

Rendering Discrete Participating Media using Geometrical Optics Approximation

Jie Guo¹, Bingyang Hu¹, Yanjun Chen¹, Yuanqi Li¹, Yanwen Guo¹✉ and Ling-Qi Yan²

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

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 several orders of magnitude greater, resulting in an appearance with distinct graininess as opposed to the smooth appearance of continuous media. One fundamental issue in the physically-based synthesis of such appearance is to determine the necessary optical properties in every local region. Since these 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 for any particle size distribution. We then use a practical Monte Carlo rendering solution to solve energy transfer in the discrete participating media. Our proposed framework is the first to simulate a wide range of discrete participating media with different levels of graininess, converging to the continuous media case as the particle concentration increases.

Keywords 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 it plausibly and efficiently [10, 53]. From the physics point of view, radiative transfer is rather complicated and must be rigorously derived from Maxwell's electromagnetic theory. To make simulation tractable, various compromises have been made in physically-based rendering over the past decades. Two main assumptions are *independent scattering* [11] and *local continuity* (or *statistical homogeneity* [6]). The first 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 by incorporating spatial correlations between scatterers into radiative transfer frameworks [5, 14, 15, 22, 33], 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 every point of the medium, resulting in locally smooth rendering. However, many participating media are composed of separate particles distributed randomly within a given volume, and scattering happens only at particle positions [6, 62]. These facts indicate that the assumption of continuous media only holds when the particle size is much smaller than the resolution of the sensor (e.g. the human eye) and the quantity is sufficiently large [3]. Otherwise, individual grains can be observed when zooming in. To overcome this constraint, the graininess of the medium should be taken into consideration.

Several recent studies in computer graphics [46, 50, 51] have noticed the graininess in rendering

1 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.

2 Department of Computer Science, UC Santa Barbara, United States. E-mail: lingqi@cs.ucsb.edu.

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

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 number of very large particles using 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 number of scatterers whose particle size distribution (PSD) ranges widely. These media are omnipresent in natural and artificial environments, such as airborne dust, blowing snow, powder suspensions, 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 must be evaluated on-the-fly. These properties are closely related to the scattering behavior of each individual particle and fluctuations in 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 [43, 48] which can provide substantial realism in rendering participating media [18]. However, numerical evaluation of the Lorenz-Mie coefficients is known to be difficult and time-consuming as the particle size increases [20]. To ameliorate this issue, we introduce geometrical optics approximation (GOA) [20, 61] and use 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 optical properties, enabling both high accuracy and low computational cost.

Due to the variations in local PSDs, 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 given 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 can 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 in 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 [10, 11] by means of Monte Carlo path integration [41, 53, 54, 59, 63], photon density estimation [4, 13, 23, 34–37], or a combination of both [40]. In our current framework, we choose Monte Carlo path integration because of its elegant simplicity, generality, and accuracy. This technique operates by stochastically constructing a large number of light paths between receivers 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 approach [34, 36, 40]. Multiple importance sampling [49, 64] 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. The recent trend in computer graphics is to try 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 can be extended with local directional dependency based on a microflake model [17, 26, 31]. It is also possible to extend the VRE to simulate the effects of spatial correlations [5, 22, 33], yielding non-exponential attenuation of light. Notably, these methods still assume the media to be statistically homogeneous within each differential volume [6], 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 rendering for fabrics and textiles [1, 38, 74–76]. Unfortunately, these methods are highly data-intensive and plagued by heavy computation. To improve the performance while maintaining good accuracy, some downsampling strategies [44, 77] have been developed. Current rendering solutions for granular materials are also based on explicit geometries and pre-captured optical properties of each individual grain [46, 50, 51]. 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. [69, 70] 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 [71]. Other methods focus specifically on capturing spatially-varying highlights from scratched surfaces, under either geometric optics [7, 47, 55] or wave optics [67]. To ease the burden of computation and storage, Kuznetsov et al. [39] proposed to learn high-frequency angular patterns from existing examples, using a generative adversarial network (GAN). Jakob et al. [32] 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 [21] and allow fast global illumination [66].

2.4 Lorenz-Mie Theory and GOA

Lorenz-Mie theory [43, 48] develops a rigorous solution to the problem of light scattering by spherical particles. It was introduced to the graphics community by Rushmeier [57] to accurately simulate the physics of light transport in participating media. Later, Callet [9] used this theory to model pigmented materials consisting of pigmented particles in a transparent solvent. Atmospheric phenomena, such as halos and rainbows, are especially well handled

Tab. 1 Abbreviations used in the paper.

