





A Discrete Polydisperse Anisotropic BSDF Model based on the Micrograin Framework

K. Xu¹  S. Lucas^{2,3}  M. Ribardière¹  B. Bringier¹  and P. Barla² 

¹Université de Poitiers, CNRS XLIM UMR 7252, France

²Inria, France

³Université Côte d'Azur, France




Figure 1: Our polydisperse micrograin-based BSDF model enables the representation of complex material appearances composed of multiple non axis-aligned microscopic ellipsoidal elements. The car bodywork material consists of (a) a smooth aluminium bulk surface painted with (b) flat, isotropic bluish particles and (c) elongated, oriented gold glitter. (d) A dusty appearance can be simulated by adding small, diffuse micrograins, which affect retro-reflection at the object's silhouette.

Abstract

We introduce a discrete polydisperse micrograin BSDF model for the rendering of porous surface materials composed of microscopic elements of different size, shape and reflectance distributed on a bulk medium. Our approach generalizes the anisotropic monodisperse micrograin model. We first reformulate it in a non-axis-aligned configuration, allowing for the later combination of different micrograin types elongated in arbitrary directions. We then extend the monodisperse model to the polydisperse case, deriving its three key components: (i) a general filling factor that controls the mix between micrograins and the bulk medium; (ii) an exact normal distribution function for surfaces composed of polydisperse micrograin distributions; and (iii) the corresponding fully-correlated shadowing and masking term. This results in an analytical single-scattering BSDF for discrete polydisperse surface materials, validated over ground truth simulations, for which we also derive a dedicated importance sampling procedure. Our model supports varying heights and anisotropy orientations of different micrograin types as input, giving additional control to simulate phenomena like retro-reflection from mixed materials, color mixture depending on lighting and observation directions, multiple directions of anisotropy, etc.

CCS Concepts

• *Computing methodologies* → *Reflectance modeling*;

1. Introduction

Many natural or man-made materials are polydisperse: they are composed of microscopic elements of different shapes, sizes and

reflectance properties. The modeling of such polydisperse media is a difficult endeavor, with most previous work in Computer Graphics focusing on volumetric materials having granular (e.g., [MPH*15, MPG*16]), or layered (e.g., [EKM01, JDJM14]) structures. In granular materials, polydispersity is handled through simulation due to the difficulty of finding analytical solutions to all the interactions that may occur among diverse microscopic elements. In layered materials, a solution is to discretize a layer in depth and to numerically combine the impact of thin sub-layers.

In this paper we focus on materials that have a porous *surface*, composed of aggregates of microscopic elements of different sizes, shapes, and reflectance, all distributed on the same bulk medium. The recent micrograin formalism [LRPB23, LRPB24, LPB25] deals with porous surfaces, but it is restricted to the mono-disperse case: micrograins are half-ellipsoids having all the same size, shape and reflectance. The main goal of this paper is to extend the micrograin formalism to the polydisperse case. As in the monodisperse case, we only focus on *single scattering* in our approach, and leave the extension to multiple scattering (e.g., see [ZSL*25] for a first monodisperse solution) to future work.

Our key idea is to combine a finite number of different types of micrograins – which we call *discrete* polydispersity – with the objective of retaining the key properties of the micrograin formalism: control over the filling factor, anisotropic normal distribution functions and fully-correlated visibility terms. In terms of appearance, the goal is to provide control over the optical mix of the different components (due to the micrograins of different types and to the bulk medium on which they lie); to allow for highlights of different sharpness, color and elongation across a surface; and to generate interesting retro-reflection and grazing angle effects. An important property of our approach is that it should be equivalent to the monodisperse case when a single type of micrograin is used.

We first recall in Section 3 the basic different terms and formula of the monodisperse micrograin model [LRPB24]. The original formulation was given for an axis-aligned configuration, assuming the full BSDF would later be aligned with local surface tangents. However, in the polydisperse case, different micrograin types may be aligned in different directions. We thus also reformulate the original monodisperse model for the non axis-aligned case in Section 4. We then extend this formulation to the polydisperse case in Section 5. This requires extending all three basic elements of the micrograin formalism: filling factor, normal distribution function and shadowing/masking term. We show the resulting polydisperse micrograin BSDF is a generalization of the monodisperse micrograin BSDF as required. We also describe an importance sampling routine for our polydisperse model in Section 6, which properly reverts to the monodisperse routine when a single type of micrograin is used. Our model is validated against reference simulations and compared to other models in Section 7, showing both its accuracy and specificity. We illustrate the appearance effects afforded by our approach in Section 8 and discuss its limitations and future extensions in Section 9.

