Efficient Rendering of Heterogeneous Polydisperse Granular Media

Thomas Müller[1,2] Marios Papas[1] Markus Gross[1,2] Wojciech Jarosz[3] Jan Novák[1]

[1]Disney Research [2]ETH Zürich [3]Dartmouth College

Figure 1: *Our method allows efficient rendering of dynamic, heterogeneous mixtures of grains, as demonstrated in the* BOWL *scene. We show two rendered frames of a simulation that was created with Houdini. The small images on the right of each frame show the individual components of our method, which add up to the full light transport. Explicit Path Tracing (EPT) simulates most of the high-frequency detail, such as glints. Proxy Path Tracing (PPT) accelerates the simulation using a precomputed aggregation of light scattering within each individual grain. Shell Tracing (ST) advances light through the medium in large steps using a precomputed aggregation of scattering events within large spherical regions. Across the entire simulation, we achieve a* 2.5–$5\times$ *speed improvement over the explicitly path traced reference.*

Abstract

We address the challenge of efficiently rendering massive assemblies of grains within a forward path-tracing framework. Previous approaches exist for accelerating high-order scattering for a limited, and static, set of granular materials, often requiring scene-dependent precomputation. We significantly expand the admissible regime of granular materials by considering *heterogeneous* and *dynamic* granular mixtures with spatially varying grain concentrations, pack rates, and sizes. Our method supports both procedurally generated grain assemblies and dynamic assemblies authored in off-the-shelf particle simulation tools. The key to our speedup lies in two complementary aggregate scattering approximations which we introduced to jointly accelerate construction of short and long light paths. For low-order scattering, we accelerate path construction using novel *grain scattering distribution functions* (GSDF) which aggregate intra-grain light transport while retaining important grain-level structure. For high-order scattering, we extend prior work on *shell transport functions* (STF) to support dynamic, heterogeneous mixtures of grains with varying sizes. We do this without a scene-dependent precomputation and show how this can also be used to accelerate light transport in arbitrary continuous heterogeneous media. Our multi-scale rendering automatically minimizes the usage of explicit path tracing to only the first grain along a light path, or can avoid it completely, when appropriate, by switching to our aggregate transport approximations. We demonstrate our technique on animated scenes containing heterogeneous mixtures of various types of grains that could not previously be rendered efficiently. We also compare to previous work on a simpler class of granular assemblies, reporting significant computation savings, often yielding higher accuracy results.

Keywords: ray tracing, global illumination, appearance modeling

Concepts: •**Computing methodologies** → **Ray tracing;** Reflectance modeling;

1 Introduction

In this paper, we address the problem of efficiently rendering massive, dynamic, and heterogeneous collections of discrete grains which can be differently sized (polydisperse). Such discrete granular materials are frequently found in natural scenes, e.g. wet and dry sand, bubbles and foam in liquids, layers of soils, snow, or a bowl with spices. These materials are characterized by fine-scale details and smooth large-scale appearance, both of which stem directly from the material properties and arrangement of individual grains. It is precisely the potential for visual complexity at either of these two extremes that makes rendering massive collections of grains a fundamentally difficult problem.

The high-frequency visual structures that are typical for grains—including geometric detail, sudden changes in color, and glints—arise from the first few interactions of light upon hitting the granular assembly. This high-frequency **low-order** transport is coherent, making it difficult to summarize with an aggregate approximation without eliminating its characteristic salient details. In granular assemblies composed of highly scattering grains, such as snow or sugar, light can scatter hundreds or even thousands of times before exiting the material. This **high-order** scattering is what leads to the characteristic smooth large-scale appearance of many granular materials. Unfortunately, such high-order scattering is

SA '16 Technical Papers,, December 05 - 08, 2016, Macao
ISBN: 978-1-4503-4514-9/16/12
DOI: http://dx.doi.org/10.1145/2980179.2982429

Table 1: *Function, algorithm, and measure acronyms in this paper.*

Acronym	Definition
TSDF	Teleporation Scattering Distribution Function
GSDF	Grain Scattering Distribution Function
STF	Shell Transport Function
EPT	Explicit Path Tracing
PPT	Proxy Path Tracing
VPT	Volumetric Path Tracing
ST	Shell Tracing
DA	Diffusion Approximation
TTUV	Time To Unit Variance
(M)RSE	(Mean) Relative Squared Error

extremely expensive to simulate using brute-force techniques like path tracing [Kajiya 1986] where each individual interaction must be simulated. Luckily, the correlation of the light paths diminishes with each grain interaction, which makes summarizing the high-order scattering with an aggregate approximation feasible.

Current state-of-the-art techniques [Moon et al. 2007; Meng et al. 2015] for rendering discrete granular materials focus their attention on efficiently aggregating high-order scattering. Moon et al. [2007] precompute *shell transport functions* (STFs; see Table 1 for a list of acronyms used throughout this paper) which summarize the spatio-angular light transport within increasingly large spherical shells of a *particular* discrete granular assembly. They use the STFs during rendering to make large jumps through the assembly, aggregating many scattering interactions into a single step. They call this procedure *shell tracing* (ST). While it significantly accelerates high-order scattering, the required precomputation is dependent on the specific arrangement of grains, making it incompatible with dynamic scenes and impractical for exploring the visual impact of different types of grains or assemblies. Meng et al. [2015] addressed this problem by converting the discrete assembly of grains to a continuous homogeneous participating medium—with transport thereof being described by the radiative transfer equation (RTE) [Chandrasekar 1960]—and simulated high-order transport using standard volumetric path tracing [Lafortune and Willems 1996] and the diffusion approximation [Stam 1995; Jensen et al. 2001; Li et al. 2005] (see Figure 2 for an illustration). Their primary contribution was a principled way to derive the appropriate RTE parameters for a dense assembly of grains, driven only by the packing density and a per-grain precomputation.

While these prior methods can substantially accelerate high-order scattering, they do not address low-order scattering, relying instead on brute-force, explicit path tracing [Kajiya 1986] to faithfully preserve the high-frequency details from low-order scattering (Figure 2(a)). These methods are in fact so successful at accelerating high-order transport that they shift the computational bottleneck to the brute-force simulation of the remaining low-order scattering component. Unfortunately, the high-order solutions are quite approximate, so they *must* be prefixed by explicit path tracing through several grains in order to retain visual accuracy. Moreover, in order to attain such impressive speedup, these prior approaches imposed severe constraints on the types of supported scenes, such as homogeneity, static scenes, procedural sphere packings, etc. These constraints make these prior approaches difficult to generalize to the more difficult case of dynamic, polydisperse, heterogeneous granular mixtures that we are interested in simulating.

In contrast to these prior approaches, we develop a method to directly tackle the coherent, low-order scattering in granular assemblies

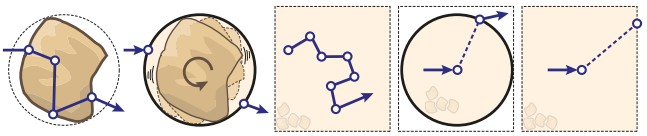

(a) *Explicit PT* **(b)** *Proxy PT* **(c)** *Volum. PT* **(d)** *Shell tracing* **(e)** *Diffusion*

Figure 2: *Overview of different techniques. Meng et al. [2015] combine* **(a)**, **(c)**, *and* **(e)**. *We introduce two new approaches* **(b, d)** *and propose to use them for the majority of low-order and high-order transport respectively, coupling them with* **(a)** *and* **(c)** *only sparingly when absolutely necessary.*

while simultaneously proposing a generalized high-order scattering solution that works with dynamic, heterogeneous and polydisperse collections of grains. Our low-order scattering approximation, *proxy path tracing* (PPT) using *GSDFs* (Figure 2(b)), is accurate enough to be used immediately on the first interaction from the camera for distant views, and typically requires handling only the very first grain using explicit path tracing for close-up views—significantly reducing the current path tracing bottleneck in low-order scattering. For high-order scattering, we transition to a continuous volumetric medium representation like Meng et al. In doing so, however, we lift their monodisperse, homogeneous, procedural sphere packing assumption and show how to derive RTE parameters for heterogeneous, polydisperse granular assemblies. Furthermore, our analysis of high-order scattering in Meng et al.'s approach revealed that its efficiency is limited by volumetric path tracing and its fidelity is limited by the severe approximations introduced by diffusion. To address these problems, we accelerate high-order scattering with a generalization of Moon et al.'s STFs (Figure 2(d)). We overcome the scene dependence of Moon et al.'s precomputation by posing our STFs within an RTE appearance space, requiring only a single precomputation to support arbitrary continuous homogeneous media. We further show how to apply our single database of precomputed STFs in heterogeneous granular mixtures while limiting visible bias. The end result is the first accelerated method that can handle heterogeneous, dynamically changing mixtures of grains like that shown in Figure 1.

2 Previous Work

Aggregating complex scattering to avoid costly on-the-fly simulations is a common practice in computer graphics. In the following, we categorize the wide spectrum of related work w.r.t. the scale of aggregation and review prior work dealing with granular materials.

Primitive-Scale Aggregation. Realistic rendering of fine assemblies often demands individual primitives—sand grains [Meng et al. 2015], hair strands [Wei et al. 2005; Jakob et al. 2009], or fibers of knitwear [Xu et al. 2001]—to be discernible in the final image. Accurately simulating light transport within each primitive, however, is generally not necessary and can be aggregated or precomputed. A portion of our method is inspired by the success of primitive-scale aggregation techniques for fibers [Marschner et al. 2003], water droplets [Sadeghi et al. 2012], human faces [Debevec et al. 2000], and BSSRDF approaches [Jensen et al. 2001; Peers et al. 2006] that propose various ways to "short-circuit" the transport internal to individual primitives. Early approaches for simulating light propagation in fibers used voxel-based approximations [Kajiya and Kay 1989; Neyret 1998], while more recent work, pioneered by Marschner et al. [2003] and Zinke and Weber [2007], aggregates chains of internal scattering events within cylindrical hair fiber geometry into a reflectance function, which is then applied to simplified curve

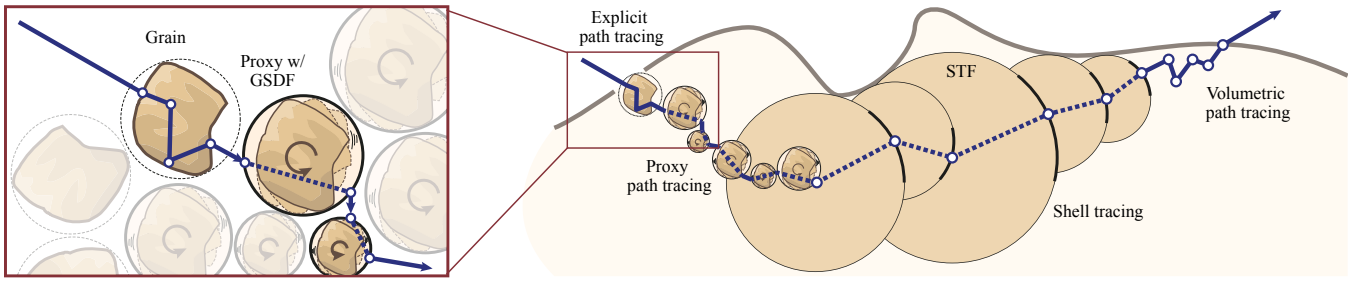

Figure 3: *We start building paths by (optionally) explicitly path tracing (EPT) to account for fine detail of grains and resolve high-frequency effects, e.g. glints. We then switch to our proxy representation (PPT) which leverages GSDFs to accelerate grain-internal transport. To quickly simulate long paths, we employ shell tracing (ST) using our extension to STFs. Whenever the smallest shell does not fit inside the boundary of the assembly, we employ volumetric path tracing (VPT).*

geometry. Since then, researchers have further investigated issues of energy conservation [d'Eon et al. 2011], non-human hair [Yan et al. 2015], and efficient importance sampling [Hery and Ramamoorthi 2012; Ou et al. 2012; d'Eon et al. 2013].