Abbreviation	Meaning
3D-DDA	3D Digital Differential Analyzer
EPT	Explicit Path Tracing
GOA	Geometrical Optics Approximation
HG	Henvey and Greenstein
PSD	Particle Size Distribution
RelMSE	Relative Mean Squared Error
VRE	Volumetric Rendering Equation

by this theory [29, 42, 52, 56]. Frisvad et al. [18] 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 [20, 28, 61, 68, 72, 73, 78] 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 including absorbing particles [73] and aspherical particles [25, 28, 45, 58]. 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

3.1 Preliminaries

We first consider light scattering by a single particle. Meanings of some abbreviations used in the paper are listed in Table 1. We suppose that the particle is approximately spherical and has a set of physical properties including its radius r and a 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 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 polarization, 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 [6]:

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

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 given by

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

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 scattering of light, an electromagnetic wave, from a homogeneous spherical particle, exact solutions of the two scattering amplitude functions are given by Lorenz-Mie theory [43, 48]. Its accuracy has been validated against real measurements in various literature [12, 19].

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 computed in the series for $S_1(\theta, r)$ and $S_2(\theta, r)$ is proportional to the size parameter α [18, 20]; it is suggested that an appropriate number of terms to sum is $\lceil |\alpha| + 4.3|\alpha|^{1/3} + 1 \rceil$ [8]. For this reason, simpler approximate expressions are useful 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) [20, 28, 61, 68, 72, 73, 78] provides a simplified but good solution.

3.2 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. its refractive index). Its amplitude functions for the forward direction are readily described by Fraunhofer




Fig. 1 Light scattering by a spherical particle according to GOA.

diffraction theory as [20]

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

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 $q = 1$ indicates that the incident ray hits the particle on the upper hemisphere and $q = -1$ on the lower, and l is an integer ensuring that the scattering angle θ is restricted to the range 0 to π .

Clearly, reflected and transmitted rays depend on the shape and composition of the particle. Their scattering amplitudes for each polarization can be derived to be [61, 62]:

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

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 get the total amplitude functions of GOA as

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

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

In GOA, an analytical expression for $C_t(r)$ can be derived:

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

where $\mathcal{P} = \{1, 3, 5, \dots\}$. This expression is fast to evaluate and captures ripple structures well [16].

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 [6]

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

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.

To our knowledge, this is the first time that GOA has been introduced into computer graphics. The benefit of GOA is its comparable accuracy to rigorous Lorenz-Mie theory and its high efficiency in describing and evaluating key optical properties (e.g. the extinction and absorption coefficients) that are necessary for rendering.

3.3 Number of Terms

An infinite summation is needed in principle in GOA to compute S_1 and S_2 . However, unlike the one needed 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 for $p = 100$, 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. Thus, p may be safely set to 3 in what follows.




Fig. 2 Impact of p on $(|S_1| + |S_2|)/2$ 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 . We compare our GOA calculation (red: $p = 3$, green: $p = 1$) to that of Lorenz-Mie theory (blue) with $\eta = 1.33$ and $\lambda = 0.6 \mu\text{m}$. The ripple structures [16] are well captured by GOA as 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 light rays with $p > 1$ contribute little to forward scattering: see Fig. 3, where the C_t curves for $p = 1$ (green) and $p = 3$ (red) are virtually indistinguishable for a very wide range of r .

3.4 GOA and Lorenz-Mie Theory

To investigate the range of validity of GOA in simulating the scattering patterns of spherical particles, we compare its results with the rigorous Lorenz-Mie results for 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 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 particle radius decreases further.

To show the influence of these discrepancies on the perception of translucent appearance, we rendered a smooth medium comprising monodisperse particles of

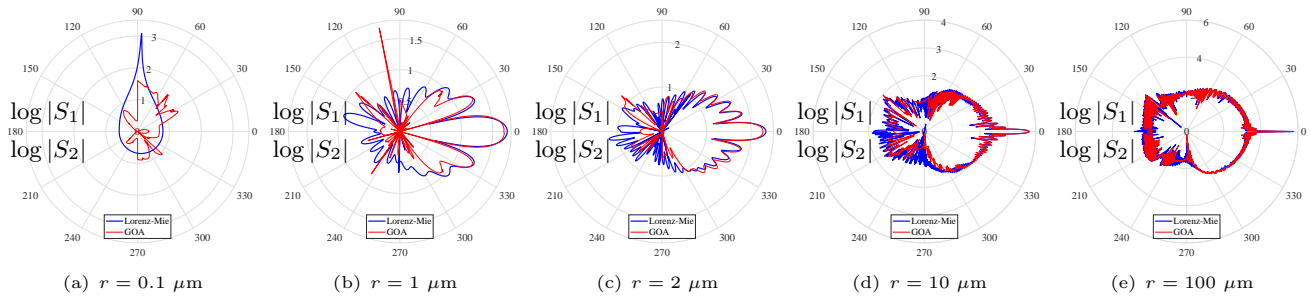


Fig. 4 Comparison of $\log |S_1|$ (upper half) and $\log |S_2|$ (lower half) computed by Lorenz-Mie theory (blue curves) and GOA (red curves) for $\eta = 1.33$ and $\lambda = 0.6 \mu\text{m}$. Particle radii shown are $r = 0.1, 1, 2, 10$ and $100 \mu\text{m}$, respectively.

