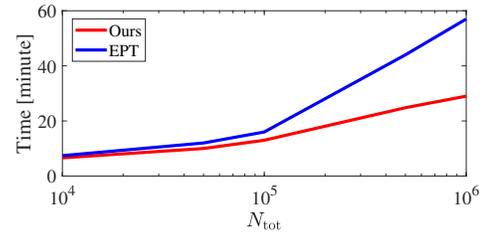


Fig. 20. Rendering time comparison between our method and EPT on the CORNELL BOX scene shown in Fig. 12.

is achieved when N_{tot} increases to 10^6 in the CORNELL BOX scene. Since only the position and the radius of each particle are required, the memory consumption of our rendering solution is affordable even when N_{tot} is very large. The storage scales with the number of particles and the resolution of the grid.

7 LIMITATIONS AND FUTURE WORK

Although our framework has successfully simulated the grainy appearance of discrete participating media, it has several limitations deserving further research.

Non-spherical particles. Our current framework focuses on discrete participating media composed of spherical particles. However, non-spherical particles are also very common, e.g., in the context of rendering rainbow [68] or large snow flakes. Extending our framework to non-spherical particles requires to derive new expressions for the scattering amplitude functions S_1 and S_2 . For some special particle shapes, determining S_1 and S_2 is rather straightforward and analytical expressions exist [61], [66], [67]. For more general shapes, precomputation would be required in practice.

Spatial correlations. As we assume the particles to be sparsely distributed, spatial correlations between particles are not considered and we directly extend the conventional VRE to support multi-scale graininess. However, for densely packed particles the conventional VRE becomes questionable due to strongly correlated scattering effects [12], [13], [14]. Recently, new radiative transfer frameworks dedicated for correlated media are available in computer graphics [5], [6], [9]. It would be an interesting future work to investigate a more general framework supporting both effects.

Mixture of different particles. Another possible direction of future work is to explore an efficient strategy to handle particle mixtures with different compositions. For instance, dusts in real world may be made up of soil particles, textile fibers, human skin cells, etc. A physically-correct participating medium should consider such heterogeneous granular mixtures. For continuous media, this is relatively simple since the concentrations of different grains are fixed [17]. However, for discrete participating media, the concentrations are dynamic and change spatially in a similar way as PSDs [14]. Therefore, determining the bulk optical properties should take spatially-varying concentrations into consideration.

8 CONCLUSION

We have developed a general and physically-based framework for modeling and rendering discrete participating media composed of massive assemblies of independent particles. Notable characteristics of these media include a wide range of PSDs and the appearance of being graininess. To faithfully simulate their appearances, we have derived a novel multi-scale VRE in which a combination of Lorenz-Mie theory and GOA is used to enable high-efficient evaluation of the important optical properties. A Monte Carlo rendering solution is developed to solve the multi-scale VRE with high accuracy and low computational cost. We have extensively evaluated our framework and compared against conventional methods, demonstrating that the proposed framework allows us to reproduce a variety of grainy appearances stemming from different discrete participating media and guarantee temporal coherence in animation. Therefore, for the first time in computer graphics, we have greatly extended the participating media rendering framework to handle a much larger range of particle size statistics. We believe the proposed framework is a further step in computer graphics to manage the details of participating media and expect a more insightful exploration of this phenomenon in the future research.