Our GSDF is conceptually similar to the aforementioned models in that it expresses the aggregate intra-primitive scattering as a reflectance function that is then applied to simplified proxy geometry. However, since the grains can be arbitrary, with appearance hard to capture in closed form, we construct the GSDF numerically, in a low-dimensional subspace, using an inexpensive precomputation step. Our approach therefore shares similarities to a tabulated reflectance field [Debevec et al. 2000] that describes the relation between incident and outgoing light fields [Levoy and Hanrahan 1996]. Directly related to our approach is the TSDF by Meng et al. [2015] that also aggregates per-grain scattering. Meng et al., however, use their TSDF to only derive parameters of a substituting continuous volume, which cannot preserve the appearance of individual grains or BRDF-like effects such as sheen, or direction-dependent shadowing and masking.

Tangentially related to ours is also the work on synthesizing and importance sampling glints [Yan et al. 2014; Jakob et al. 2014b] due to small faceted structures. These methods, however, operate on surfaces and their extension to granular assemblies is an open problem beyond the scope of this paper; we currently handle glints by explicit path tracing.

Local-Scale Aggregation. Whenever distinguishing individual primitives is not necessary, one can summarize their light interactions by statistical means. This is typical for micro-facet models derived analytically [Torrance and Sparrow 1967], measured from real materials [Matusik et al. 2003], or simulated using virtual surfaces [Westin et al. 1992; Ashikhmin et al. 2000; Kimmel and Baranoski 2007; Schröder et al. 2011]. Pharr and Hanrahan [2000] presented a formal framework for aggregating local interactions and demonstrated the benefits on subsurface scattering. Sadeghi et al. [2013] developed a phenomenological model for cylindrical assemblies such as cloth, and Jakob et al. [2014a] proposed a framework for aggregating scattering across layers. Zhao et al. [2014] studied visual equivalence of materials with different parameters and proposed to leverage similarity relations to accelerate rendering.

Most relevant to ours, however, are methods that approximate the meso-scale assembly by a continuous medium with a comparable scattering behavior. We build on the contributions of Meng et al. [2015] who combined work on densely-packed materials [Dixmier 1978; Dullien 1991; Skoge et al. 2006; Song et al. 2008] to develop a principled way of converting monodisperse, homogeneous mixtures of grains into a continuous volume expressed with RTE parameters. We derive RTE parameters for the more general case of inhomogeneous granular mixtures and further enhance our model by factoring in observations of Torquato and Lu [1993] for polydisperse random media.

Large-Scale Aggregation. One of our goals is to efficiently capture and synthesize high-order transport—possibly over large distances—that gives high-albedo granular materials their characteristic smooth appearance. Meng et al. [2015] accelerate high-order scattering by using a diffusion approximation combining multiple approaches [Li et al. 2005; d'Eon and Irving 2011; Habel et al. 2013] to quickly approximate multiple scattering from a point inside an approximating continuous volume. Unfortunately, the approximate boundary conditions of existing diffusion methods often lead to significant bias. Aggregating transport between possibly distant parts, while acknowledging the scene geometry, is common to many precomputed radiance transport techniques, which, however, typically require scene-dependent precomputation; see Ramamoorthi [2009] for an overview. Two notable exceptions are the modular approach by Loos et al. [2011] for fast indirect illumination in indoor environments, and the work of Zhao et al. [2013] for efficiently simulating volumetric transport in large assemblies of cloth fibers. In the latter, the authors precompute internal transport in a number of box exemplars to later use them to construct large steps through the assembly. This is conceptually similar to our approach, but in terms of parameterization, our aggregates are closer to the spherical STFs introduced by Moon et al. [2007].

In a precomputation, Moon et al. [2007] tabulate the amount of light due to an infinitesimal pencil beam that reaches spherical shells of increasing radii. Assuming a homogeneous granular medium, this function can be used to accelerate the render time by aggregating many complex grain interactions with a simple evaluation of a STF. This provides high-quality approximations of high-order scattering at the cost of scene-dependent precomputations. A very similar construct has been employed by Lee and O'Sullivan [2007] directly on homogeneous continuous volumes, also facilitating scene-dependent precomputations. In our work we show how to extend STFs for evaluating large-scale transport in arbitrary continuous media with low bias and without the need for any scene-dependent precomputation. Further, we leverage the continuous volume approximation of Meng et al. [2015] such that our extended STFs can be used to accelerate light transport in granular media.

3 Method Overview

We target Monte Carlo path-tracing frameworks that estimate light transport by means of averaging multiple randomly constructed

light paths starting at the camera. Similar to previous works on rendering discrete media, we aim for a multi-scale approach that trades accuracy for speed whenever possible without sacrificing visual quality. The granular assembly must therefore be represented at each scale with all necessary, but no extraneous detail required for accurately simulating its visual appearance. In contrast to Meng et al. [2015] and Moon et al. [2007], our goal is to support polydisperse mixtures, i.e. random mixtures with varying grain sizes, and general inhomogeneous assemblies that consist of multiple base materials. Figure 3 illustrates the progression of high-level algorithmic steps a single light path would undergo in our approach (from left to right), which we explain and motivate in the following.

Explicit Path Tracing. Since constructing light paths is costly, we consider the full detail only—when necessary—for the first intersected grain of the path. This preserves fine-scale structure and high frequency effects such as glints. We adopt the terminology of Meng et al. and refer to this part as *explicit path tracing* (EPT).

Proxy Path Tracing. Instead of removing grain-level detail by switching directly to a continuous volumetric approximation like Meng et al., we introduce a new low-order scattering acceleration technique that retains important grain-level structure. As soon as possible, we replace the individual grains by tightly bounding spherical proxies that summarize internal scattering using a novel *grain scattering distribution function* (GSDF). The GSDF approximately relates the incident and outgoing radiance functions on the spherical proxy bounding the grain (Section 4). We call this approach *proxy path tracing* (PPT). Since our low-order scattering approximation retains grain-level details, we can use it early on, significantly reducing the bottleneck of the EPT prefix of Meng et al., and postponing the smoothing effects of a continuous volumetric approximation.

Volumetric Path Tracing. Only after proxy path tracing a certain depth into the medium—i.e. Euclidean distance from the surface—we begin using a continuous volumetric approximation. Meng et al. [2015] derived RTE parameters for such an approximation exclusively for a rigid procedural packing of equal-sized (monodisperse) grains, which is too restrictive for our purposes. To support dynamic scenes with heterogeneous, polydisperse, and arbitrarily varying packing densities of grains, our system accepts both i) procedurally defined granular assemblies with heterogeneously textured mixing ratios, and ii) an explicit list of the grain centers, bounding radii, and grain type (e.g. a dynamic simulation). In both cases, we only require that the grain bounding spheres do not overlap. In Section 5 we describe our new method for deriving heterogeneous RTE parameters that approximate such a generic collection of polydisperse, heterogeneous grains. Once we use a continuous volumetric approximation, we can rely on volumetric path tracing (VPT) like previous work [Meng et al. 2015].

Shell Tracing. To further reduce the cost of high-order scattering, however, we employ shell tracing to summarize large amounts of scattering interactions in single steps. Leveraging STFs within a continuous medium like Lee and O'Sullivan [2007] and not a discrete one like Moon et al. [2007] allows us to circumvent their scene-dependent precomputation by generating a database of STFs over an appearance space of RTE parameters (Section 6) in a global one-time precomputation. This single database not only allows us to accelerate high-order scattering in our polydisperse granular mixtures, but it is also readily applicable to rendering *any* continuous homogeneous volume governed by the RTE. While we precompute STFs for arbitrary homogeneous media, we also devise a way to leverage them with minimal visual error for the case of heterogeneously varying RTE parameters.

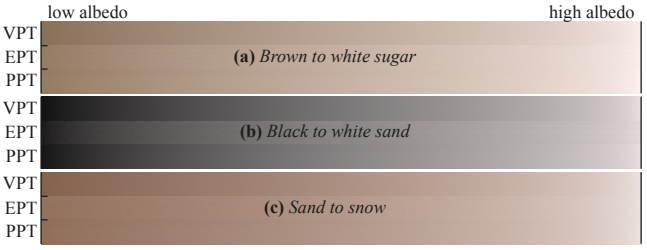

Figure 4: *Examples showing the error of a continuous approximation for three gradually changing mixtures of tiny (sub-pixel) grains. Compared to the reference (middle) the VPT solution (top) is noticeably darker than in the low-albedo region (left). Our PPT solution (bottom) approximates the reference more faithfully.*

With the GSDF and STF aggregations available, we show how all the pieces of our approach fit together (Section 7) to ensure high visual fidelity at low computational cost. We present implementation details (Section 8), rendered results (Section 9), and conclude by discussing limitations and remaining challenges (Section 10).

4 Intra-Grain Scattering Aggregation

Previous approaches [Moon et al. 2007; Meng et al. 2015] handle low-order scattering via expensive explicit path tracing (EPT) that quickly becomes the computational bottleneck. A trivial solution would be to approximate the discrete medium by a continuous one, however, in addition to losing high-frequency detail, a continuous approximation of low-albedo grains suffers from incorrect estimation of brightness; see Figure 4. Our goal is to retain the visual fidelity of individual grains necessary to preserve the aforementioned effects.

We propose to accelerate low-order transport by approximating the grains using tight spherical proxies with largely identical internal transport; up to the accuracy of the grain scattering distribution function (GSDF), denoted f_g, which models the grain's appearance. The GSDF is effectively a BSSRDF [Nicodemus et al. 1992] that relates the incident radiance to differential outgoing radiance on the surface of the proxy sphere:[1]

$$f_g(\mathbf{x}_i, \vec{\omega}_i, \mathbf{x}_o, \vec{\omega}_o) = \frac{\mathrm{d}L_o(\mathbf{x}_o, \vec{\omega}_o)}{L_i(\mathbf{x}_i, \vec{\omega}_i)\mathrm{d}\mathbf{x}_i\mathrm{d}\vec{\omega}_i}. \tag{1}$$

We define the GSDF as the sum of two terms, one for uncollided-flux (i.e. light that hits the grain's bounding sphere, but misses the grain itself—a delta component; superscript 0) and one for scattered transport (superscript $+$):

$$\begin{aligned} f_g(\mathbf{x}_i, \vec{\omega}_i, \mathbf{x}_o, \vec{\omega}_o) = {} & \alpha_g^0(\mathbf{x}_o, \vec{\omega}_o)\, \delta(\vec{\omega}_i + \vec{\omega}_o) \\ & + \alpha_g^+(\mathbf{x}_o, \vec{\omega}_o)\, p_g(\mathbf{x}_i, \vec{\omega}_i | \mathbf{x}_o, \vec{\omega}_o). \end{aligned} \tag{2}$$

The probability density function (PDF) $p_g(\mathbf{x}_i, \vec{\omega}_i | \mathbf{x}_o, \vec{\omega}_o)$ describes the distribution of spatio-angular locations of incident radiance w.r.t. a given location of the outgoing radiance. In other words, p_g represents the shape of the incident light field contributing to a given coordinate $(\mathbf{x}_o, \vec{\omega}_o)$ in the outgoing light field.