2. Previous work

Initially introduced by Torrance and Sparrow [TS67], the microfacet theory models the appearance of a rough microsurface

by considering stochastically distributed microfacets, with orientations following a Normal Distribution Function (NDF). First introduced to the computer graphics community by Cook and Torrance [CT82], the model assumes mirror-like microfacets, though the theory can also be extended to Lambertian microfacets [ON94, HD15]. Masking and Shadowing between microfacets are accounted by the Geometrical Attenuation Factor (GAF), which is typically computed using Smith's derivation [Smi67]. In recent years, both the NDF [WMLT07, DWMG16, RBMS17] and GAF [vGSK98, APS00, RDP05, Hei14, TRM*25] have been the subject of extensive research aimed at broadening the range of appearances that microfacet theory can represent. Although microfacet theory is limited to single scattering, various approaches have been developed to account for multiple scattering among microfacets [HHdD16, XH18, LJJ*18, Bd22].

One limitation of basic microfacet theory is its assumption that all microfacets exhibit identical optical behavior properties. To address this issue, a straightforward approach involves blending of multiple microfacet-based models while ensuring energy conservation [KSK01]. This blending can be implemented vertically, as a layered stack, where rough interfaces are each modeled using a microfacet-based model (e.g., [JDJM14, WW07, MBT*18, GHZ18]). Lateral blending has also been explored, yielding hazy gloss effects [BPV18]. It may also be used for the representation of porous materials [MDG02], in which case pores are modeled as cylindrical cavities with a diffuse reflectance, with the base surface and pores combined using an empirical porosity parameter. D'Eon et al. [dBWZ23] employed a height-based mixture of multiple microfacet distributions to enable retro-reflective effects at grazing angles.

Microfacet theory can be further refined to simulate other types of effects. Some approaches adopt a volumetric description, replacing microfacets with microflakes [JDJM14, DHd16, WJHY22]. However, these volumetric approaches do not reproduce the grazing-angle visibility effects observed in porous surface-based models. Recently, Lucas et al. [LRPB23, LRPB24, LPB25] introduced a novel microsurface representation based on micrograin theory. Half-ellipsoidal micrograins, each sharing the same BSDF and axis aligned elongation, are deposited atop a bulk surface. The authors derive an exact, height-dependent, and fully-correlated GAF that accounts for correlations between both incoming and outgoing directions. Additionally, they proposed an accurate weighting scheme for mixing contributions of the bulk surface and the micrograins. Zhou et al. [ZSL*25] took an initial step toward incorporating multiple scattering effects. However, their model remains limited to a biased version of micrograin theory, relying on the Smith GAF rather than the fully correlated formulation.

3. Micrograin formalism

A fully-correlated anisotropic micrograin BSDF model has been proposed to render porous materials by Lucas et al. [LRPB24]: It improves their own model [LRPB23] by adding an anisotropic NDF and an exact shadowing/masking term. The model assumes a uniform random distribution of identical micrograins (opaque half-ellipsoids) laid on a base bulk medium of arbitrary (potentially transparent) appearance. The BSDF $f(\mathbf{i}, \mathbf{o})$ is then defined as a mix-

ture of the micrograins BRDF $f_g(\mathbf{i}, \mathbf{o})$ and bulk BSDF $f_B(\mathbf{i}, \mathbf{o})$:

$$f(\mathbf{i}, \mathbf{o}) = \tau_0 f_g(\mathbf{i}, \mathbf{o}) + V_p(\mathbf{i}, \mathbf{o}) f_B(\mathbf{i}, \mathbf{o}) \quad (1)$$

where $\tau_0 \in [0, 1)$ is the surfacic filling factor that controls the coverage of micrograins, $V_p(\mathbf{i}, \mathbf{o}) = (1 - \tau_0) G^{\text{dist}}(\mathbf{i}, \mathbf{o}, 0)$ is the visibility of the bulk interface in both incident direction \mathbf{i} and outgoing direction \mathbf{o} , with $G^{\text{dist}}(\mathbf{i}, \mathbf{o}, h)$ the height-dependent distant bidirectional GAF evaluated in this case at $h = 0$.