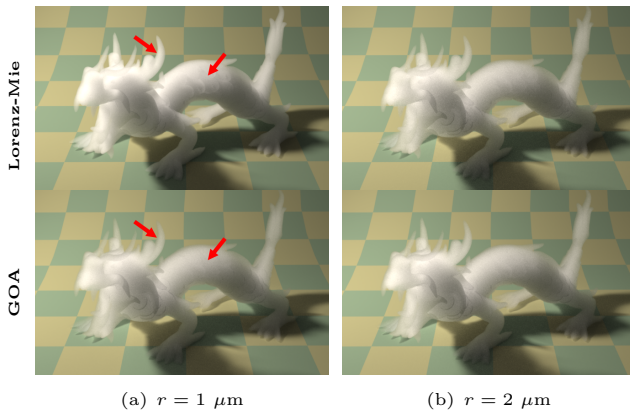


Fig. 5 Rendering a smooth medium with phase functions derived from Lorenz-Mie theory (top row) and GOA (bottom row), respectively. Here, $\eta = 1.33$. The red arrows highlight the differences.

radius r . We determined the phase function using either Lorenz-Mie theory or GOA according to r ; these phase functions were precomputed and stored in tables. The extinction coefficient was set to a constant 0.6 for a fair comparison. The synthesized images are presented in Fig. 5 with r set 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 inconsistent appearance. However, this inconsistency almost disappears completely as r increases to $2 \mu\text{m}$, although there are some mismatches on the backward peaks of $S_1(\theta)$ (see Fig. 4(c); $S_2(\theta)$ is similar). Figs. 4 and 5 together validate choosing GOA to compute $S_1(\theta)$ and $S_2(\theta)$ when $r \geq 2 \mu\text{m}$.

To further show the similarity between Lorenz-Mie theory and GOA in computing $S_1(\theta)$ and $S_2(\theta)$, we report their relative mean squared error (RelMSE) in Fig. 6(left). RelMSE is computed on $(|S_1(\theta)| + |S_2(\theta)|)/2$ for different relative refractive indexes. The results confirm that subtle errors exist when r is large: despite some fluctuations, these calculations of GOA




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 .

exhibit errors of less than 0.01 as compared with exact Lorenz-Mie calculations when $r \geq 2 \mu\text{m}$.

When calculating $(|S_1(\theta)| + |S_2(\theta)|)/2$, 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 to 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 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 an accurate method when required.

4 Bulk Optical Properties with Graininess

Next, we consider light scattering by a cloud of spherical particles of the same composition but 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 much larger. 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 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. $N(r) dr$ is the total concentration (number of particles per unit volume) of particles with sizes in the range $[r, r + dr]$. The total particle 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 $N(r)/N_0$.

It is generally reported that particle sizes closely follow a log-normal distribution [18, 58]:

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

in which σ_g is the geometric standard deviation and \bar{r}_g is the geometric mean radius. Obviously, this statistical tendency stems from 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 generated 2000 particles in a box with uniformly distributed positions and log-normal distributed radii. A visualization of these particles and their PSD are shown in Fig. 7(above). We then extracted three small patches from the box and estimate their actual PSDs by binning. As expected, the PSDs plotted in Fig. 7(below) 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 Fig. 7(above). 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 coefficient¹ σ_t of this volume is

¹In rendering literature, the symbol σ refers to the cross section sometimes, while using μ for the coefficient.




Fig. 7 Spatially-varying PSD. Above: 2000 polydisperse particles and their global PSD. Below: Actual PSDs for particles in three small patches.

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 $\mathcal{V}_{\mathbf{x}}$, and $r_{\min}(\mathcal{V}_{\mathbf{x}})$ and $r_{\max}(\mathcal{V}_{\mathbf{x}})$ are 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. The scattering coefficient and the absorption coefficient can be defined in a similar way by replacing $C_t(r)$ with $C_s(r)$ and $C_a(r)$, respectively². 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} (which is similar to the log-normal distribution with a very large σ_g). 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 to 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

²These 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 $\sigma_a(\mathbf{x}) = \sigma_a(\mathcal{V}_{\mathbf{x}})$.

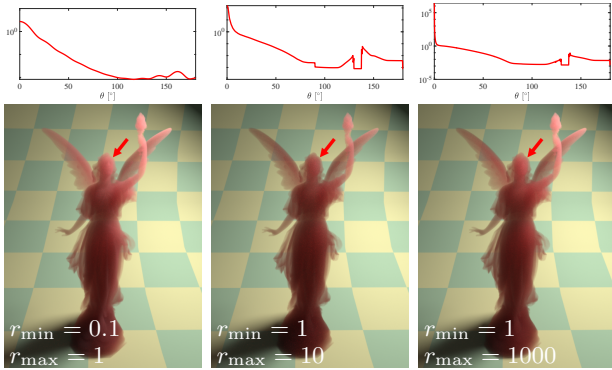


Fig. 8 Above: Phase functions estimated by our method. Below: Corresponding renderings of a smooth homogeneous medium using these phase functions, keeping other properties fixed. Red arrows highlight differences.

light beam of length s as $T_{\mathcal{A}}(s)$

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

in which $\mathcal{C} = \mathcal{A} \times s$ represents a small cylinder around the light beam with cross section \mathcal{A} .