4.1 Dimensionality Reduction

In its monochromatic form, p_g is eight-dimensional making tabulation of the GSDF for arbitrary grains impractical. To this end, we

[1] We use subscripts i and o for *incident* and *outgoing*, respectively, in the sense of light propagation. By convention $\vec{\omega}_o$ points in the direction of light propagation whereas $\vec{\omega}_i$ points in the opposite direction.

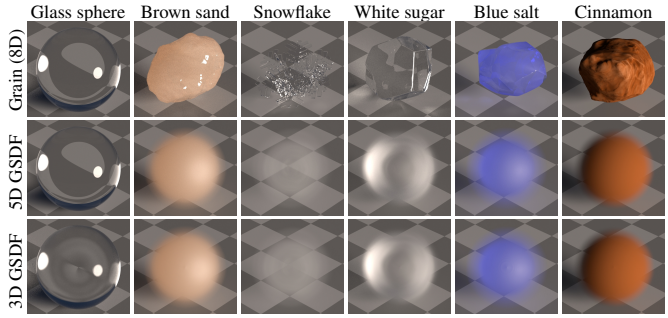

Figure 5: *Appearance of various grains (top row) and their corresponding GSDFs (bottom row). The middle row shows the impact of the random-orientation assumption without assuming separability of the directional and spatial components of the GSDF.*

introduce a number of assumptions that reduce the 8D space to two independent 3D subspaces. To ensure that the reduced-dimensional GSDF does not impair visual quality, we use the proxies only for grains viewed from a distance or "through" other grains; we elaborate on the criterion for using PPT in Section 4.3 and 7.1.

Random-Orientation Assumption. First, we assume that all light paths interacting with a proxy experience the grain each under a different random orientation. The blurry visual effect of this is best understood by looking at the middle row in Figure 5. It is important to note that all orientation-dependent effects can no longer be resolved under this assumption.

Mathematically, this approximation averages the grain appearance over all sphere surface points $\mathbf{x}_o$ and rotations about their respective normal $\vec{n}_o$. Thus, the GSDF no longer depends on $\mathbf{x}_o$ and the azimuthal component γ_o of $\vec{\omega}_o$; see Figure 6 for an illustration. In terms of the *outgoing* light field, the GSDF varies only in the inclination angle $\beta_o = \cos^{-1}(\vec{\omega}_o \cdot \vec{n}_o)$.

$$f_g(\mathbf{x}_i, \vec{\omega}_i, \mathbf{x}_o, \vec{\omega}_o) \approx \alpha_g^0(\beta_o)\,\delta(\vec{\omega}_i + \vec{\omega}_o) + \alpha_g^+(\beta_o)\,p_g(\mathbf{x}_i, \vec{\omega}_i|\beta_o). \tag{3}$$

The distribution of scattered radiance $p_g(\mathbf{x}_i, \vec{\omega}_i|\beta_o)$ is now reduced to five dimensions, which still requires excessive storage, especially if the scene contains multiple different grain types. To this end, we further reduce the dimensionality by decoupling the spatial and angular dimensions of the *incident* light field and tabulating them independently.

Separability Assumption. In order to decouple the spatial and angular components, we assume that the distribution of scattered radiance is separable, i.e.:

$$p_g(\mathbf{x}_i, \vec{\omega}_i|\beta_o) \approx p_g^{\mathbf{x}}(\mathbf{x}_i|\beta_o)\,p_g^{\vec{\omega}}(\vec{\omega}_i|\beta_o). \tag{4}$$

Introducing the separability is motivated by the fact that grains are relatively small and most incident illumination can often be assumed as purely directional, i.e. $L_i(\mathbf{x}_i, \vec{\omega}_i) = L_i^{\vec{\omega}}(\vec{\omega}_i)$. Under directional illumination, the reflection equation simplifies, allowing

Table 2: *Five tabulated dimensions of the GSDF with the respective range, discretization spacing, and resolution.*

Dim.	Domain	Spacing	Resolution
β_o	$\in [0, \pi/2]$	linear in $\cos\beta_o$	50 bins
β_i	$\in [0, \pi]$	linear in $\cos\beta_i$	360 bins
γ_i	$\in [0, 2\pi]$	linear in γ_i	180 bins
θ_i	$\in [0, \pi]$	linear in $\cos\theta_i$	180 bins
ϕ_i	$\in [0, 2\pi]$	linear in ϕ_i	180 bins

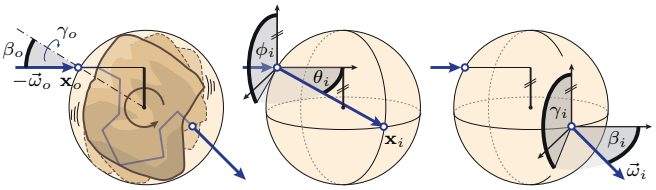

Figure 6: *Our reduced-dimensional GSDF preserves the inclination angle β_o of light paths (the remaining dimensions of the outgoing light field are neglected), and decomposes the spatial $\mathbf{x}_i$ and directional $\vec{\omega}_i$ components of the incident light field.*

pre-integration over the spatial domain:

$$\begin{aligned} L_o(\mathbf{x}_o, \vec{\omega}_o) &= \int_{S_{\vec{\omega}}} \int_{S_{\mathbf{x}}} f_g(\mathbf{x}_i, \vec{\omega}_i, \mathbf{x}_o, \vec{\omega}_o)\, L_i(\mathbf{x}_i, \vec{\omega}_i)\, \mathrm{d}\mathbf{x}_i\, \mathrm{d}\vec{\omega}_i \\ &= \int_{S_{\vec{\omega}}} L_i^{\vec{\omega}}(\vec{\omega}_i) \int_{S_{\mathbf{x}}} f_g(\mathbf{x}_i, \vec{\omega}_i, \mathbf{x}_o, \vec{\omega}_o)\, \mathrm{d}\mathbf{x}_i\, \mathrm{d}\vec{\omega}_i \\ &= \int_{S_{\vec{\omega}}} L_i^{\vec{\omega}}(\vec{\omega}_i)\, \hat{f}_g(\vec{\omega}_i, \mathbf{x}_o, \vec{\omega}_o)\, \mathrm{d}\vec{\omega}_i, \end{aligned} \tag{5}$$

where the distribution of scattered radiance from Equation (3) can be represented, without changing the reflection integral, by a product of the directional component $p_g^{\vec{\omega}}(\vec{\omega}_i|\beta_o)$ and a constant. Evaluating the reflection equation thus requires only a 3D precomputation.

In reality, the incident illumination is not perfectly directional, e.g. light arriving from nearby grains. We thus do not assume perfectly directional lighting, but only employ the separability assumption, i.e. we replace the aforementioned constant by a 3D spatial distribution $p_g^{\mathbf{x}}(\mathbf{x}_i|\beta_o)$, as shown in Equation (4). This preserves some of the dependence of the GSDF on the spatial component of the incident light field without increasing the dimensionality of the tabulation.

Figure 5 shows the type of blurring that the random-orientation and separability assumptions induce. While the first is responsible for the loss of geometric detail, the second assumption may result in further blurring of details seen through the grain (see the bottom image of the glass sphere in Figure 5). Despite the introduced approximations, the reduced-dimensional GSDF still preserves most of the important directional effects, such as local brightening and darkening due to reflection and refraction that are important to preserve characteristic visual features, such as sheen.

4.2 Precomputation of the GSDF

When parsing the scene, we take each grain type and precompute a 3D GSDF. We discretize its individual components according to Table 2 and populate it in the following way. We shoot a light path from a randomly chosen direction towards the grain, choosing the origin of the path such that the grain's bounding sphere's cross section is hit with a uniform probability density. After resolving all interactions with the grain, we record the exit location of the

path in two 3D RGB histograms: a spatial one parameterized by $(\beta_o, \mathbf{x}_i)$ and a directional one parameterized by $(\beta_o, \vec{\omega}_i)$. After we trace enough paths—each effectively experiencing the grain under a different random orientation—we convert the histograms into a set of piece-wise constant conditional PDFs $p_g^{\mathbf{x}}(\mathbf{x}_i|\beta_o)$ and $p_g^{\vec{\omega}}(\vec{\omega}_i|\beta_o)$ by normalizing 2D β_o-slices of the respective histograms, facilitating trivial importance sampling of the GSDF. We also keep track of the relative fractions of energy that exited the proxy after zero and after at least one scattering event to represent $\alpha_g^0(\beta_o)$ and $\alpha_g^+(\beta_o)$, respectively. We describe the use of the precomputed GSDF during path tracing in Section 7.

On our workstation running two Intel Xeon E5-2680v3 CPUs (24 cores in total) a GSDF's precomputation takes between 2 min and 10 min depending on how many scattering events occur within the given grain. With our discretization defined in Table 2 stored as an array of 4 byte floating-point numbers, a GSDF occupies roughly 60 MB of memory in uncompressed form. We did not attempt to reduce the GSDF's memory footprint, and we chose its resolution conservatively.

4.3 Directional Approximation Error of the GSDF

During rendering, we need to decide in which situations the GSDF provides a sufficiently good approximation, and when to employ explicit path tracing, e.g. to resolve glints. To address this, we estimate the GSDF's approximation error in the directional domain (i.e. when averaging grain appearances over all incident and outgoing path locations $\mathbf{x}_i$ and $\mathbf{x}_o$) during its precomputation as:

$$e_{f_g}(\vec{\omega}_i, \vec{\omega}_o) = |f_g(\vec{\omega}_i, \vec{\omega}_o) - f_g(\beta_i)|. \tag{6}$$

For a given pair of incident and outgoing directions, the error amounts to the absolute difference between the ground-truth GSDF averaged over $\mathbf{x}_i$ and $\mathbf{x}_o$, $f_g(\vec{\omega}_i, \vec{\omega}_o)$, and our GSDF with the random-orientation assumption and also averaged spatially, $f_g(\beta_i)$. Note, that spatial averaging collapses our approximate GSDF into a rotationally symmetric phase function, depending only on a single one-dimensional parameter $\beta_i = \cos^{-1}(-\vec{\omega}_i \cdot \vec{\omega}_o)$. We provide visualizations and statistics of e_{f_g} in the supplementary material. To decide when to switch to proxy path tracing, we use the maximum approximation error

$$E_{f_g} = \max_{\vec{\omega}_i, \vec{\omega}_o} e_{f_g}(\vec{\omega}_i, \vec{\omega}_o), \tag{7}$$

which directly corresponds to the worst-case appearance difference between a grain and its 3D GSDF; we discuss details in Section 7.1.

4.4 Analysis and Discussion

One could theoretically reduce the 8D GSDF to an arbitrary low-dimensional subset. Our choice of dimensions to keep and decouple is motivated by the intuitive physical interpretation, and by the relative importance of individual visual features. If, for instance, we flipped the dimensions neglecting the incident dimensions and decoupling the outgoing ones, we would preserve the geometric detail of grains, but they would appear lit by a constant (averaged) environment map. For the applications we target, maintaining directional lighting effects is more important than preserving fine-scale geometric detail (these can be preserved using EPT), and we chose the low-dimensional subset accordingly. An additional benefit of choosing the subset empirically, in contrast to subspaces obtained by SVD or PCA, is that the visual impact can be well understood allowing to easily design switching criteria for using the GSDF.