Filling factor The filling factor τ is a fundamental property of the micrograin model. In its general formulation, it defines a height- and direction-dependent coverage of micrograins. Lucas et al. [LRPB24] have shown that it is directly related to the projection area σ of a single micrograin at height h from a direction \mathbf{d} :

$$\tau(h, \mathbf{d}) = 1 - \exp(-\rho\sigma(h, \mathbf{d})), \quad (2)$$

with ρ is the the micrograin density. The surface filling factor τ_0 of Equation 1 is defined as the micrograin coverage at the bulk interface (i.e., at $h = 0$), projected along the macrosurface normal \mathbf{n} : in case of half-sphere micrograins, $\tau_0 = \tau(0, \mathbf{n})$ with $\sigma(0, \mathbf{n}) = \pi r^2$ where r is the micrograin radius.

Normal distribution function

The NDF provided in [LRPB24] Section 4.1 allows to stretch the micrograin shape anisotropically along two axis-aligned directions, controlled by two roughness parameters β_x and β_y . Considering a normal \mathbf{m}' on a spherical micrograin surface (i.e. $\beta_x = \beta_y = 1$), the NDF $D(\mathbf{m}')$ is defined as:

$$D(\mathbf{m}') = \frac{-\ln(1 - \tau_0)(1 - \tau_0)^{\sin^2 \theta_{\mathbf{m}'}}}{\pi \tau_0} \quad (3)$$

The subsequent non-axis-aligned micrograin (Section 4) and the polydisperse extension (Section 5) are built upon this definition.

Shadowing and Masking term In the original micrograin model, the GAF is classically composed of local and distant terms: $G(\mathbf{i}, \mathbf{o}, \mathbf{m}, h) = G^{\text{local}}(\mathbf{i}, \mathbf{o}, \mathbf{m}) G^{\text{dist}}(\mathbf{i}, \mathbf{o}, h)$, with the specificity that the distant GAF also depends on the height h . For conciseness, we also adopt the unidirectional notation $G^{\text{dist}}(\mathbf{v}, h)$ when a single direction is involved, following the conventions of the original micrograin formulation. The distant GAF is also used in the V_p term of Equation 1 for $h = 0$. To account for anisotropy, Lucas et al. compute visibility terms on spherical micrograins after applying the inverse of the stretching transform $M = \text{diag}(\frac{1}{\beta_x}, \frac{1}{\beta_y}, 1)$, and stretch the result back to account for anisotropy [LRPB24] Section 4.2.

4. Non axis-aligned monodisperse micrograin

Here we introduce a more general form of non-axis-aligned anisotropy, on which our polydisperse micrograin BSDF model will be based.

Normal distribution function We extend NDF through a 3×3 homogeneous matrix M , which performs a non-axis-aligned anisotropic stretching of the surface, with $\det M \neq 0$ to avoid a vanishing of micrograins.

Following the established property of the filling factor recalled in Equation 2, we rewrite the surface filling factor τ_0 as:

$$\tau_0 = 1 - \exp(-\rho\pi r^2 |\det M|) \quad (4)$$

where we have used $\sigma(0, \mathbf{n}) = \pi r^2 |\det M|$ to account for anisotropy.

The non-axis-aligned anisotropic NDF is then expressed as:

$$D(\mathbf{m}) = \frac{|\det M|}{\|M^T \mathbf{m}\|^4} D(\mathbf{m}'), \quad (5)$$

where $D(\mathbf{m}')$ is defined in Equation 3, $\mathbf{m}' = \frac{M^T \mathbf{m}}{\|M^T \mathbf{m}\|}$ is the normal transformed back to spherical micrograin space, and the stretched surface filling factor of Equation 4 is used for τ_0 for proper normalization. Derivation details and the corresponding importance sampling routine are given in Appendices A and B respectively. The axis-aligned NDF is retrieved when M is a diagonal matrix.

Also note that since τ_0 is exposed as a material parameter, neither ρ nor r need to be explicitly provided in the monodisperse case.

For shadowing and masking term, the process remains unchanged when using a non-diagonal stretching matrix M .

5. Polydisperse micrograin BSDF model

We extend the monodisperse micrograin model to the polydisperse case by considering that the microsurface is covered with a finite number N of different types of micrograins, each having their own size, shape and reflectance properties. An illustration is given in Figure 2[a,b], for $N = 2$ types of micrograins.