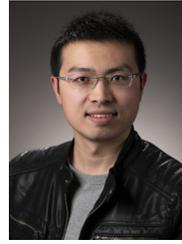
REFERENCES

- [1] E. Cerezo, F. Pérez, X. Pueyo, F. J. Seron, and F. X. Sillion, "A survey on participating media rendering techniques," *The Visual Computer*, vol. 21, no. 5, pp. 303–328, Jun 2005.
- [2] 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)*, vol. 37, no. 2, may 2018.
- [3] S. Chandrasekhar, *Radiative transfer*. Dover, 1960.
- [4] C. F. Bohren and D. R. Huffman, *Absorption and Scattering of Light by Small Particles*. Wiley, 1983.
- [5] A. Jarabo, C. Aliaga, and D. Gutierrez, "A radiative transfer framework for spatially-correlated materials," *ACM Trans. Graph.*, vol. 37, no. 4, pp. 83:1–83:13, Jul. 2018.
- [6] B. Bitterli, S. Ravichandran, T. Müller, M. Wrenninge, J. Novák, S. Marschner, and W. Jarosz, "A radiative transfer framework for non-exponential media," *ACM Trans. Graph.*, vol. 37, no. 6, pp. 225:1–225:17, Dec. 2018.
- [7] E. d'Eon, "A reciprocal formulation of nonexponential radiative transfer. 1: Sketch and motivation," *Journal of Computational and Theoretical Transport*, vol. 47, no. 1-3, pp. 84–115, 2018.
- [8] —, "A reciprocal formulation of nonexponential radiative transfer. 2: Monte carlo estimation and diffusion approximation," 2018.
- [9] 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 Trans. Graph.*, vol. 38, no. 4, pp. 45:1–45:13, Jul. 2019.
- [10] H. C. van de Hulst, *Light scattering by small particles*. Dover, 1981.
- [11] J. Arvo, "Transfer equations in global illumination," in *Global Illumination, SIGGRAPH 93 Course Notes*, 1993.
- [12] J. T. Moon, B. Walter, and S. R. Marschner, "Rendering discrete random media using precomputed scattering solutions," in *Proceedings of the 18th Eurographics Conference on Rendering Techniques*, ser. EGSR'07, 2007, pp. 231–242.
- [13] 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.*, vol. 34, no. 4, pp. 49:1–49:13, Jul. 2015.
- [14] T. Müller, M. Papas, M. Gross, W. Jarosz, and J. Novák, "Efficient rendering of heterogeneous polydisperse granular media," *ACM Trans. Graph.*, vol. 35, no. 6, pp. 168:1–168:14, Nov. 2016.
- [15] L. Lorenz, "Lysbevægelser i og uden for en af plane lysbølger belyst kugle," *Det kongelige danske Videnskabernes Selskabs Skrifter*, pp. 2–62, 1890.
- [16] G. Mie, "Beiträge zur optik trüber medien, speziell kolloidaler metallösungen," *Annalen der Physik*, vol. 330, no. 3, pp. 377–445, 1908.
- [17] J. R. Frisvad, N. J. Christensen, and H. W. Jensen, "Computing the scattering properties of participating media using lorenz-mie theory," in *ACM SIGGRAPH 2007 Papers*, ser. SIGGRAPH '07, 2007.
- [18] W. J. Glantschnig and S.-H. Chen, "Light scattering from water droplets in the geometrical optics approximation," *Appl. Opt.*, vol. 20, no. 14, pp. 2499–2509, Jul 1981.
- [19] A. Ungut, G. Grehan, and G. Gouesbet, "Comparisons between geometrical optics and lorenz-mie theory," *Appl. Opt.*, vol. 20, no. 17, pp. 2911–2918, Sep 1981.
- [20] E. P. Lafortune and Y. D. Willems, "Rendering participating media with bidirectional path tracing," in *EGWR*, Vienna, Jun. 1996, pp. 91–100.
- [21] E. Veach, "Robust monte carlo methods for light transport simulation," Ph.D. dissertation, Stanford, CA, USA, 1997.
- [22] M. Pauly, T. Kollig, and A. Keller, "Metropolis light transport for participating media," in *EGWR*, Vienna, 2000, pp. 11–22.
- [23] L. Szirmay-Kalos, M. Magdics, and M. Sbert, "Multiple scattering in inhomogeneous participating media using rao-blackwellization and control variates," *Computer Graphics Forum*, vol. 37, no. 2, pp. 63–74, 2018.
- [24] H. W. Jensen and P. H. Christensen, "Efficient simulation of light transport in scenes with participating media using photon maps," in *SIGGRAPH*, Jul. 1998, pp. 311–320.
- [25] W. Jarosz, M. Zwicker, and H. W. Jensen, "The beam radiance estimate for volumetric photon mapping," *Computer Graphics Forum*, vol. 27, no. 2, pp. 557–566, Apr. 2008.
- [26] 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.*, vol. 30, no. 1, pp. 5:1–5:19, Feb. 2011.
- [27] W. Jarosz, D. Nowrouzezahrai, R. Thomas, P.-P. Sloan, and M. Zwicker, "Progressive photon beams," *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, vol. 30, no. 6, Dec. 2011.
- [28] T. Hachisuka, W. Jarosz, I. Georgiev, A. Kaplanyan, and D. Nowrouzezahrai, "State of the art in photon density estimation," in *ACM SIGGRAPH Asia Courses*, nov 2013.
- [29] B. Bitterli and W. Jarosz, "Beyond points and beams: Higher-dimensional photon samples for volumetric light transport," *SIGGRAPH*, vol. 36, no. 4, pp. 112:1–112:12, Jul. 2017.
- [30] X. Deng, S. Jiao, B. Bitterli, and W. Jarosz, "Photon surfaces for robust, unbiased volumetric density estimation," *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, vol. 38, no. 4, Jul. 2019.