It is worth noting that the 3D GSDFs are not reciprocal. Reciprocity could be enforced by treating the corresponding incident and outgoing dimensions equally. For instance, one could neglect both incident and outgoing spatial dimensions and preserve all directional ones; this corresponds to reducing the appearance of the grain to a disk of constant color. We found it to be visually beneficial to treat the incident and outgoing dimensions independently. This provides higher quality where needed in our applications, but arguably complicates integration of our GSDFs into bidirectional algorithms.

Several other papers propose to aggregate local scattering. Meng et al. [2015] use a 2D teleportation scattering distribution function (TSDF) that aggregates intra-grain scattering. In contrast to our GSDF, which relates light fields around the grain, they use the TSDF during volumetric path tracing to "only" statistically account for the effect of non-point scatterers; the actual positions of grains are thus ignored. The TSDF is defined over a cylindrical volume that is larger than the bounding sphere of the proxy. The teleportation can thus in practice lead to light leaks due to crossing of the volume boundary and possible penetration of another object; these need to be checked to avoid artifacts. Our GSDF does not suffer from such issues as we "teleport" between points on the proxy of the grain.

5 Continuous-Volume Approximation

Following, we describe our methodology for converting a heterogeneous granular medium to an equivalent continuous volume. We extend the work of Meng et al. [2015], which derives RTE (σ_t, α, f_p) parameters for homogeneous, monodisperse mixtures, to assemblies that are polydisperse, i.e. mixtures with differently sized grains, where all grains, except for a scaling factor, are identical. We assume that the density of the grain's interior medium scales with the size of the grain. This simplifies the derivation, allows for an artist-friendly, scale-independent modeling in normalized object space, and enables using a single GSDF for all sizes of the grain. Finally, we demonstrate how to handle heterogeneous mixtures by discretizing the RTE parameters into a voxel grid.

The conversion to RTE parameters is necessary for efficiently mapping the appearance of arbitrary granular assemblies into a single unified *appearance space* of continuous volumes. This appearance space is the cornerstone of our large-scale aggregation procedure, which is detailed in Section 6.

5.1 Homogeneous polydisperse Mixtures

Phase Function. We obtain the phase function of the continuous approximation analogously to Meng et al.: we observe that scattering distributions due to randomly oriented grains can be approximated well by one-dimensional functions of $(\vec{\omega}_i \cdot \vec{\omega}_o)$. We thus extract f_p from the directional component of the GSDF $p_g^{\vec{\omega}}(\vec{\omega}_i|\beta_o)$ by averaging over the azimuthal dimension of $\vec{\omega}_i$. Since p_g is normalized, the 1D average is already a valid PDF.

Albedo. We compute the albedo of the continuous approximation α_v from the albedo of the scattered component of the GSDF by integrating area-weighted $\alpha_g^+(\beta_o)$ over $\langle 0, \pi/2 \rangle$.

Extinction Coefficient. The extinction coefficient σ_t of the medium is the reciprocal of the mean free path λ_t of the granular mixture:

$$\lambda_t = \lambda_c + \alpha_v \lambda_v, \tag{8}$$

where λ_c corresponds to the mean outer chord length[2] of the original grain mixture. It can be computed with the following closed-form

[2] In granular media literature, *outer chord length* describes the length of an inter-grain path segment; its analog in continuous media is *free path*.

solution to an infinite series [Meng et al. 2015]:

$$\lambda_c = (\lambda_s + \lambda_\delta)\frac{1-c}{c} + \lambda_s, \tag{9}$$

where c is the conditional probability that a random ray passing through a grain's bounding sphere will hit the grain and λ_s is the mean outer chord length obtained by replacing each grain with an opaque version of its bounding sphere. Meng et al. compute λ_s with an approximation proposed by Dixmier [1978] for monodisperse mixtures:

$$\lambda_s^m = \frac{4R}{3}\frac{1-f}{f}, \tag{10}$$

where $4R/3$ corresponds to the volume-to-surface ratio for a sphere with radius R, and f is the packing rate of the bounding sphere assembly. We instead use a generalized version of this approximation proposed by Torquato and Lu [1993] for polydisperse mixtures:

$$\lambda_s^p = \frac{4\langle R^3\rangle}{3\langle R^2\rangle}\frac{1-f}{f}, \tag{11}$$

where $\langle . \rangle$ denotes the expected value along a ray passing through the corresponding *infinite* polydisperse mixture. The ratio of expected values in Equation (11) can be interpreted as the radius of the "average" sphere that the ray experiences.

The mean spatial offset of an unscattered ray λ_δ and the length of the mean teleportation distance λ_v are precomputed by shooting many rays against the grain; please refer to [Meng et al. 2015] for details. Both lengths linearly depend on the grain. For polydisperse mixtures, we compute them on grains scaled to have unit-radius bounding spheres and then scale them by $\langle R^3\rangle/\langle R^2\rangle$, i.e. the radius of the "average" sphere.

5.2 Heterogeneous Mixtures

In order to support varying pack rates and grain concentrations, we allow the RTE parameters to change spatially and discretize them into a voxel grid. Our goal is to derive a heterogeneous continuous approximation without making any assumptions about how the discrete assembly was created, thereby supporting procedural, dynamically simulated, or fully explicit modeling approaches.

To construct the heterogeneous medium, we first create a voxel grid that will hold the RTE parameters of the continuous approximation; we use $2\times$ the diameter of the largest grain as the voxel size in all our examples. Since the voxels are relatively small, we can safely assume the assembly to be homogeneous within each voxel; this allows deriving the local RTE parameters using recipes described in Section 5.1. To collect reliable statistics of grains in voxel v, we enumerate all grains in a small neighborhood of $3 \times 3 \times 3$ voxels centered around v and compute: the mean GSDF, the mean squared and cubed radii $\langle R^2\rangle$, $\langle R^3\rangle$ of the bounding spheres, respectively, and their packing rate, i.e. the fraction of the voxel volume they occupy. These provide necessary information to derive the (homogeneous) RTE parameters for the voxel. Since the grains' types are not correlated with their sizes in our scenes, size-agnostic averaging of the GSDFs is optimal. However, cross-section weighted averaging is more appropriate in the general case in order to capture the correlation between a grain's size and its contribution to the appearance of the volume. The voxel-grid representation is sufficient to perform volumetric path tracing, which we accelerate using shell transport functions described next.

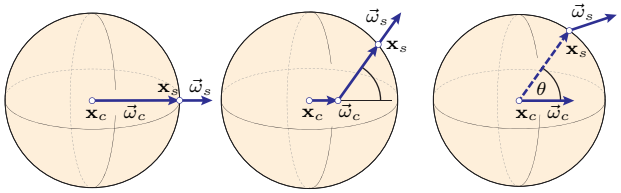

(a) *Uncollided flux* **(b)** *Single scattering* **(c)** *Multiple scattering*

Figure 7: *The parametrization of the shell transport function. We split the function into non-scattering paths* **(a)***, single scattering paths* **(b)** *and multiple scattering paths* **(c)***. The dashed arrow from* $\mathbf{x}_c$ *to* $\mathbf{x}_s$ *summarizes all scattering happening within the shell.*

6 Large-Scale Aggregation

Our goal in this section is to aggregate many grain interactions in one step to accelerate the evaluation of high-order transport in granular media with high albedo. We build on the idea of approximating the granular assembly by a continuous medium [Meng et al. 2015], but instead of using the error-prone diffusion approximation, we accelerate construction of paths using STFs [Moon et al. 2007]. We first briefly review the approach by Moon et al. and then present our modified STFs that operate on a continuous approximation of heterogeneous assemblies.

6.1 Shell Transport Functions

A shell transport function [Moon et al. 2007] aggregates the transport in a spherical region containing many grains. In contrast to the GSDF, which relates incident and outgoing light fields, the STF captures only the outgoing light field due to an infinitesimal pencil beam emitter at the center of the shell. In the monochromatic case, the STF is ten-dimensional parameterized by the location $\mathbf{x}_c$ and direction $\vec{\omega}_c$ of the emitter, and by the position $\mathbf{x}_s$ and direction $\vec{\omega}_s$ on the surface of the shell (see Figure 7 for an illustration). Given a photon originating from this pencil beam emitter, the STF quantifies the probability of the photon first exiting the shell at $\mathbf{x}_s$ in direction $\vec{\omega}_s$:

$$f_s(\mathbf{x}_c, \vec{\omega}_c, \mathbf{x}_s, \vec{\omega}_s) = \alpha_s p_s(\mathbf{x}_s, \vec{\omega}_s | \mathbf{x}_c, \vec{\omega}_c), \tag{12}$$

where α_s and $p_s(\mathbf{x}_s, \vec{\omega}_s | \mathbf{x}_c, \vec{\omega}_c)$ are the fraction and spatio-angular distribution of light that reached the surface of the shell, respectively. We will refer to α_s as the shell albedo.

By assuming a homogeneous assembly of grains with optical properties invariant with respect to position or direction, Moon et al. make the STF agnostic to the position and orientation of the emitter, and rotationally symmetric about $\vec{\omega}_c$. The function thus reduces to four dimensions and can be conveniently parameterized by the radius $r = \|\mathbf{x}_s - \mathbf{x}_c\|$, elevation angle of the surface point $\theta = \cos^{-1}(\vec{\omega}_c \cdot \vec{n}_{\mathbf{x}_s})$, and outgoing direction $\vec{\omega}_s$; see Figure 7 for an illustration. To precompute the STF, Moon et al. [2007] trace a large number of randomly seeded photons through an infinite homogeneous granular assembly. When the photon first reaches distance r from the origin, they record its flux, location, and direction to later reconstruct the numerical STF using density estimation.

6.2 Appearance Parameterization

In order to efficiently support mixtures with spatially varying appearance, we strive for a general function, tabulated only *once*, that can be used for arbitrary sets of grains, including those not known at the time of precomputation. To that end, we use the RTE parameters (σ_t, α, f_p) of the continuous-volume approximation (Section 5) as