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

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

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

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. We also develop a Monte Carlo sampling based method to solve the VRE; it only requires the position and size of each particle, avoiding 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 used for acceleration.

5.1 Multi-scale VRE

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

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




Fig. 9 A query cylinder and active voxels in our ray traversal algorithm. Only the particles contained in active voxels are queried.

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, this is a thin cylinder centered around $\mathbf{x} \rightarrow \mathbf{y}$, as illustrated in Fig. 9. We call such a cylinder 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, making it too small one may cause

numerical issues and large variance. Conversely, bias will be introduced in $T_{\mathcal{A}}(\mathbf{x}, \mathbf{y})$ if \mathcal{A} is too 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

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

Here, S_{pix} is the pixel's size and k can be viewed as the proportion of the pixel's footprint at 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 tests for every particle in the medium. A particle with position \mathbf{p} and radius r_p is 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 traversal in the medium using a 3D digital differential analyzer (3D-DDA) [2, 65]. Specifically, we construct a uniform grid for the medium and adopt a ray traversal algorithm, similar to that in [2], to find the active voxels intersected by the query cylinder. Fig. 9(right) illustrates all active voxels corresponding to the orange query cylinder. Only those particles inside the active voxels are 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 over 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 [36]. After collecting all 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 spatial subdivision.

5.4 Computing $Q(\theta)$

To solve the multi-scale VRE, we also have to compute $Q(\theta)$ at each 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 $Q(\theta)$ using 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

Following the simplifications used in rendering surface glints [21, 32], 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 during preprocessing, assuming the entire bounding box of the medium to be $\mathcal{V}_{\mathbf{x}}$ in Eqs. (16) and (17).

6 Results

6.1 Setting

We have implemented our rendering solution on top of the Mitsuba renderer [30], with spectral rendering enabled. We use 8 spectral samples in the range of the visible spectrum at equally-spaced locations [71]. After rendering, we convert the spectral values to the sRGB color space. All synthesized images were created on a PC with an Intel 16-core i7-6900K CPU and 16 GB 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 ($r_{\text{max}} = 2000$) to avoid unreasonably large particles which are unusual and unsuitable for treatment 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 for the STAIRCASE scene, the refractive index is set using data for ice from [18], and the host medium is set to air with $\eta_m = 1$.

6.2 Comparison to Explicit Path Tracing

We first compare our method to traditional path tracing. Previous methods simulating the grainy appearance of discrete participating media mostly rely on explicit path tracing (EPT), possibly with approximations to simplify the computation of high-order scattering [46, 50, 51]. However, EPT and other




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




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

approximations are restricted to geometric optics. This means that only surface reflection and refraction are properly handled. In principle, a mesh should be explicitly generated for every particle and costly ray-object intersections are required.

Compared to EPT, our method offers at least two benefits. First, particle scattering is considered which includes Fraunhofer diffraction and phase differences, etc. 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 produced and rendered 5×10^5 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 (above), the energy inherently belonging to Fraunhofer diffraction is not correctly captured by EPT, resulting in overly sparse and specular volumetric glints in this LAMP scene. The importance of Fraunhofer diffraction is shown in Fig. 11, which plots the percentage of energy contributed by Fraunhofer diffraction at different scattering angles and for different sized particles. Obviously, Fraunhofer diffraction cannot be ignored, especially for small particles. Moreover, EPT easily misses many small particles that are hard gather along




Fig. 12 Comparisons between EPT and our method with or without Fraunhofer diffraction (-FD) when rendering 10^5 large particles of the same size ($1000 \mu\text{m}$). The impact of 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 [27]. The scattering model generated by Lorenz-Mie theory (green curve, right) is also provided for reference. Due to obvious discrepancies between our estimated phase function (red curve, right) and the HG phase function (blue curve, right), the renderings differ slightly (see the close-ups).

an ordinary light path, leading to a slow convergence rate. Also, even when substantially increasing the sampling rate or using ‘fat ray’ tracing as in our method, the appearance is still quite sparse. On the contrary, our method (below) preserves the energy from Fraunhofer diffraction and produces a smoother appearance that is closer to reality, thanks to the Airy pattern³ caused by 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 a 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 becomes applicable and serves as a valid approximation to particle scattering. This is further verified in the third column of Fig. 11. We see that the contribution of Fraunhofer diffraction is concentrated in a very narrow angle in this case, making it hard to observe. Consequently, if we remove Fraunhofer diffraction from our model and leave

³Fraunhofer diffraction distorts light passing through the particle, making it appear larger—the smaller the particle, the larger the light’s distribution.