- [31] 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.*, vol. 33, no. 4, pp. 103:1–103:13, Jul. 2014.
- [32] E. Veach and L. J. Guibas, "Optimally combining sampling techniques for monte carlo rendering," in *Proceedings of SIGGRAPH '95*, ser. SIGGRAPH '95. Association for Computing Machinery, 1995, pp. 419–428.
- [33] B. Müller, I. Georgiev, and W. Jarosz, "A null-scattering path integral formulation of light transport," *ACM Trans. Graph.*, vol. 38, no. 4, Jul. 2019.
- [34] W. Jakob, A. Arbree, J. T. Moon, K. Bala, and S. Marschner, "A radiative transfer framework for rendering materials with anisotropic structure," *ACM Trans. Graph.*, vol. 29, no. 4, pp. 53:1–53:13, Jul. 2010.
- [35] E. Heitz, J. Dupuy, C. Crassin, and C. Dachsbacher, "The sgx microflake distribution," *ACM Trans. Graph.*, vol. 34, no. 4, pp. 48:1–48:11, Jul. 2015.
- [36] J. Dupuy, E. Heitz, and E. d'Eon, "Additional Progress Towards the Unification of Microfacet and Microflake Theories," in *Eurographics Symposium on Rendering - Experimental Ideas & Implementations*, E. Eisemann and E. Fiume, Eds. The Eurographics Association, 2016.
- [37] S. Zhao, W. Jakob, S. Marschner, and K. Bala, "Building volumetric appearance models of fabric using micro ct imaging," *ACM Trans. Graph.*, vol. 30, no. 4, pp. 44:1–44:10, Jul. 2011.
- [38] —, "Structure-aware synthesis for predictive woven fabric appearance," *ACM Trans. Graph.*, vol. 31, no. 4, pp. 75:1–75:10, Jul. 2012.
- [39] P. Khungurn, D. Schroeder, S. Zhao, K. Bala, and S. Marschner, "Matching real fabrics with micro-appearance models," *ACM Trans. Graph.*, vol. 35, no. 1, 2016.
- [40] S. Zhao, F. Luan, and K. Bala, "Fitting procedural yarn models for realistic cloth rendering," *ACM Trans. Graph.*, vol. 35, no. 4, Jul. 2016.
- [41] C. Aliaga, C. Castillo, D. Gutierrez, M. A. Otaduy, J. Lopez-Moreno, and A. Jarabo, "An appearance model for textile fibers," *Computer Graphics Forum*, vol. 36, no. 4, pp. 35–45, 2017.
- [42] S. Zhao, L. Wu, F. Durand, and R. Ramamoorthi, "Downsampling scattering parameters for rendering anisotropic media," *ACM Trans. Graph.*, vol. 35, no. 6, pp. 166:1–166:11, Nov. 2016.
- [43] G. Loubet and F. Neyret, "A new microflake model with microscopic self-shadowing for accurate volume downsampling," *Computer Graphics Forum*, vol. 37, no. 2, pp. 111–121, May 2018.
- [44] L.-Q. Yan, M. Hašan, W. Jakob, J. Lawrence, S. Marschner, and R. Ramamoorthi, "Rendering glints on high-resolution normal-mapped specular surfaces," *ACM Trans. Graph.*, vol. 33, no. 4, pp. 116:1–116:9, Jul. 2014.
- [45] L.-Q. Yan, M. Hašan, S. Marschner, and R. Ramamoorthi, "Position-normal distributions for efficient rendering of specular microstructure," *ACM Trans. Graph.*, vol. 35, no. 4, pp. 56:1–56:9, Jul. 2016.
- [46] L.-Q. Yan, M. Hašan, B. Walter, S. Marschner, and R. Ramamoorthi, "Rendering specular microgeometry with wave optics," *ACM Trans. Graph.*, vol. 37, no. 4, 2018.
- [47] S. Mérillou, J. Dischler, and D. Ghazanfarpour, "Surface scratches: measuring, modeling and rendering," *The Visual Computer*, vol. 17, no. 1, pp. 30–45, Feb. 2001.
- [48] C. Bosch, X. Pueyo, S. Mérillou, and D. Ghazanfarpour, "A physically-based model for rendering realistic scratches," *Computer Graphics Forum*, vol. 23, no. 3, pp. 361–370, 2004.
- [49] B. Raymond, G. Guennebaud, and P. Barla, "Multi-scale rendering of scratched materials using a structured sv-brdf model," *ACM Trans. Graph.*, vol. 35, no. 4, Jul. 2016.
- [50] S. Werner, Z. Velinov, W. Jakob, and M. B. Hullin, "Scratch iridescence: Wave-optical rendering of diffractive surface structure," *ACM Trans. Graph.*, vol. 36, no. 6, Nov. 2017.
- [51] A. Kuznetsov, M. Hašan, Z. Xu, L.-Q. Yan, B. Walter, N. K. Kalantari, S. Marschner, and R. Ramamoorthi, "Learning generative models for rendering specular microgeometry," *ACM Trans. Graph.*, vol. 38, no. 6, 2019.
- [52] W. Jakob, M. Hašan, L.-Q. Yan, J. Lawrence, R. Ramamoorthi, and S. Marschner, "Discrete stochastic microfacet models," *ACM Trans. Graph.*, vol. 33, no. 4, pp. 115:1–115:10, Jul. 2014.
- [53] J. Guo, Y. Chen, Y. Guo, and J. Pan, "A Physically-based Appearance Model for Special Effect Pigments," *Computer Graphics Forum*, 2018.
- [54] B. Wang, L. Wang, and N. Holzschuch, "Fast Global Illumination with Discrete Stochastic Microfacets Using a Filterable Model," *Computer Graphics Forum*, 2018.