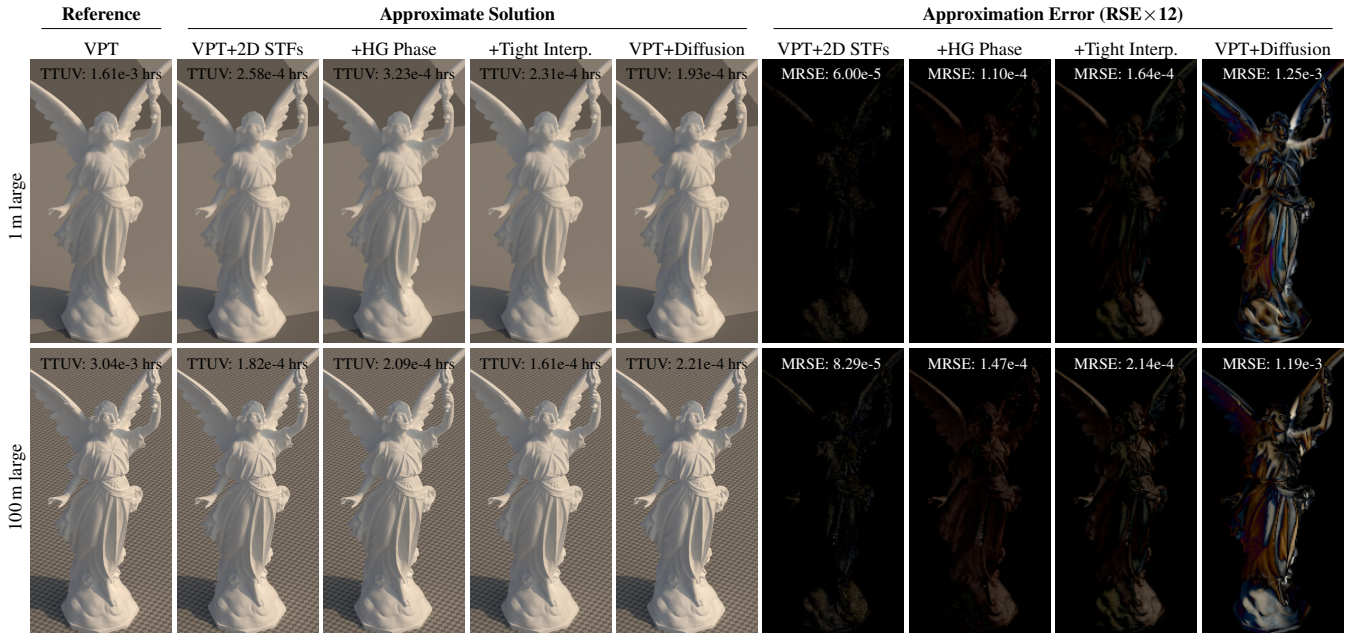

Figure 8: *We demonstrate the error introduced by each of the approximations we apply to STFs in the* LUCY *scene, which contains the Lucy statue filled with a homogeneous continuous participating medium derived from densely packed snow grains. The first column shows volumetrically path traced (VPT) reference images, which can be matched perfectly by the 4D STF in the case of homogeneous continuous media. The following three columns demonstrate the visual impact of our approximations: 2D STFs approximate $p_s^{m,\vec{\omega}}$ as a cosine-weighted hemispherical distribution, using HG phase function significantly reduces the dimensionality of continuous-volume-appearance space to permit dense tabulation, and interpolation between tabulated STFs allows fitting STF tightly into the mesh. The last column shows the diffusion approximation employed by Meng et al. [2015], which yields higher error than all our approximations combined.*

the basis for classifying any mixture of grains. This enables the abstraction of the actual grain geometry and allows us to focus only on parameters relevant to high-order transport.

Since the STF, as proposed by Moon et al., is already 4D, adding the RTE parameters naively as extra dimensions would make precomputation and storage intractable. Fortunately, the extinction coefficient of the medium can be implicitly accounted for by expressing the shell's radius r in units of mean free paths $1/\sigma_t$. Furthermore, we can reduce the phase function f_p to a single parameter, the mean cosine g of scattered light in the Henyey-Greenstein phase function [Henyey and Greenstein 1941]. The appearance of any homogeneous grain assembly can thus be classified using (α, g). Adding these to the STF—six dimensions in total—is unfortunately still impractical to tabulate. To this end, we study the directional distribution of flux exiting the shell and propose a decomposition that allows reducing the tabulation to only four dimensions.

6.3 Separation of Uncollided and Once-Scattered Flux

We define our STF database as a sum of uncollided, single-scattered, and multi-scattered components:

$$\begin{aligned} f_s(\mathbf{x}_c, \vec{\omega}_c, \mathbf{x}_s, \vec{\omega}_s, \alpha, g) = {} & \alpha_s^0(r)\,\delta(\vec{\omega}_c - \vec{\omega}_s) \\ & + \alpha_s^1(\alpha, g, r)\,p_s^1(\mathbf{x}_s, \vec{\omega}_s | g, r) \\ & + \alpha_s^m(\alpha, g, r)\,p_s^m(\mathbf{x}_s, \vec{\omega}_s | \alpha, g, r), \end{aligned} \quad (13)$$

where the superscripts 0, 1, m mark quantities of the respective components, $p_s^1(\mathbf{x}_s, \vec{\omega}_s, r)$ is the distribution of single-scattered photons. The joint PDF $p_s^m(\mathbf{x}_s, \vec{\omega}_s | \alpha, g, r)$ is defined as the product of positional and directional components:

$$p_s^m(\mathbf{x}_s, \vec{\omega}_s | \alpha, g, r) = p_s^{m,\mathbf{x}}(\theta | \alpha, g, r)\; p_s^{m,\vec{\omega}}(\vec{\omega}_s | \alpha, g, r, \theta). \quad (14)$$

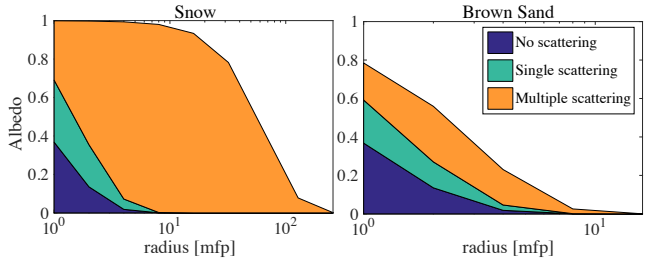

Figure 9: *Normalized contributions of individual components (uncollided, single-scatterd, and multi-scattered flux) to the total albedo vary with the radius of the shell (measured in mean free paths).*

Separating out the uncollided and single-scattered flux, and integrating them on the fly using Monte Carlo estimators, results in a significantly smoother function to be tabulated. We also found that the multi-scattering lobe $p_s^{m,\vec{\omega}}$ can be approximated well by a cosine-weighted hemispherical distribution oriented about normal $\vec{n}_{\mathbf{x}_s}$; Figure 8 and the supplementary material demonstrate the approximation error. This reduces the tabulation of p_s^m to acceptable four dimensions (α, g, r, θ). Figure 9 provides a component breakdown of the shell albedo as a function of distance from the shell center for highly scattering and absorbing grains.

6.4 Precomputation of the STF Database

To construct the STF database p_s^m, we discretize its individual components according to Table 3 and populate it in the following way. For each value of the mean cosine g_j, we create an infinite homo-

Table 3: *STF precomputation discretization spacing, range, and resolution of the four tabulated dimensions. We only precompute a given radius r if the corresponding shell albedo (α_s) is larger than 0.5. Furthermore, we add an additional bin for $\alpha = 1$.*

Dim.	Domain	Spacing	Resolution
g	$\in [-1, 1]$	linear in g	200 bins
α	$\in [0, 0.999]$	linear in $log(\frac{1}{1-\alpha})$	1000 bins
r	$\in [1, 256]$ ($\alpha_s > 0.5$)	linear in $log(r)$	≤ 9 bins
θ	$\in [0, \pi]$	linear in $\cos\theta$	45 bins

geneous volume with the corresponding HG phase function and $\alpha = 1$; our goal is to precompute for all values of α at once. We trace N photons and first tabulate the shell with the smallest radius. The radius for the k-th shell is computed as $r_k = 2^{k-1}$ mean free paths; we start with $k = 1$. For each spatial bin θ_l, we construct a histogram H_l, where the b-th bin counts photons exiting the shell after exactly b bounces. This allows computing the multi-scattered flux for any value of the discretized albedo α_i as:

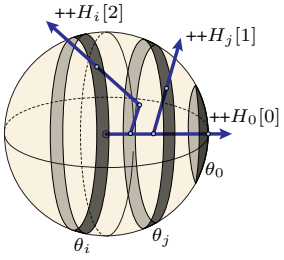

$$p_s^{m,\mathbf{x}}(\theta_k|\alpha_i, g_j, r_k) = \frac{1}{N}\sum_{b=2}^{|H_l|}\alpha_i^b H_l[b] \tag{15}$$

We also use the histogram to compute the albedos of the individual components as:

$$\alpha_s^0(r_k) = \exp(-r_k), \tag{16}$$

$$\alpha_s^1(\alpha_i, g_j, r_k) = \frac{1}{N}\sum_l \alpha_i H_l[1], \tag{17}$$

$$\alpha_s^m(\alpha_i, g_j, r_k) = \frac{1}{N}\sum_l\sum_{b=2}^{|H_l|}\alpha_i^b H_l[b]. \tag{18}$$

Equations (15–18) express all tabulated terms for a shell of radius r_k mean free paths. We use the smaller shells to accelerate tracing of paths for larger shells; this significantly reduces the number of traced path segments, cutting down the precomputation of the entire STF database to around 10 hours on our workstation running two Intel Xeon E5-2680v3 CPUs (24 cores in total). With our discretization defined in Table 3 stored as an array of 4 byte floating-point numbers, the entire STF database occupies roughly 400 MB of memory in uncompressed form. We did not attempt to reduce the database's memory footprint, and we chose its resolution conservatively.

6.5 Discussion

Figure 8 shows the Lucy statue filled with a highly scattering continuous volume that was derived from a tight homogeneous packing of sugar grains with a radius of 1 mm. The renderings use various instances of our STFs to compute high-order transport and demonstrate the error of the multiple approximations we introduce. In order to evaluate our method on varying medium densities, we show the statue at two different sizes—1 m and 100 m large. We compare the VPT ground truth to:[3]

- STFs with a cosine hemisphere as their directional component,
- STFs with the additional assumption that the volume has a HG phase function,
- STFs additionally interpolated logarithmically in size such that they tightly fit into the volume boundary,
- the diffusion approximation used by Meng et al. [2015].

The results suggest that the cosine hemisphere approximation introduces minimum error compared to diffusion while maintaining a similar rendering speed. The next step—introducing the HG phase assumption—again adds little error, but permits the tabulation of STFs over a dense appearance space, which can be re-used for all possible sets of RTE parameters. However, the lookups of the correct STF from this database during path tracing end up increasing render times slightly. With our last approximation—interpolation of STFs in order to achieve tight fits in the bounding mesh—we introduce further error, but gain speed again.

7 Multi-Scale Rendering

Our rendering algorithm is based on a standard forward path-tracing approach. Rays originating from the camera traverse through the scene and query for the nearest grain bounding sphere intersection. In the following we describe how to switch between explicit, proxy, shell, and volumetric path tracing and leverage the previously introduced GSDFs and STFs to accelerate light-path construction.

7.1 Rendering Low-Order using EPT and PPT

Choice between EPT and PPT. When a ray intersects the first grain bounding sphere, we decide whether to instantiate the full geometry (EPT) or its corresponding proxy (PPT). This decision is governed by a criterion that depends on how close the grain is to the camera and how good of an approximation our GSDF is for that particular grain's type:

$$\text{instantiate}\begin{cases}\text{explicit grain} & \text{if } \Omega_g > 4\Omega_{px} \text{ or } E_{f_g} > 0.1,\\ \text{proxy} & \text{otherwise,}\end{cases}$$

where the error of the GSDF E_{f_g} is computed using the method explained in Section 4.3, Ω_g is the solid angle occupied by the grain's bounding sphere when viewed from the camera, and Ω_{px} is approximately the solid angle of the corresponding pixel. This allows our method to switch to PPT immediately for distant grains that are well approximated by their proxies. Furthermore, this criterion is only used for the first grain interaction of each light path. After interacting with a grain or its proxy once, further interactions will no longer instantiate explicit grains.[4]

Proxy Path Tracing. Upon intersecting a grain's proxy, we first probabilistically decide whether to sample uncollided or scattered transport proportional to α_g^0 and α_g^+, respectively. In the first case, the path continues through the proxy in a straight line. Otherwise, we importance sample $p_g^{\mathbf{x}}(\mathbf{x}_i|\beta_o)$ and $p_g^{\vec{\omega}}(\vec{\omega}_i|\beta_o)$ to determine the location and direction for continuing the path. We additionally generate a shadow connection by importance sampling light sources from the exit location and combine the light-source and the directional GSDF sampling via MIS [Veach and Guibas 1995]. We provide pseudocode of this process in the supplementary material.