Fig. 14 Visual comparisons with continuous media for 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 for a very dense medium ($r = 400 \mu\text{m}$ and $N_{\text{tot}} = 10^8$).

only reflection and refraction (see the third column of Fig. 12), we achieve a grainy appearance similar to that of our full model. However, there are still subtle differences (highlighted in the difference image) contributed by Fraunhofer diffraction.

The second strength of our rendering solution lies in its 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 rendering time.

6.3 Comparison to 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 [27], because of its simplicity and well-defined behavior. However, the HG model is not very accurate, as pointed out by various works [19, 24, 60]. In Fig. 13 we compare our estimated phase function with the best-fit HG phase function when 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 from the best-fit HG phase function is slightly different from ours (see insets). Recall that our rendering is close to that generated using EPT for this specific scene with large particles.

6.4 Comparison to Continuous Media

We also compared 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 volume V . In this configuration, it is easy to derive the extinction coefficient and the scattering coefficient for 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 becomes hazy, and the rendering result becomes closer to that of the corresponding continuous medium. When N_{tot} is sufficiently large (e.g. $N_{\text{tot}} = 10^7$), the discrete medium and its continuous counterpart provide quite similar appearance. In Fig. 15, a much denser medium ($N_{\text{tot}} = 10^8$) is rendered by our method, which again produces a smooth appearance matching that from a continuous medium. These pair-wise comparisons demonstrate that our rendering solution converges to traditional volumetric rendering of continuous media in the limit of particle concentration.




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




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

6.5 Query Cylinder 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 choosing its value in the range $[0.5, 1]$ which yields a reasonable grainy appearance as shown Fig. 16(2nd column). Values in this range allow us to faithfully capture almost all grains in scenes with little bias. Generally, setting k too large produces undesirable aliasing as shown in Fig. 16(columns 3,4). In these two cases, although the overall brightness is similar to that for $k = 0.5$, the volumetric glints are overly blurred. On the other hand, a too small k will miss many particles during querying 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 may sometimes be acceptable. For instance, the first image in the bottom row of Fig. 16 is rendered with $k = 0.2$ which achieves a similar effect to that for $k = 0.5$ for this relatively dense medium ($N_{\text{tot}} = 10^5$).

6.6 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 Fig. 17(right), the grid resolution (Res) influences the performance. Although no analytical analysis can be provided to select the best resolution, we empirically observe that a grid resolution yielding roughly one particle per voxel achieves the optimum solution for most scenes. The BENCH scene in Fig. 17(left) contains 2×10^5 particles and best performance is achieved 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 $60 \times 60 \times 60$ has very poor performance since too many particles reside in each voxel. For other scenes, a similar conclusion can be drawn.

6.7 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 with the imaginary part of the complex refractive index, η_i . The extinction cross section must be slightly modified according to [73]. Results of doing so are shown in Fig. 18. This STAIRCASE scene aims to simulate floating dust in a dirty room lit by a local area light through the window. As expected, increasing η_i increasingly reduces the intensity of scattering.

6.8 Impact of Global PSD

Fig. 19 analyzes the impact of the global PSD on the appearance of discrete participating media. Here, we generated $N_{\text{tot}} = 4 \times 10^6$ particles with different global PSDs in the HOUSE scene. This scene is designed to simulate the appearance of blowing snow 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 particle radius. The geometric standard deviation σ_g 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 a log-normal distribution. Other distributions also work. For instance, in the third row of Fig. 19, we show a cloud of blowing snow following a bimodal log-normal distribution. We generated 1.99×10^7 small particles with $\bar{r}_{g,1} = 100$ (or $\bar{r}_{g,1} = 200$) and $\sigma_{g,1} = 1$, and also generated 10^5 large particles with $\bar{r}_{g,2} = 1000$,




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 $N_{\text{tot}} = 3 \times 10^5$, $\bar{r}_g = 800$ and $\sigma_g = 4$.

Tab. 2 Rendering time (in minutes) and memory consumption for six typical scenes shown in this paper. Absorption is not included in the computation.

Scene	N_{tot}	SPP	Mem	Time	
				EPT	Ours
LAMP	5×10^5	1024	49M	44	36
CORNELL BOX	10^5	1024	60M	16	13
DRAGON	10^5	2048	397M	—	16
BENCH	2×10^5	2048	52M	—	20
STAIRCASE	3×10^6	4096	138M	—	53
HOUSE	4×10^6	1024	545M	—	38

$\sigma_{g,2} = 1.2$, leading to 2×10^7 particles in total. Using this complex distribution, we can observe both haziness from massive small particles and graininess from a handful of large particles.