- [55] H. Rushmeier, "Input for participating media," in *ACM SIGGRAPH 1995 Courses*, ser. SIGGRAPH '95, 1995.
- [56] P. Callet, "Pertinent data for modelling pigmented materials in realistic rendering," *Computer Graphics Forum*, vol. 15, no. 2, pp. 119–127, 1996.
- [57] D. Jackèl and B. Walter, "Modeling and rendering of the atmosphere using mie-scattering," *Computer Graphics Forum*, vol. 16, no. 4, pp. 201–210, 1997.
- [58] T. Nishita and Y. Dobashi, "Modeling and rendering of various natural phenomena consisting of particles," in *Computer Graphics International 2001 (CGI'01)*, 2001, pp. 149–156.
- [59] K. Riley, D. S. Ebert, M. Kraus, J. Tessendorf, and C. Hansen, "Efficient rendering of atmospheric phenomena," in *Proceedings of the Fifteenth Eurographics Conference on Rendering Techniques*, ser. EGSR'04, 2004, pp. 375–386.
- [60] P. Laven, "Simulation of rainbows, coronas, and glories by use of mie theory," *Appl. Opt.*, vol. 42, no. 3, pp. 436–444, Jan 2003.
- [61] E. A. Hovenac, "Calculation of far-field scattering from nonspherical particles using a geometrical optics approach," *Appl. Opt.*, vol. 30, no. 33, pp. 4739–4746, Nov 1991.
- [62] X. Zhou, S. Li, and K. Stamnes, "Geometrical-optics code for computing the optical properties of large dielectric spheres," *Appl. Opt.*, vol. 42, no. 21, pp. 4295–4306, Jul 2003.
- [63] L. Wu, H. Yang, X. Li, B. Yang, and G. Li, "Scattering by large bubbles: Comparisons between geometrical-optics theory and debye series," *Journal of Quantitative Spectroscopy and Radiative Transfer*, vol. 108, no. 1, pp. 54 – 64, 2007.
- [64] H. Yu, J. Shen, and Y. Wei, "Geometrical optics approximation of light scattering by large air bubbles," *Particology*, vol. 6, no. 5, pp. 340 – 346, 2008.
- [65] —, "Geometrical optics approximation for light scattering by absorbing spherical particles," *Journal of Quantitative Spectroscopy and Radiative Transfer*, vol. 110, no. 13, pp. 1178 – 1189, 2009.
- [66] H. He, W. Li, X. Zhang, M. Xia, and K. Yang, "Light scattering by a spheroidal bubble with geometrical optics approximation," *Journal of Quantitative Spectroscopy and Radiative Transfer*, vol. 113, no. 12, pp. 1467 – 1475, 2012.
- [67] Y. F. Lu, Y. P. Han, J. J. Wang, and Z. W. Cui, "Geometrical optics approximation for forward light scattering by a large chiral sphere," *Journal of Quantitative Spectroscopy and Radiative Transfer*, vol. 228, pp. 90 – 96, 2019.
- [68] I. Sadeghi, A. Munoz, P. Laven, W. Jarosz, F. Seron, D. Gutierrez, and H. W. Jensen, "Physically-based simulation of rainbows," *ACM Transactions on Graphics*, vol. 31, no. 1, pp. 3:1–3:12, 2012.
- [69] I. Gkioulekas, S. Zhao, K. Bala, T. Zickler, and A. Levin, "Inverse volume rendering with material dictionaries," *ACM Trans. Graph.*, vol. 32, no. 6, 2013.
- [70] A. Dal Corso, J. R. Frisvad, T. K. Kjeldsen, and J. A. Bærentzen, "Interactive Appearance Prediction for Cloudy Beverages," in *Workshop on Material Appearance Modeling*, R. Klein and H. Rushmeier, Eds. The Eurographics Association, 2016.
- [71] V. E. Cachorro and L. L. Salcedo, "New improvements for mie scattering calculations," *Journal of Electromagnetic Waves and Applications*, pp. 913–926, 1991.
- [72] R. A. Dobbins and T. I. Eklund, "Ripple structure of the extinction coefficient," *Applied Optics*, vol. 16, pp. 281–282, Feb. 1977.
- [73] J. Amanatides and A. Woo, "A fast voxel traversal algorithm for ray tracing," in *In Eurographics '87*, 1987, pp. 3–10.
- [74] I. Wald, T. Ize, A. Kensler, A. Knoll, and S. G. Parker, "Ray tracing animated scenes using coherent grid traversal," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 485–493, Jul. 2006.
- [75] W. Jakob, "Mitsuba renderer," 2010, <http://www.mitsuba-renderer.org>.
- [76] L. G. Henyey and J. L. Greenstein, "Diffuse radiation in the Galaxy," *The Astrophysical Journal*, vol. 93, pp. 70–83, Jan 1941.
- [77] D. Toublanc, "Henyey-greenstein and mie phase functions in monte carlo radiative transfer computations," *Appl. Opt.*, vol. 35, no. 18, pp. 3270–3274, Jun 1996.
- [78] T. Hawkins, P. Einarsson, and P. Debevec, "Acquisition of time-varying participating media," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 812–815, Jul. 2005.