Since micrograins of different types may intersect, our model remains a mix between a micrograins BRDF (of different types) and a bulk BSDF, having a form very similar to Equation 1:

$$f(\mathbf{i}, \mathbf{o}) = \bar{\tau}_0 \bar{f}_g(\mathbf{i}, \mathbf{o}) + \bar{V}_p(\mathbf{i}, \mathbf{o}) f_B(\mathbf{i}, \mathbf{o}). \quad (6)$$

The only difference with the monodisperse case is that $\bar{\tau}_0$, \bar{f}_g and \bar{V}_p now respectively denote the surface filling factor, micrograins BRDF and bulk visibility due to the global distribution of micrograins (i.e., considering all types of micrograins). We begin by deriving $\bar{\tau}_0$ in Section 5.1. Next, in Sections 5.2 and 5.3, we introduce the two main ingredients of \bar{f}_g , namely the polydisperse NDF and GAF; the latter being used to define \bar{V}_p . Both the NDF and GAF are lastly incorporated in the diffuse and specular terms of \bar{f}_g as detailed in Section 5.4.

5.1. Global filling factor

The filling factor remains a fundamental property of the polydisperse micrograin model. We refer to it as the global filling factor to denote the coverage made by micrograins of all types.

Denoting by ρ the combined density of all different types of micrograins, the density of the i th type is defined as $\omega_i \rho$, where ω_i is a weighting factor with $\sum_{i=1}^N \omega_i = 1$. Using a weighted sum of the micrograin areas σ_i of different types in Equation 2, we obtain (see Appendix C for details):

$$\bar{\tau}(h, \mathbf{d}) = 1 - \exp\left(-\rho \sum_{i=1}^N \omega_i \sigma_i(h, \mathbf{d})\right). \quad (7)$$

The filling factor $\tau_i(h, \mathbf{d})$ of each type of micrograin may be derived in the same way, yielding $\tau_i(h, \mathbf{d}) = 1 - \exp(-\rho \omega_i \sigma_i(h, \mathbf{d}))$.

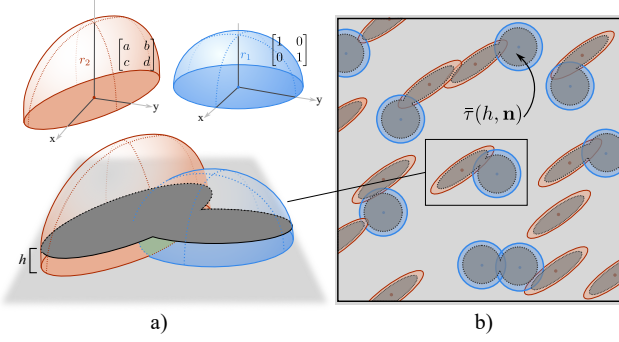


Figure 2: (a) Multiple, potentially non axis-aligned micrograins are distributed across a bulk interface. An example of a slice at height h where two different micrograins intersect is also shown. (b) The global filling factor, denoted as $\bar{\tau}(h, \mathbf{n})$, is evaluated for a slice at height h , encompassing all the micrograins.

As a result, the global filling factor may be rewritten in terms of τ_i :

$$\bar{\tau}(h, \mathbf{d}) = 1 - \prod_{i=1}^N (1 - \tau_i(h, \mathbf{d})). \quad (8)$$

Intuitively, Equation 8 means that the area covered by all micrograins is obtained by the complement of the area *not* covered by *any* type of micrograin, as illustrated in Figure 2[b].

The global surface filling factor $\bar{\tau}_0$ is then trivially given by:

$$\bar{\tau}_0 = 1 - \prod_{i=1}^N (1 - \tau_{i,0}), \quad (9)$$

where $\tau_{i,0} := \tau_i(0, \mathbf{n})$ is the surface filling factor of the i th micrograin type:

$$\tau_i(0, \mathbf{n}) = 1 - \exp(-\rho\omega_i\pi r_i^2 |\det M_i|) \quad (10)$$

with r_i and M_i the micrograin radius and stretching matrix for type i respectively. As required, when $N = 1$ we get $\bar{\tau}_0 = \tau_{1,0}$ as in the monodisperse case.