[3] The full 4D STFs as described by Moon et al. [2007]—if applied to homogeneous continuous media—are an exact solution up to a quantization error, and are thus not listed in Figure 8.

[4] Certain assemblies, e.g. showflakes, may require handling also the second grain using EPT to properly capture glints seen through reflection or refraction. We, however, always use proxies after the first grain for the simplicity of implementation.

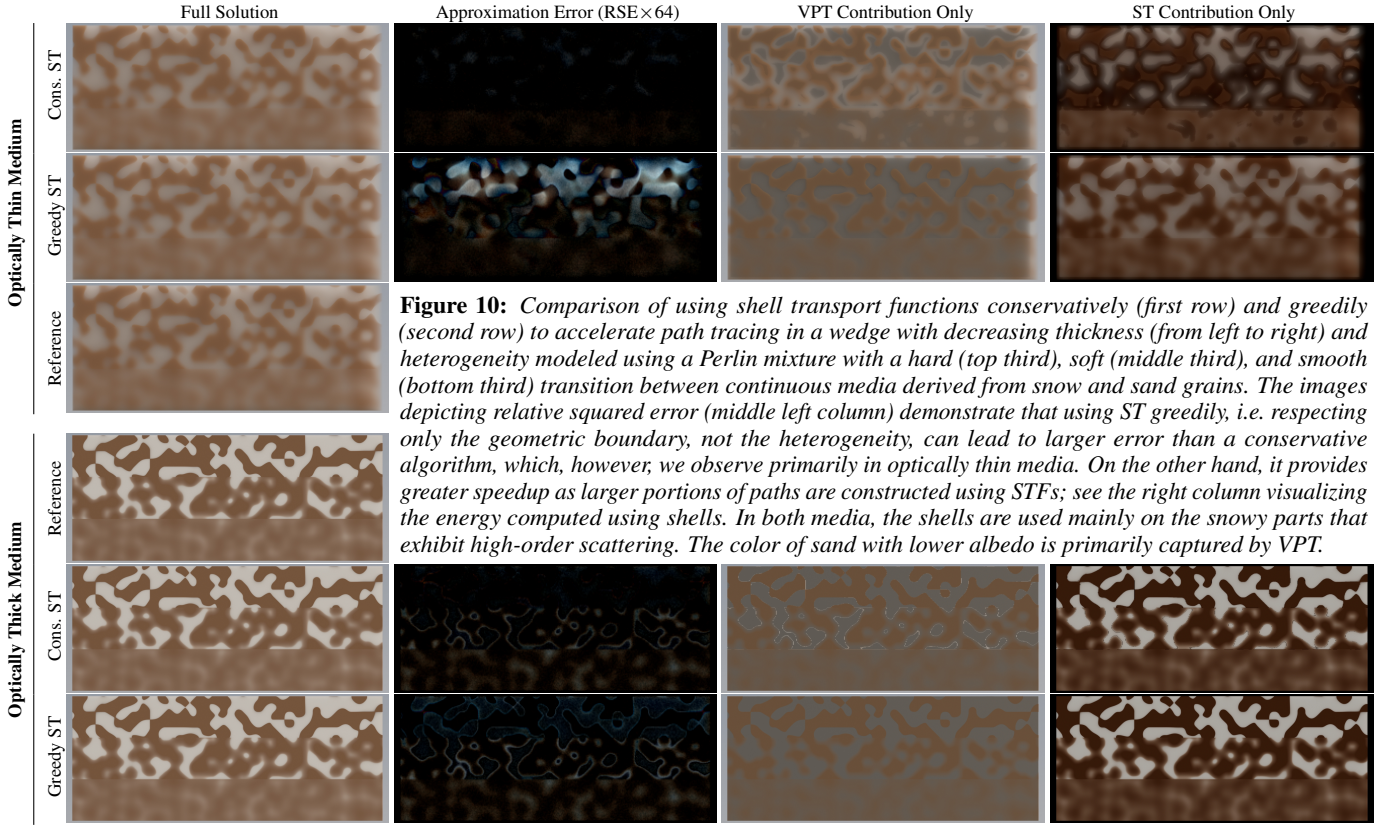

Figure 10: *Comparison of using shell transport functions conservatively (first row) and greedily (second row) to accelerate path tracing in a wedge with decreasing thickness (from left to right) and heterogeneity modeled using a Perlin mixture with a hard (top third), soft (middle third), and smooth (bottom third) transition between continuous media derived from snow and sand grains. The images depicting relative squared error (middle left column) demonstrate that using ST greedily, i.e. respecting only the geometric boundary, not the heterogeneity, can lead to larger error than a conservative algorithm, which, however, we observe primarily in optically thin media. On the other hand, it provides greater speedup as larger portions of paths are constructed using STFs; see the right column visualizing the energy computed using shells. In both media, the shells are used mainly on the snowy parts that exhibit high-order scattering. The color of sand with lower albedo is primarily captured by VPT.*

7.2 Rendering High-Order with Shell Tracing

We switch from the discrete, proxy representation to a continuous volume approximation if the GSDF albedo is above 0.9, and if the path is deeper than 1 mean free path inside the medium—i.e. as soon as the smallest precomputed shell fully fits inside the bounding mesh. We then accelerate path construction using STFs, and only resort to volumetric path tracing near the boundary.

Homogeneous Mixtures. Suppose a light path is to be continued from vertex $\mathbf{x}_c$ in direction $\vec{\omega}_c$. We first look up the STF with the best-matching g and α parameters from our database. Then, we use the largest radial component of this STF that fits inside the mesh to advance the light path. Next, we randomly pick one of the components proportional to their respective α_s. If the uncollided component is chosen, we simply continue the path forward in r mean free path units, i.e. $\mathbf{x}_s = \mathbf{x}_c + r\vec{\omega}_c$, $\vec{\omega}_s = \vec{\omega}_c$. For single scattering, we first sample the free-flight distance to the *one* scattering location, $\mathbf{x}$, then generate direction $\vec{\omega}_s$ by phase function importance sampling, and set $\mathbf{x}_s$ to the shell surface point seen from $\mathbf{x}$ in $\vec{\omega}_s$. To sample multiple scattering, we sample $\mathbf{x}_s$ on the surface of the shell according to the tabulated distribution $p_s^{m,\mathbf{x}}$ and choose $\vec{\omega}_s$ by cosine-weighted sampling of the hemisphere at $\mathbf{x}_s$. We provide pseudocode of this process in the supplementary material.

Heterogeneous Mixtures. In heterogeneous mixtures we choose a STF the same way as in homogeneous mixtures, except that we use the RTE parameters of the voxel containing the *current* path vertex. The difference is in choosing the size (radius) of the STF. We tested two techniques: *greedy shell selection*, which ignores the heterogeneity and selects the shell respecting only the boundary of the assembly (like in homogeneous mixtures), and *conservative shell selection*, limiting the shell size to a locally homogeneous region.

Conservative STF Selection. When conservatively selecting STF sizes, we limit the amount of variation of σ_t and α within the volume of the selected STF. More specifically, we choose an upper bound on the STF radius r by enforcing the following constraints on all locations $\mathbf{x}$ inside the STF:

$$\left|\log_2\left(\frac{\sigma_t(\mathbf{x})}{\sigma_t(\mathbf{x}_c)}\right)\right| < 0.5, \tag{19}$$

$$\left|\log_2\left(\frac{1-\alpha(\mathbf{x})+\epsilon}{1-\alpha(\mathbf{x}_c)+\epsilon}\right)\right| < 0.5. \tag{20}$$

Constraint (19) is motivated by the exponential nature of transmittance $\tau(\sigma_t, t) = \sigma_t e^{-\sigma_t t}$. Constraint (20) is based on the observation that the diffuse reflectance of a continuous volume is roughly proportional to $\log(\frac{1}{1-\alpha})$ up to values very close to 1, at which point the relation becomes sublinear. We set $\epsilon = 0.001$ in order to avoid the singularity at $\alpha = 1$, and to capture the sublinear aspect.

Figure 10 shows a comparison of the greedy and conservative shell selection on a heterogeneous wedge with a variety of optical configurations. The difference images illustrate, how greedy shell tracing introduces additional error when compared to conservative shell tracing, i.e. respecting the heterogeneity. This effect is especially pronounced in optically *thin* media. In optically *thick* media, however, there is relatively small difference between greedy and conservative shell tracing. Since the primary benefit of shell tracing is the acceleration of high-order scattering in optically thick media, we found little merit in using conservative shells and instead select shells greedily in all other renderings in the paper. Nevertheless, we expect that in heterogeneous media with higher spatial density variation as found in our scenes (e.g. an avalanche) the bias from greedy shell selection may increase, in which case conservative selection could offer a superior trade-off between performance and error.

8 Implementation

We implemented our method as a forward path-tracer in the open source Mitsuba renderer [Jakob 2010]. Unlike Meng et al. [2015] we allow grains to be loaded explicitly—e.g. from a simulation—in addition to procedural instantiation. Our only two constraints on the grains are that they have to be oriented uniformly at random and that their bounding spheres must not overlap. Furthermore, we slightly modify the procedural grain instantiation as described by Meng et al. [2015] in order to accelerate ray-grain intersection queries; details can be found in the supplementary material. When comparing our method to the method of Meng et al., we use this optimization in both cases, such that our method does not have an unfair advantage that stems from implementation details.

Multi-channel rendering. STFs introduce complications in multi-channel rendering. Since we compute STFs at discrete radii, which are multiples of the mean-free-path of the medium, the sizes of shells in each channel do not match up if the extinction of the medium varies across channels. This means that we cannot trivially evaluate STFs across channels. We solve this problem by approximating STFs of arbitrary radii by linearly interpolating precomputed STFs with neighboring radii in logarithmic space. For example, if $f_s(r_1)$ and $f_s(r_2)$ were precomputed, and $r_1 \leq r < r_2$, then

$$f_s(r) \approx (1 - w_r) f_s(r_1) + w_r f_s(r_2) \tag{21}$$

$$w_r = \frac{\log(r) - \log(r_1)}{\log(r_2) - \log(r_1)}. \tag{22}$$

The usage of the logarithmic space is motivated by the exponential nature of continuous media, and by empirical experiments, e.g. as shown in Figure 8.

9 Results

In this section we compare our method to reference images rendered with EPT and previous work. To best illustrate the bias each method introduces we present results that are as converged as possible. Unfortunately, some reference images converge extremely slowly (e.g. the SNOWMAN reference in Figure 12 took 13 core-years to compute) and thus still exhibit residual noise. Due to the varying degrees of convergence across images, we do not report absolute rendering times, but instead use *time to unit variance* (TTUV). Additionally, we report *mean relative squared error* (MRSE). A comprehensive summary of the aforementioned metrics on the scenes presented in this paper can be found in Table 4. We provide more detailed comparisons on a larger set of scenes in the supplementary material.

Time to Unit Variance. We define TTUV as the product of the time T it takes to render an image on a single CPU core and its mean pixel variance $V = \frac{1}{n}\sum_{i=1}^{n} V_p[p_i]$ where $V_p[p_i]$ is the variance of the i-th pixel of the image. With the TTUV it becomes trivial to compute the rendering time required for achieving a desired variance $T(V) = \text{TTUV}/V$ and vice versa, the variance achieved after rendering for a desired time $V(T) = \text{TTUV}/T$.