Jie Guo received his bachelor's degree and Ph.D. degree in computer science from Nanjing University in 2008 and 2013, respectively. Currently, he is an assistant researcher at State Key Laboratory for Novel Software Technology of Nanjing University. He was a research intern at Microsoft Research Asia. His research interests include appearance modeling, real-time rendering, and virtual reality.



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.



Bingyang Hu is a graduate student in the Department of Computer Science and Technology, Nanjing University. His research interests include realistic rendering and appearance modeling.



Yanjun Chen is a graduate student in the Department of Computer Science and Technology, Nanjing University. His research interests include realistic rendering and appearance modeling.



Yuanqi Li is a PhD student in the Department of Computer Science and Technology, Nanjing University. He received his bachelor degree from the same university in 2018. His research interests include realistic rendering and 3D vision.



Yanwen Guo received the Ph.D. degree in applied mathematics from the State Key Lab of CAD&CG, Zhejiang University, China, in 2006. He is currently a Full Professor with the National Key Laboratory for Novel Software Technology, Department of Computer Science and Technology, Nanjing University, Jiangsu, China. He was a Visiting Professor with the Department of Computer Science and Engineering, The Chinese University of Hong Kong, in 2006 and 2009, respectively, the Department of Computer Science, The University of Hong Kong, in 2008, 2012, and 2013, respectively, and a Visiting Scholar with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, from 2013 to 2015. His research interests include image and video processing, vision, and computer graphics.

ence, The University of Hong Kong, in 2008, 2012, and 2013, respectively, and a Visiting Scholar with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, from 2013 to 2015. His research interests include image and video processing, vision, and computer graphics.