5.2. Normal distribution function

As derived in Appendix D, the polydisperse NDF is given by:

$$D(\mathbf{m}) = \frac{-1}{\bar{\tau}_0\pi} \sum_{i=1}^N \frac{|\det M_i|}{\|M_i^T \mathbf{m}\|^4} \ln(1 - \tau_{i,0}) \times \prod_{j=1}^N (1 - \tau_{j,0}) \left[1 - \left(\frac{r_i \cos \theta_{\mathbf{m}'_i}}{r_j} \right)^2 \right]^+, \quad (11)$$

where $\mathbf{m}'_i = \frac{M_i^{-1} \mathbf{m}}{\|M_i^{-1} \mathbf{m}\|}$ is the normal mapped back to the spherical micrograin space of type i , with $\cos \theta_{\mathbf{m}'_i} = \frac{\cos \theta_{\mathbf{m}}}{\|M_i^T \mathbf{m}\|}$. The clamping operator $[\cdot]^+$ ensures that its argument remains within $[0, +\infty)$. Here, $r_i \cos \theta_{\mathbf{m}'_i}$ represents the geometric height reached by the transformed normal \mathbf{m}'_i on micrograin type i , see Figure 3. Dividing by r_j gives its normalized height relative to micrograin type j .

Accordingly, the factor of the product in the second line of Equation 11 acts as an attenuation term expressing how much type j occludes or intersects the contribution of type i at that elevation. Since all micrograin types are assumed to be randomly distributed and may mutually overlap, the contribution of type i is reduced proportionally to the overlap probability with type j at the same height (Figure 3[a]). When the current height exceeds the extent of type j , i.e. $\left(\frac{r_i \cos \theta_{\mathbf{m}'_i}}{r_j} \right)^2 \geq 1$, the exponent becomes zero, the attenuation term reduces to 1, and type j ceases to affect the contribution from type i , as no intersections can occur beyond that elevation (Figure 3[b]).

Note that when $N = 1$, only the case $i = j = 1$ remains in the second line of Equation 11; hence we have $(\sin \theta_{\mathbf{m}'_1})^2 \geq 0$ for the exponent and we recover the monodisperse NDF of Equation 5 as required.

In terms of material parameters, we choose to expose all surface filling factors $\tau_{i,0}$, from which the global surface filling factor $\bar{\tau}_0$ is derived using Equation 9. As in the monodisperse case, the density ρ does not need to be provided; however the different micrograin radii r_i must be specified, up to a scaling factor since only ratios of radii are involved in Equation 11.




Figure 3: Surface points of micrograin type i at height $r_i \cos \theta_{\mathbf{m}'_i}$ are shown as a green ring. (a) When $r_i \cos \theta_{\mathbf{m}'_i} < r_j$, these points can be occluded by type j . (b) When $r_i \cos \theta_{\mathbf{m}'_i} \geq r_j$, these points can never be occluded by type j .

5.3. Shadowing and masking term

As in the monodisperse case, the polydisperse GAF is a product of a local and distant terms: $G(\mathbf{i}, \mathbf{o}, \mathbf{m}) = G^{\text{local}}(\mathbf{i}, \mathbf{o}, \mathbf{m}) G^{\text{dist}}(\mathbf{i}, \mathbf{o}, h)$.

The local term remains unchanged as it only depends on local vectors: $G^{\text{local}}(\mathbf{i}, \mathbf{o}, \mathbf{m}) = \chi^+(\mathbf{i} \cdot \mathbf{m}) \chi^+(\mathbf{o} \cdot \mathbf{m})$, where $\chi^+(x)$ is the Heaviside step function, equal to 1 if $x > 0$ and 0 otherwise.

The distant term is more involved, since for *any* micrograin type,

we need to account for the visibility according to *all* micrograin types. We first introduce the following effective distant GAF for a micrograin of *any* type i :

$$\tilde{G}_i^{\text{dist}}(\mathbf{i}, \mathbf{o}, \mathbf{m}) = \prod_{j=1}^N G_j^{\text{dist}}\left(\mathbf{i}, \mathbf{o}, \left[\frac{h_i(\mathbf{m})}{r_j}\right]^1\right), \quad (12)$$

where the point under consideration is assumed to lie on a micrograin of type i at height $h_i(\mathbf{m}) = \left(\frac{M_i^T \mathbf{m}}{\|M_i^T \mathbf{m}\|} \cdot \mathbf{n}\right) r_i$. Each factor is based on the monodisperse distant GAF G_j^{dist} applied to micrograins of type j (Section 3). Clamping by $[\cdot]_0^1$ ensures there is no shadowing nor masking (i.e., $G_j^{\text{dist}}(\mathbf{i}, \mathbf{o}, 1) = 1$) when the considered point is above all micrograins of type j (i.e., $h_i(\mathbf{m}) > r_j$). The product over all micrograin types implies that, to be visible, a point must not be masked by any other micrograin type. Using this notation, the polydisperse distant GAF writes compactly as:

$$G^{\text{dist}}(\mathbf{i}, \mathbf{o}, \mathbf{m}) = \sum_{i=1}^N p(i|\mathbf{m}) \tilde{G}_i^{\text{dist}}(\mathbf{i}, \mathbf{o}, \mathbf{m}). \quad (13)$$