6.9 Performance

The runtime performance for some test scenes can be found in Table 2. As mentioned previously, our rendering solution achieves roughly $1.2\times$ the speed of EPT for the LAMP scene and the CORNELL BOX scene. As the number of particles N_{tot} increases, the improvement becomes more evident. As shown in Fig. 20, twice the speed is achieved when N_{tot} is increased to 10^6 for 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 successfully simulates the grainy appearance of discrete participating media, it has several limitations deserving further research.

7.1 Aspherical particles

Our current framework focuses on discrete participating media composed of spherical particles. However, aspherical particles are also very common, e.g., in the context of rendering rainbows [58] or large snow flakes. Extending our framework to aspherical particles requires the derivation of 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 [25, 28, 45]. For more general shapes, precomputation would be required in practice.

7.2 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 [46, 50, 51]. Recently, new radiative transfer frameworks dedicated for correlated media have become available in computer graphics [5, 22, 33]. It would be interesting future work to investigate a more general framework supporting both effects.

7.3 Mixed particles

Another possible direction for future work is to explore an efficient strategy to handle particle mixtures with different compositions. For instance, dust in the 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 [18]. However, for discrete participating media, the concentrations are dynamic and change spatially as for PSDs [51]. Determining the




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

bulk optical properties should take spatially-varying concentrations into consideration.

8 Conclusion

We have developed a general 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 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-efficiency evaluation of the important optical properties. A Monte Carlo rendering solution is used to solve the multi-scale VRE with high accuracy and low computational cost. We have extensively evaluated our framework and compared it to conventional methods, demonstrating that the proposed framework allows us to reproduce a variety of grainy appearances stemming from different discrete participating media. We can guarantee temporal coherence in animation. We have greatly extended the participating media rendering framework to handle




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

a much larger range of particle sizes. We believe the proposed framework is a further step in computer graphics to manage the details of participating media and expect further insightful exploration of this phenomenon in future research.

Acknowledgements

We would like to thank the reviewers for their valuable suggestions. This work was supported by National Natural Science Foundation of China (Grant Nos. 61972194 and 62032011).

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

References

- [1] C. Aliaga, C. Castillo, D. Gutierrez, M. A. Otaduy, J. Lopez-Moreno, and A. Jarabo. An appearance model for textile fibers. *Computer Graphics Forum*, 36(4):35–45, 2017.
- [2] J. Amanatides and A. Woo. A fast voxel traversal algorithm for ray tracing. In *In Eurographics '87*, pages 3–10, 1987.
- [3] J. Arvo. Transfer equations in global illumination. In *Global Illumination, SIGGRAPH 93 Course Notes*, 1993.
- [4] B. Bitterli and W. Jarosz. Beyond points and beams: Higher-dimensional photon samples for volumetric light transport. *SIGGRAPH*, 36(4):112:1–112:12, July 2017.
- [5] 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.*, 37(6):225:1–225:17, Dec. 2018.
- [6] C. F. Bohren and D. R. Huffman. *Absorption and Scattering of Light by Small Particles*. Wiley, 1983.
- [7] C. Bosch, X. Pueyo, S. Mérillou, and D. Ghazanfarpour. A physically-based model for