Mean Relative Squared Error. The MRSE is defined as the mean $\frac{1}{n}\sum_{i=1}^{n}(p_i - \hat{p}_i)^2/(\hat{p}_i^2 + \epsilon)$ where p_i and $\hat{p}_i$ are the values of the i-th approximate and reference pixels, respectively, and ϵ is set to 0.01 in order to not over-emphasize very dark pixels and to prevent division by 0. We downscale images by a factor of 4 in both width and height via bilinear filtering before computing MRSE to reduce the influence of residual Monte Carlo noise in the reference image.

9.1 DUNES Scene

The DUNES scene consists of quadrillions of procedurally instantiated, polydisperse sand grains using the tile-based approach by Meng et al. [2015]. Each tile was generated by placing grains of normally-distributed sizes using a dart-throwing algorithm. The resulting distribution of grain radii has a the mean of 0.28 mm, standard deviation of 0.19 mm, and the minimum and maximum of 0.023 mm and 1 mm, respectively.

Each grain has a dielectric boundary and is filled with one of four distinct continuous homogeneous media. These grain types are uniformly distributed throughout the entire scene. In Figure 11 we show renderings of the DUNES scene with just EPT, with the method by Meng et al. [2015] (leveraging our extensions for polydisperse mixtures), and with our method. The scene showcases both the high-frequency detail of grains close to the camera and the large-scale appearance of grains at a large distance. For a reference of scale, the dinosaur skeleton is about 12 m long, and the large dunes in the distance are circa 1 km away.

Both alternatives to EPT are close to the ground truth, but the previous approach produces visibly darker colors. In terms of performance, however, our method converges $5\times$ faster than the algorithm by Meng et al., and $13\times$ faster than EPT. In fact, most of the variance in our method comes from the region in the bottom-right where the individual grains are close and we can not use PPT from the beginning. When cropping away the bottom region of the image, or when neglecting glints, our method becomes about $80\times$ faster than the method by Meng et al.

9.2 SNOWMAN Scene

The SNOWMAN scene by Meng et al. [2015] contains a 1 m tall snowman, made up from a homogeneous mixture of procedural mono-disperse snow grains, which have a radius of 0.1 mm. We render this scene in Figure 12 using our full method (EPT+PPT+ST) and the full method of Meng et al. [2015] (EPT+VPT+DA). Our method is able to render the scene $37.5\times$ faster thanks to the automatic switching criterion that skips EPT and immediately employs PPT. The previous approach is not able to switch to VPT as quickly due to the poor reproduction of color of the continuous volume approximation. We also compare the error introduced by each method versus an explicitly path traced reference. Our method introduces over $18\times$ less MRSE, which can be mostly attributed to the error-prone diffusion approximation.

9.3 TWO PILES Scene

The TWO PILES scene shows piles of blue and pink salt. The grains in each pile are polydisperse and distributed heterogeneously in layers based on a simulation which we performed with Houdini. The resulting distribution of grain radii has a the mean of 1.31 mm, standard deviation of 0.33 mm, and the minimum and maximum of 0.44 mm and 2.97 mm, respectively.

This scene demonstrates simultaneous usage of all our novel features, which the previous state of the art does not support: polydispersivity, acceleration of low-order light transport, and acceleration of high-order light transport in heterogeneous mixtures.

Figure 13 compares the convergence rates and appearances of EPT, PPT, and our full method (EPT+PPT+ST). PPT alone is $28.6\times$ faster than EPT, whereas our full method only achieves a $4.14\times$ speed improvement. The comparison illustrates how capable PPT is at capturing the overall appearance of granular media, only losing the rotation-dependent high-frequency detail such as glints, as can be seen at the top of the blue pile.

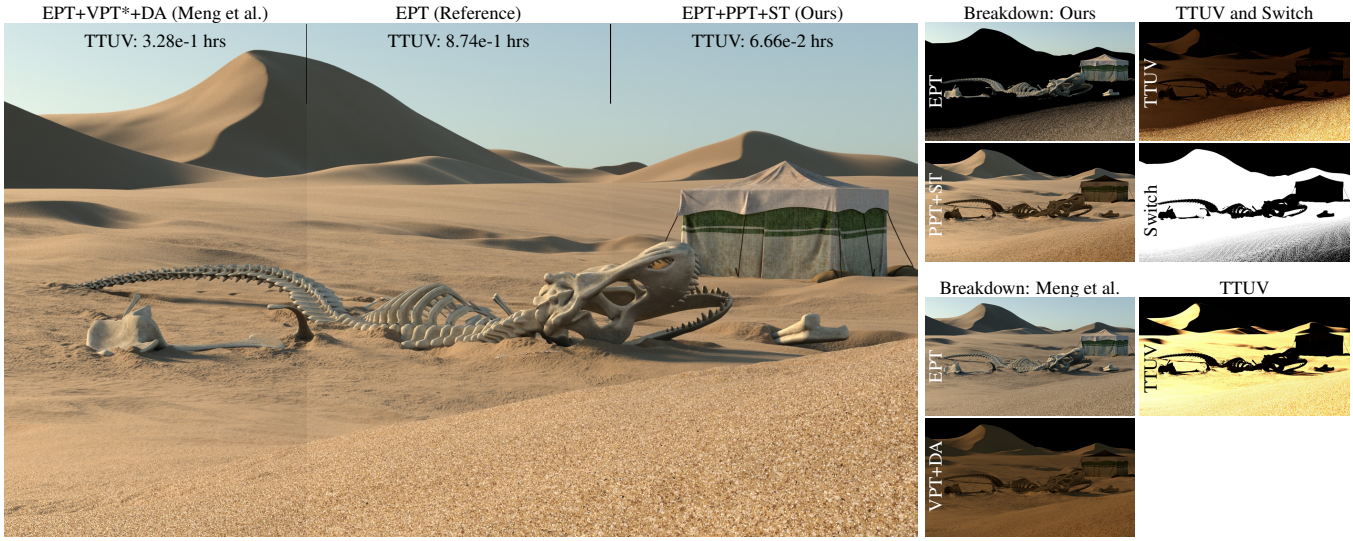

Figure 11: *The* DUNES *scene is a massive static desert landscape consisting of quadrillions of procedurally instantiated polydisperse sand grains with average radius of* 1 *mm. We render the scene with the method of Meng et al. [2015], with explicit path tracing (EPT) for reference, and with our method. Our technique provides a speedup of* $13\times$ *over EPT, whereas the method of Meng et al. accelerates rendering only by factor of* $2.5\times$*. The images on the right depict contributions of EPT and approximations used by each method, per-pixel visualizations of TTUV, and the switching criterion for our method (black: EPT, white: PPT on the first bounce). Our method uses EPT only for the foreground grains where glints can occur, employing PPT for the vast majority of transport. The method of Meng et al. uses EPT for significantly more bounces, but still switches to the continuous approximation too early, leading to small inaccuracies in brightness.*

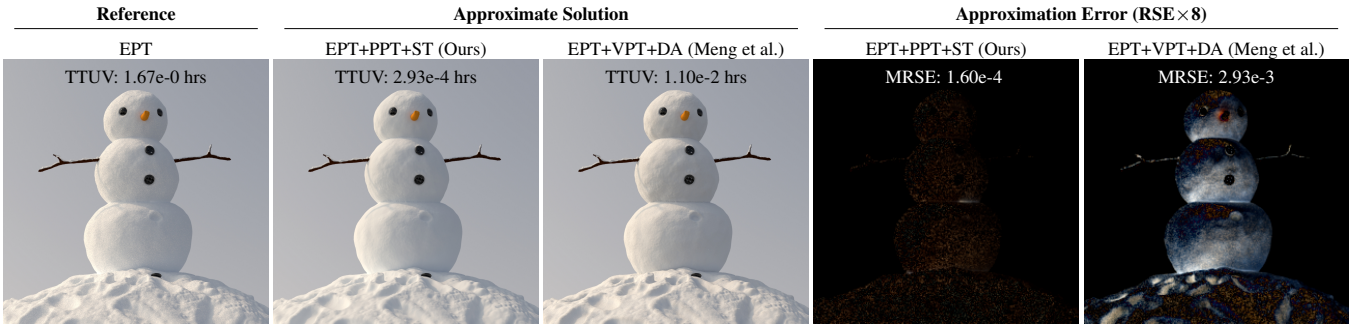

Figure 12: *We render the* SNOWMAN *scene by Meng et al. [2015] using their full method and our full method, and compare the error and time to unit variance. Our method achieves significantly faster convergence* ($37.5\times$) *than previous work, because, in contrast, it can skip the expensive EPT prefix using the automatic switching criterion described in Section 7.1. Additionally, ST introduces significantly less bias than the diffusion approximation of Meng et al. [2015], which we quantify as MRSE and visualize on the right.*

Figure 13: *The* TWO PILES *scene consists of two heterogeneous piles made out of polydisperse salt grains. PPT (top) does not capture glints (see the boundary at top cap of the blue pile) and achieves a speedup of* $28.6\times$ *compared to the EPT reference. Our full method (EPT+PPT+STF) has an overall lower bias than PPT and is able to capture glints, with a speedup of* $4.14\times$ *over the EPT reference. On the right we provide a breakdown of the individual components of our full method. EPT contributes almost exclusively glints, PPT preserves most of the residual high-frequency granular structure, and ST captures low-frequency, high-order light transport.*

Table 4: *We compare the time to unit variance (TTUV) and mean relative squared error (MRSE) for each scene presented in this paper. We show results for a reference solution, the method by Meng et al. [2015], and our method. VPT* (note the asterisk) means that we replaced the volumetric approximation of Meng et al. with our extended version for heterogeneous, polydisperse granular media. The original method by Meng et al. by itself is not able to render these scenes, and their diffusion approximation can not be used in heterogeneous media.*

		Reference		Ours			Meng et al. [2015]		
Scene	Resolution	Method	TTUV	Methods	TTUV (speedup)	MRSE	Methods	TTUV (speedup)	MRSE
DUNES	2.07 MP	EPT	8.74e-1 hrs	EPT+PPT+ST	6.66e-2 hrs (13.1×)	7.25e-4	EPT+VPT*+DA	3.28e-1 hrs (2.53×)	7.82e-3
SNOWMAN	1.04 MP	EPT	1.67e-0 hrs	EPT+PPT+ST	2.93e-4 hrs (5699×)	1.60e-4	EPT+VPT+DA	1.10e-2 hrs (152.×)	2.93e-3
TWO PILES	1.15 MP	EPT	2.97e-1 hrs	EPT+PPT+ST	7.17e-2 hrs (4.14×)	7.86e-4	EPT+VPT*	8.30e-2 hrs (3.58×)	2.60e-3
BOWL (frame 100)	1.31 MP	EPT	4.47e-3 hrs	EPT+PPT+ST	7.88e-4 hrs (5.67×)	1.11e-4	EPT+VPT*	9.36e-4 hrs (4.78×)	5.01e-4
BOWL (frame 1000)	1.31 MP	EPT	2.17e-3 hrs	EPT+PPT+ST	7.79e-4 hrs (2.78×)	1.19e-4	EPT+VPT*	7.27e-4 hrs (2.98×)	5.53e-4
LUCY (1m large)	0.52 MP	VPT	1.61e-3 hrs	VPT+ST	2.31e-4 hrs (6.96×)	1.64e-4	VPT+DA	1.93e-4 hrs (8.35×)	1.25e-3
LUCY (100m large)	0.52 MP	VPT	3.04e-3 hrs	VPT+ST	1.61e-4 hrs (18.9×)	2.14e-4	VPT+DA	2.21e-4 hrs (13.8×)	1.19e-3

10 Discussion and Future Work

Variance and Per-sample Cost. The time per sample of our full method is only at most 5× larger than the per-sample time of rendering corresponding assemblies as *opaque diffuse* surfaces, suggesting there is only a relatively small potential for further optimizing our implementation. The per-sample variance of our method, however, is far higher than that of a diffuse surface. Future work should thus focus on better importance sampling techniques.