The $p(i|\mathbf{m})$ term characterizes the probability of the considered point to be on a micrograin of type i , knowing that it has a normal \mathbf{m} . Following Bayes' rule, it is given by $p(i|\mathbf{m}) = \frac{p(i, \mathbf{m})}{p(\mathbf{m})}$. Here $p(\mathbf{m}) = D(\mathbf{m}) \cos(\theta_{\mathbf{m}})$, while $p(i, \mathbf{m})$ has a similar form, except that it is restricted to the i th micrograin type (see Equation 31):

$$p(i, \mathbf{m}) = \frac{-\cos \theta_{\mathbf{m}} |\det M_i|}{\bar{\tau}_0 \pi \|M_i^T \mathbf{m}\|^4} \ln(1 - \tau_{i,0}) \times \prod_{j=1}^N (1 - \tau_{j,0}) \left[1 - \left(\frac{r_i \cos \theta_{\mathbf{m}_i}}{r_j}\right)^2\right]^+ \quad (14)$$

The visibility term $V_p(\mathbf{i}, \mathbf{o})$ of Equation 6 has a simpler form since it only considers points on the bulk interface (i.e., at a height $h = 0$, with a coverage of $1 - \bar{\tau}_0$). Yet it still has to account for visibility according to *all* types of micrograins, yielding:

$$\bar{V}_p(\mathbf{i}, \mathbf{o}) = (1 - \bar{\tau}_0) \prod_{j=1}^N G_j^{\text{dist}}(\mathbf{i}, \mathbf{o}, 0). \quad (15)$$

5.4. Polydisperse micrograin BRDF

The polydisperse micrograin BRDF \tilde{f}_g is composed of a diffuse and a specular term, as in the monodisperse case: $\tilde{f}_g(\mathbf{i}, \mathbf{o}) = \tilde{f}_g^s(\mathbf{i}, \mathbf{o}) + \tilde{f}_g^d(\mathbf{i}, \mathbf{o})$. However, polydispersity allows for micrograins of different types to have different reflectance properties. This requires a slight adjustment of the common microfacet equations.

Starting with the specular term, we separate the local and distant GAFs in order to treat both visibility and reflectance per type of micrograin, yielding:

$$\tilde{f}_g^s(\mathbf{i}, \mathbf{o}) = \frac{D(\mathbf{h}) G^{\text{local}}(\mathbf{i}, \mathbf{o}, \mathbf{h})}{4|\mathbf{i} \cdot \mathbf{n}||\mathbf{o} \cdot \mathbf{n}|} \sum_{i=1}^N p(i|\mathbf{h}) \mathcal{G}_i^s(\mathbf{i}, \mathbf{o}, \mathbf{h}), \quad (16)$$

where $\mathcal{G}_i^s(\mathbf{i}, \mathbf{o}, \mathbf{h}) = \tilde{G}_i^{\text{dist}}(\mathbf{i}, \mathbf{o}, \mathbf{h}) F_i(\mathbf{i} \cdot \mathbf{h}, \mathbf{h})$ encapsulates both the distant GAF and the Fresnel reflectance for micrograins of type i . Note that we also retain the possibility of having a varying reflectance according to the location on the micrograin, here identified by the halfway vector $\mathbf{h} = \frac{\mathbf{i} + \mathbf{o}}{\|\mathbf{i} + \mathbf{o}\|}$, used as an additional parameter to F_i .

The definition of the diffuse term follows a similar separation:

$$\tilde{f}_g^d(\mathbf{i}, \mathbf{o}) = \int_{\Omega^+} \frac{|\mathbf{i} \cdot \mathbf{m}| |\mathbf{o} \cdot \mathbf{m}|}{|\mathbf{i} \cdot \mathbf{n}| |\mathbf{o} \cdot \mathbf{n}|} D(\mathbf{m}) G^{\text{local}}(\mathbf{i}, \mathbf{o}, \mathbf{m}) \times \sum_{i=1}^N p(i|\mathbf{m}) \mathcal{G}_i^d(\mathbf{i}, \mathbf{o}, \mathbf{m}) d\mathbf{m} \quad (17)$$