- rendering realistic scratches. *Computer Graphics Forum*, 23(3):361–370, 2004.
- [8] V. E. Cachorro and L. L. Salcedo. New improvements for mie scattering calculations. *Journal of Electromagnetic Waves and Applications*, pages 913–926, 1991.
- [9] P. Callet. Pertinent data for modelling pigmented materials in realistic rendering. *Computer Graphics Forum*, 15(2):119–127, 1996.
- [10] E. Cerezo, F. Pérez, X. Pueyo, F. J. Seron, and F. X. Sillion. A survey on participating media rendering techniques. *The Visual Computer*, 21(5):303–328, Jun 2005.
- [11] S. Chandrasekhar. *Radiative transfer*. Dover, 1960.
- [12] A. Dal Corso, J. R. Frisvad, T. K. Kjeldsen, and J. A. Bærentzen. Interactive Appearance Prediction for Cloudy Beverages. In R. Klein and H. Rushmeier, editors, *Workshop on Material Appearance Modeling*. The Eurographics Association, 2016.
- [13] 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), July 2019.
- [14] E. d’Eon. A reciprocal formulation of nonexponential radiative transfer. 1: Sketch and motivation. *Journal of Computational and Theoretical Transport*, 47(1-3):84–115, 2018.
- [15] E. d’Eon. A reciprocal formulation of nonexponential radiative transfer. 2: Monte carlo estimation and diffusion approximation, 2018.
- [16] R. A. Dobbins and T. I. Eklund. Ripple structure of the extinction coefficient. *Applied Optics*, 16:281–282, Feb. 1977.
- [17] J. Dupuy, E. Heitz, and E. d’Eon. Additional Progress Towards the Unification of Microfacet and Microflake Theories. In E. Eisemann and E. Fiume, editors, *Eurographics Symposium on Rendering - Experimental Ideas & Implementations*. The Eurographics Association, 2016.
- [18] 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, SIGGRAPH ’07*, 2007.
- [19] I. Gkioulekas, S. Zhao, K. Bala, T. Zickler, and A. Levin. Inverse volume rendering with material dictionaries. *ACM Trans. Graph.*, 32(6), 2013.
- [20] W. J. Glantschnig and S.-H. Chen. Light scattering from water droplets in the geometrical optics approximation. *Appl. Opt.*, 20(14):2499–2509, Jul 1981.
- [21] J. Guo, Y. Chen, Y. Guo, and J. Pan. A Physically-based Appearance Model for Special Effect Pigments. *Computer Graphics Forum*, 2018.
- [22] 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.*, 38(4):45:1–45:13, July 2019.
- [23] 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.
- [24] T. Hawkins, P. Einarsson, and P. Debevec. Acquisition of time-varying participating media. *ACM Trans. Graph.*, 24(3):812–815, July 2005.
- [25] 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*, 113(12):1467 – 1475, 2012.
- [26] E. Heitz, J. Dupuy, C. Crassin, and C. Dachsbacher. The sggx microflake distribution. *ACM Trans. Graph.*, 34(4):48:1–48:11, July 2015.
- [27] L. G. Henyey and J. L. Greenstein. Diffuse radiation in the Galaxy. *The Astrophysical Journal*, 93:70–83, Jan 1941.
- [28] E. A. Hovenac. Calculation of far-field scattering from nonspherical particles using a geometrical optics approach. *Appl. Opt.*, 30(33):4739–4746, Nov 1991.
- [29] D. Jackèl and B. Walter. Modeling and rendering of the atmosphere using mie-scattering. *Computer Graphics Forum*, 16(4):201–210, 1997.
- [30] W. Jakob. Mitsuba renderer, 2010. <http://www.mitsuba-renderer.org>.
- [31] 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.*, 29(4):53:1–53:13, July 2010.
- [32] W. Jakob, M. Hašan, L.-Q. Yan, J. Lawrence, R. Ramamoorthi, and S. Marschner. Discrete stochastic microfacet models. *ACM Trans. Graph.*, 33(4):115:1–115:10, July 2014.
- [33] A. Jarabo, C. Aliaga, and D. Gutierrez. A radiative transfer framework for spatially-correlated materials. *ACM Trans. Graph.*, 37(4):83:1–83:13, July 2018.
- [34] 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, Feb. 2011.
- [35] W. Jarosz, D. Nowrouzezahrai, R. Thomas, P.-P. Sloan, and M. Zwicker. Progressive photon beams. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 30(6), Dec. 2011.
- [36] W. Jarosz, M. Zwicker, and H. W. Jensen. The beam radiance estimate for volumetric photon mapping. *Computer Graphics Forum*, 27(2):557–566, Apr. 2008.
- [37] 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, July 1998.
- [38] P. Khungurn, D. Schroeder, S. Zhao, K. Bala, and S. Marschner. Matching real fabrics with micro-appearance models. *ACM Trans. Graph.*, 35(1), 2016.
- [39] 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.*, 38(6), 2019.