Reciprocity. Since we target forward path-tracing, we purposely did not enforce Helmholtz reciprocity to maximize the performance by retaining only the most visually salient features. As such, some parts of our tracing pipeline, e.g. the GSDFs, as well as the switching criteria and the order of individual accelerators would have to be revisited when incorporated into bidirectional algorithms.

Simulating Orientation of Grains. Currently, all dynamic simulations of granular materials that we know of do not simulate the rotational dynamics of grains. In our experience, however, the dynamic behavior of grain orientations is in fact quite important for low-order detail to look convincingly realistic. In the BOWL scene we hallucinate such rotations based on the absolute movement of grains, which yields only moderate realism. Ideally, the input simulation data would include physically plausible grain orientations.

Beyond Non-overlapping Bounding Spheres. By modeling grains within non-overlapping bounding spheres grains cannot be packed tighter than spheres. This assumption is not a big problem for approximately spherical grains, but it breaks down for anisotropic grains such as rice. Without the assumption of non-overlapping bounding spheres, neighboring grains influence the orientations of each other, breaking our assumption of randomly oriented grains.

11 Conclusion

We presented a multi-scale approach for rendering dynamic, heterogeneous, polydisperse assemblies of grains. We proposed to handle the two main appearance components—the high-frequency detail stemming from low-order scattering, and the translucency due to high-order transport—separately, using low-dimensional functions, the GSDFs and STFs, respectively. The conversion of discrete assemblies to corresponding continuous volumes is done using the approach of Meng et al. [2015] that we extended to polydisperse mixtures. Our complete algorithm satisfies current movie-production requirements, yet still outperforms prior art, oftentimes significantly.

Acknowledgements

The authors would like to thank Maurizzio Nitti and Alessia Marra for modeling the scenes and individual grains, Romain Prévost for help with illustrations, and the Stanford 3D Scanning Repository for providing the model used in the LUCY scene.

References

ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. S. 2000. A microfacet-based BRDF generator. In *Proc. SIGGRAPH*, 65–74.

CHANDRASEKAR, S. 1960. *Radiative Transfer*. Dover Publications.

DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *Proc. SIGGRAPH*, 145–156.

D'EON, E., AND IRVING, G. 2011. A quantized-diffusion model for rendering translucent materials. *ACM TOG 30*, 4 (July), 56:1–56:14.

D'EON, E., FRANCOIS, G., HILL, M., LETTERI, J., AND AUBRY, J.-M. 2011. An energy-conserving hair reflectance model. In *Proc. EGSR*, Eurographics Association, 1181–1187.

D'EON, E., MARSCHNER, S., AND HANIKA, J. 2013. Importance sampling for physically-based hair fiber models. In *SIGGRAPH Asia 2013 Technical Briefs*, SA '13, 25:1–25:4.

DIXMIER, M. 1978. Une nouvelle description des empilements aléatoires et des fluides denses. *Le Journal de Physique 39*, 873–895.

DULLIEN, F. A. L. 1991. *Porous Media: Fluid Transport and Pore Structure*, 2nd ed. Academic Press Inc.

HABEL, R., CHRISTENSEN, P. H., AND JAROSZ, W. 2013. Photon beam diffusion: A hybrid monte carlo method for subsurface scattering. *Computer Graphics Forum 32*, 4 (June).

HENYEY, L. G., AND GREENSTEIN, J. L. 1941. Diffuse radiation in the galaxy. *The Astrophysical Journal 93*, 70–83.

HERY, C., AND RAMAMOORTHI, R. 2012. Importance sampling of reflection from hair fibers. *Journal of Computer Graphics Techniques*.

JAKOB, W., MOON, J. T., AND MARSCHNER, S. 2009. Capturing hair assemblies fiber by fiber. *ACM TOG 28*, 5 (Dec.), 164:1–164:9.

JAKOB, W., D'EON, E., JAKOB, O., AND MARSCHNER, S. 2014. A comprehensive framework for rendering layered materials. *ACM TOG 33*, 4 (July), 118:1–118:14.

JAKOB, W., HAŠAN, M., YAN, L.-Q., LAWRENCE, J., RAMAMOORTHI, R., AND MARSCHNER, S. 2014. Discrete stochastic microfacet models. *ACM TOG 33*, 4 (July), 115:1–115:10.

JAKOB, W., 2010. Mitsuba renderer. http://mitsuba-renderer.org.

JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proc. SIGGRAPH*, 511–518.

KAJIYA, J. T., AND KAY, T. L. 1989. Rendering fur with three dimensional textures. In *Computer Graphics*, 271–280.

KAJIYA, J. T. 1986. The rendering equation. *Computer Graphics 20*, 143–150.

KIMMEL, B. W., AND BARANOSKI, G. V. G. 2007. A novel approach for simulating light interaction with particulate materials: application to the modeling of sand spectral properties. *Optics Express 15*, 15 (July), 9755–9777.

LAFORTUNE, E. P., AND WILLEMS, Y. D. 1996. Rendering participating media with bidirectional path tracing. In *Proc. EGWR*, 91–100.

LEE, R., AND O'SULLIVAN, C. 2007. Accelerated light propagation through participating media. In *Proc. Eurographics / Ieee VGTC Conference on Volume Graphics*, 17–23.

LEVOY, M., AND HANRAHAN, P. M. 1996. Light field rendering. In *Proc. SIGGRAPH*, 31–42.

LI, H., PELLACINI, F., AND TORRANCE, K. E. 2005. A hybrid Monte Carlo method for accurate and efficient subsurface scattering. In *Proc. EGSR*, 283–290.

LOOS, B. J., ANTANI, L., MITCHELL, K., NOWROUZEZAHRAI, D., JAROSZ, W., AND SLOAN, P.-P. 2011. Modular radiance transfer. *ACM TOG 30*, 6 (Dec.).

MARSCHNER, S. R., JENSEN, H. W., CAMMARANO, M., WORLEY, S., AND HANRAHAN, P. 2003. Light scattering from human hair fibers. *ACM TOG 22*, 3 (July), 780–791.

MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM TOG 22*, 3 (July), 759–769.

MENG, J., PAPAS, M., HABEL, R., DACHSBACHER, C., MARSCHNER, S., GROSS, M., AND JAROSZ, W. 2015. Multi-scale modeling and rendering of granular materials. *ACM TOG 34*, 4 (July), 49:1–49:13.

MOON, J. T., WALTER, B., AND MARSCHNER, S. R. 2007. Rendering discrete random media using precomputed scattering solutions. In *Proc. EGSR*, 231–242.

NEYRET, F. 1998. Modeling, animating, and rendering complex scenes using volumetric textures. *IEEE Trans. on Visualization and Computer Graphics 4*, 1 (Jan./Mar.), 55–70.

NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. 1992. Geometrical considerations and nomenclature for reflectance. In *Radiometry*. 94–145.

OU, J., XIE, F., KRISHNAMACHARI, P., AND PELLACINI, F. 2012. ISHair: Importance Sampling for Hair Scattering. *Computer Graphics Forum 31*, 4, 1537–1545.

PEERS, P., VOM BERGE, K., MATUSIK, W., RAMAMOORTHI, R., LAWRENCE, J., RUSINKIEWICZ, S., AND DUTRÉ, P. 2006. A compact factored representation of heterogeneous subsurface scattering. *ACM TOG 25*, 3 (July), 746–753.

PHARR, M., AND HANRAHAN, P. M. 2000. Monte carlo evaluation of non-linear scattering equations for subsurface reflection. In *Proc. SIGGRAPH*, 75–84.

RAMAMOORTHI, R. 2009. Precomputation-based rendering. *Found. Trends. Comput. Graph. Vis. 3*, 4 (Apr.), 281–369.

SADEGHI, I., MUÑOZ, A., LAVEN, P., JAROSZ, W., SERON, F., GUTIERREZ, D., AND JENSEN, H. W. 2012. Physically-based simulation of rainbows. *ACM TOG 31*, 1 (Feb.), 3:1–3:12.

SADEGHI, I., BISKER, O., DE DEKEN, J., AND JENSEN, H. W. 2013. A practical microcylinder appearance model for cloth rendering. *ACM TOG 32*, 2 (Apr.), 14:1–14:12.

SCHRÖDER, K., KLEIN, R., AND ZINKE, A. 2011. A volumetric approach to predictive rendering of fabrics. *Computer Graphics Forum 30*, 4 (July), 1277–1286.

SKOGE, M., DONEV, A., STILLINGER, F. H., AND TORQUATO, S. 2006. Packing hyperspheres in high-dimensional Euclidean spaces. *Physical Review E 74*, 4, 041127.

SONG, C., WANG, P., AND MAKSE, H. A. 2008. A phase diagram for jammed matter. *Nature*, 7195, 629–632.

STAM, J. 1995. Multiple scattering as a diffusion process. *Proc. EGWR*, 41–50.

TORQUATO, S., AND LU, B. 1993. Chord-length distribution function for two-phase random media. *Physical Review E 47* (Apr.), 2950–2953.

TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America 57*, 9 (Sept.), 1105–1112.

VEACH, E., AND GUIBAS, L. J. 1995. Optimally combining sampling techniques for monte carlo rendering. In *Proc. SIGGRAPH*, 419–428.

WEI, Y., OFEK, E., QUAN, L., AND SHUM, H.-Y. 2005. Modeling hair from multiple views. *ACM TOG 24*, 3 (July), 816–820.

WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *Computer Graphics*, 255–264.

XU, Y.-Q., CHEN, Y., LIN, S., ZHONG, H., WU, E., GUO, B., AND SHUM, H.-Y. 2001. Photorealistic rendering of knitwear using the lumislice. In *Proc. SIGGRAPH*, 391–398.

YAN, L.-Q., HAŠAN, M., JAKOB, W., LAWRENCE, J., MARSCHNER, S., AND RAMAMOORTHI, R. 2014. Rendering glints on high-resolution normal-mapped specular surfaces. *ACM TOG 33*, 4 (July), 116:1–116:9.

YAN, L.-Q., TSENG, C.-W., JENSEN, H. W., AND RAMAMOORTHI, R. 2015. Physically-accurate fur reflectance: Modeling, measurement and rendering. *ACM TOG 34*, 6 (Oct.), 185:1–185:13.

ZHAO, S., HAŠAN, M., RAMAMOORTHI, R., AND BALA, K. 2013. Modular flux transfer: efficient rendering of high-resolution volumes with repeated structures. *ACM TOG 32*, 4 (July), 131:1–131:12.

ZHAO, S., RAMAMOORTHI, R., AND BALA, K. 2014. High-order similarity relations in radiative transfer. *ACM TOG 33*, 4 (July), 104:1–104:12.

ZINKE, A., AND WEBER, A. 2007. Light scattering from filaments. *IEEE Transactions on Visualization and Computer Graphics 13*, 2, 342–356.