where $\mathcal{G}_i^d(\mathbf{i}, \mathbf{o}, \mathbf{m}) = \tilde{G}_i^{\text{dist}}(\mathbf{i}, \mathbf{o}, \mathbf{m}) \frac{K_i^d(\mathbf{m})}{\pi} T_i(\mathbf{i}, \mathbf{o}, \mathbf{m})$ encapsulates both the distant GAF and the effective diffuse reflectance for micrograins of type i , which is given by the product of the diffuse albedo K_i^d and the bidirectional transmittance $T_i(\mathbf{i}, \mathbf{o}, \mathbf{m}) = (1 - F_i(\mathbf{i} \cdot \mathbf{m}, \mathbf{m}))(1 - F_i(\mathbf{o} \cdot \mathbf{m}, \mathbf{m}))$ at the surface of a micrograin. We also retain the possibility of having a varying diffuse reflectance according to the location on the micrograin, here identified by the vector \mathbf{m} used as a parameter to K_i^d .

6. Importance sampling

The sampling of Equation 6 first requires to choose which one of its two components, \tilde{f}_g or f_B , should be sampled. This is done according to the probability p_B of sampling the bulk medium:

$$p_B(\mathbf{i}) = \frac{(1 - \bar{\tau}_0) \prod_{i=1}^N G_i^{\text{dist}}(\mathbf{i}, 0)}{\bar{\tau}_0 + (1 - \bar{\tau}_0) \prod_{i=1}^N G_i^{\text{dist}}(\mathbf{i}, 0)}. \quad (18)$$

Compared to the monodisperse case, p_B must account for the visibility due to all types of micrograins at the bulk interface. This is done through the unidirectional GAFs $G_i^{\text{dist}}(\mathbf{i}, 0)$ (Section 3). Then the probability to sample the polydisperse micrograin component \tilde{f}_g is $1 - p_B(\mathbf{i})$. At normal incidence, since $G_{i,1}^{\text{dist}}(\mathbf{n}, 0) = 1$ (i.e., no shadow), $p_B(\mathbf{n}) = 1 - \bar{\tau}_0$ as expected.




Figure 4: We order the micrograin types with increasing radii for importance sampling. This simplifies the partitioning of heights into N slices S_k (here $N = 3$), circumventing discontinuity issues.

Sampling of the micrograins BRDF As the micrograins BRDF is based on the microfacet theory, a common approach is to sample its NDF. However, this is a challenging task as the NDF of Equation 11 is not smooth: the different sizes of micrograins cause C^1 discontinuities in the NDF (illustrated in Figure 3[a-b]). Our solution is to order the micrograin types by sizes, which allows us to divide the micro-surface into N height slices, denoted S_k in Figure 4 (here we use $N = 3$). Now the height distribution inside a given slice is continuous. As detailed in Appendix E, our sampling process then works in four main steps: (1) we sample a slice

S_k according to the probability $p(S_k)$; (2) we sample a micrograin type i among the ones appearing in S_k according to the probability $p(i|S_k)$; (3) we sample the diffuse or specular component depending on the probabilities $p_i^d(\mathbf{i})$ and $1 - p_i^d(\mathbf{i})$; (4) if the specular component is chosen, we sample a normal \mathbf{m} in the NDF restricted to slice S_k and micrograin type i according to the PDF $p(\mathbf{m}|i, S_k)$. In case of diffuse component, the normal is sampled via cosine-weighted hemisphere sampling.

7. Validation

We validate our model in two steps. First, we assess the accuracy of the polydisperse NDF of Section 5.2. Next, we evaluate the full BRDF of Section 5.4 (with an opaque bulk medium). In addition, we qualitatively compare our model to related existing approaches.

We construct an explicit mesh for the polydisperse micrograin surface. We sample its normals to validate our NDF, and use it in light transport BRDF simulations [BRMS20] to validate the full BRDF. A subset of validations are shown here; additional results are provided in supplementary materials.

Validation of the NDF We construct four monodisperse surfaces, each built from a distinct configuration of non-axis-aligned micrograins (denoted as A , B , C and D) and with:

- Config. A : $r_A = 0.66$ and $\tau_{A,0} = 0.2$ (conductor aluminium (Al))
- Config. B : $r_B = 1$ and $\tau_{B,0} = 0.15$ (conductor gold (Au))
- Config. C : $r_C = 0.66$ and $\tau_{C,0} = 0.1$ (bluish plastic)
- Config. D : $r_D = 0.53$ and $\tau_{D,0} = 0.2$ (greenish diffuse).