- [40] 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.*, 33(4):103:1–103:13, July 2014.
- [41] E. P. Lafortune and Y. D. Willems. Rendering participating media with bidirectional path tracing. In *EGWR*, pages 91–100, Vienna, June 1996.
- [42] P. Laven. Simulation of rainbows, coronas, and glories by use of mie theory. *Appl. Opt.*, 42(3):436–444, Jan 2003.
- [43] L. Lorenz. Lysbevægelser i og uden for en af plane lysbølger belyst kugle. *Det kongelig danske Videnskabernes Selskabs Skrifter*, pages 2–62, 1890.
- [44] G. Loubet and F. Neyret. A new microflake model with microscopic self-shadowing for accurate volume downsampling. *Computer Graphics Forum*, 37(2):111–121, May 2018.
- [45] 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*, 228:90–96, 2019.
- [46] 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.
- [47] S. Mérillou, J. Dischler, and D. Ghazanfarpour. Surface scratches: measuring, modeling and rendering. *The Visual Computer*, 17(1):30–45, Feb 2001.
- [48] G. Mie. Beiträge zur optik trüber medien, speziell kolloidaler metallösungen. *Annalen der Physik*, 330(3):377–445, 1908.
- [49] B. Miller, I. Georgiev, and W. Jarosz. A null-scattering path integral formulation of light transport. *ACM Trans. Graph.*, 38(4), July 2019.
- [50] 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*, EGSR'07, pages 231–242, 2007.
- [51] T. Müller, M. Papas, M. Gross, W. Jarosz, and J. Novák. Efficient rendering of heterogeneous polydisperse granular media. *ACM Trans. Graph.*, 35(6):168:1–168:14, Nov. 2016.
- [52] T. Nishita and Y. Dobashi. Modeling and rendering of various natural phenomena consisting of particles. In *Computer Graphics International 2001 (CGI'01)*, pages 149–156, 2001.
- [53] 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.
- [54] M. Pauly, T. Kollig, and A. Keller. Metropolis light transport for participating media. In *EGWR*, pages 11–22, Vienna, 2000.
- [55] B. Raymond, G. Guennebaud, and P. Barla. Multi-scale rendering of scratched materials using a structured sv-brdf model. *ACM Trans. Graph.*, 35(4), July 2016.
- [56] 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*, EGSR'04, pages 375–386, 2004.
- [57] H. Rushmeier. Input for participating media. In *ACM SIGGRAPH 1995 Courses*, SIGGRAPH '95, 1995.
- [58] 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*, 31(1):3:1–3:12, 2012.
- [59] L. Szirmay-Kalos, M. Magdics, and M. Sbert. Multiple scattering in inhomogeneous participating media using rao-blackwellization and control variates. *Computer Graphics Forum*, 37(2):63–74, 2018.
- [60] D. Toubanc. Henyey–greenstein and mie phase functions in monte carlo radiative transfer computations. *Appl. Opt.*, 35(18):3270–3274, Jun 1996.
- [61] A. Ungut, G. Grehan, and G. Gouesbet. Comparisons between geometrical optics and lorenz-mie theory. *Appl. Opt.*, 20(17):2911–2918, Sep 1981.
- [62] H. C. van de Hulst. *Light scattering by small particles*. Dover, 1981.
- [63] E. Veach. *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford, CA, USA, 1997.
- [64] E. Veach and L. J. Guibas. Optimally combining sampling techniques for monte carlo rendering. In *Proceedings of SIGGRAPH '95*, SIGGRAPH '95, pages 419–428. Association for Computing Machinery, 1995.
- [65] I. Wald, T. Ize, A. Kensler, A. Knoll, and S. G. Parker. Ray tracing animated scenes using coherent grid traversal. *ACM Trans. Graph.*, 25(3):485–493, July 2006.
- [66] B. Wang, L. Wang, and N. Holzschuch. Fast Global Illumination with Discrete Stochastic Microfacets Using a Filterable Model. *Computer Graphics Forum*, 2018.
- [67] S. Werner, Z. Velinov, W. Jakob, and M. B. Hullin. Scratch iridescence: Wave-optical rendering of diffractive surface structure. *ACM Trans. Graph.*, 36(6), Nov. 2017.
- [68] 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*, 108(1):54–64, 2007.
- [69] 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.*, 33(4):116:1–116:9, July 2014.
- [70] L.-Q. Yan, M. Hašan, S. Marschner, and R. Ramamoorthi. Position-normal distributions for efficient rendering of specular microstructure.

- ACM Trans. Graph.*, 35(4):56:1–56:9, July 2016.
- [71] L.-Q. Yan, M. Hasan, B. Walter, S. Marschner, and R. Ramamoorthi. Rendering specular microgeometry with wave optics. *ACM Trans. Graph.*, 37(4), 2018.
- [72] H. Yu, J. Shen, and Y. Wei. Geometrical optics approximation of light scattering by large air bubbles. *Particuology*, 6(5):340 – 346, 2008.
- [73] H. Yu, J. Shen, and Y. Wei. Geometrical optics approximation for light scattering by absorbing spherical particles. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 110(13):1178 – 1189, 2009.
- [74] S. Zhao, W. Jakob, S. Marschner, and K. Bala. Building volumetric appearance models of fabric using micro ct imaging. *ACM Trans. Graph.*, 30(4):44:1–44:10, July 2011.
- [75] S. Zhao, W. Jakob, S. Marschner, and K. Bala. Structure-aware synthesis for predictive woven fabric appearance. *ACM Trans. Graph.*, 31(4):75:1–75:10, July 2012.
- [76] S. Zhao, F. Luan, and K. Bala. Fitting procedural yarn models for realistic cloth rendering. *ACM Trans. Graph.*, 35(4), July 2016.
- [77] S. Zhao, L. Wu, F. Durand, and R. Ramamoorthi. Downsampling scattering parameters for rendering anisotropic media. *ACM Trans. Graph.*, 35(6):166:1–166:11, Nov. 2016.
- [78] X. Zhou, S. Li, and K. Stamnes. Geometrical-optics code for computing the optical properties of large dielectric spheres. *Appl. Opt.*, 42(21):4295–4306, Jul 2003.



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



Bingyang Hu is a graduate student in the Department of Computer Science and Technology, Nanjing University, China. 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 Ph.D. student in the Department of Computer Science and Technology, Nanjing University, where he received his bachelor's degree in 2018. His research interests include realistic rendering and 3D vision.



Yanwen Guo received a 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. He was a Visiting Professor with the Department of Computer Science and Engineering at The Chinese University of Hong Kong, in 2006 and 2009, and the Department of Computer Science, The University of Hong Kong, in 2008, 2012, and 2013. He was a Visiting Scholar at 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.



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