Next, we combine these configurations pairwise to create two additional surfaces (denoted as AB and CD). Finally, we generate a final configuration that integrates all four initial configurations (denoted as $ABCD$), based on predetermined radius parameters. The NDFs obtained from these configurations are shown in Figure 5. Our analytical NDF model exhibits near-perfect agreement with simulation results. Notably, the polydisperse NDF exhibits C^1 discontinuities (see Figure 2[c,d]). Our analytical formula fully captures this characteristic.

Validation of the full BRDF To further validate our approach, we assign distinct BRDFs to each micrograin within the previously studied polydisperse configurations (AB , CD and $ABCD$), deposited on absorbing bulk. We then compare the resulting BRDFs against light transport simulations, demonstrating precise agreement between our model and simulation data. Results are shown in Figure 6. Each image represents a BRDF slice parametrized according to azimuthal (x -axis) and polar (y -axis) outgoing angles. The relative errors are computed by taking the difference between the model and the reference, and normalizing by the reference. The artefacts that appear at grazing angles mainly stem from finite-patch boundary visibility discontinuities and increased Monte Carlo variance at grazing angles; moreover, relative-error normalization amplifies small absolute deviations when the reference BRDF value becomes small, which often leads to inaccuracies in simulations under such extreme conditions. Nonetheless, our model faithfully reproduces the fused retro-reflection effects observed at these angles, capturing the complex interactions between micrograin materials.

Comparison between models We compare our model against

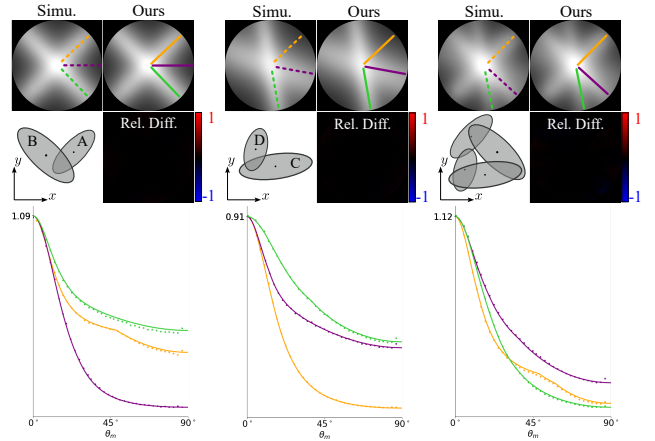


Figure 5: Validation of the NDF (solid lines) against simulations (dots) as a function of the incident angle θ_i and for 3 polydisperse non axis-aligned anisotropic configurations, from left to right : 2 categories AB , 2 categories BC and 4 categories $ABCD$.

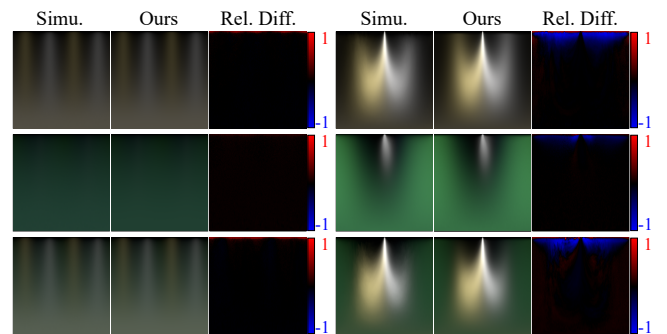


Figure 6: Validation of BRDF slices (x -axis : $\phi_0 \in [0, 2\pi)$ and y -axis : $\theta_0 \in [0, \frac{\pi}{2}]$)

for 2 incident angles (left: $\theta_i = 0^\circ$, right: $\theta_i = 87^\circ$) considering various micrograins arrangement and materials from Figure 5: AB (first row), CD (second row) and $ABCD$ (third row).

the most closely related approaches in the literature in Figure 7. Few BRDF models are capable of blending multiple material appearances to accurately represent color variations depending on both observation and incident light directions. We parameterize our polydisperse model with two categories of isotropic micrograins ($r_1 = 0.5$, gold material for the first; $r_2 = 1$, copper material for the second), deposited onto a black bulk surface. We observe variations in appearance by adjusting the filling factor. Our analysis compares these configurations with the model proposed by d'Eon et al. [dBWZ23], which can represent asymmetric surface height-fields by combining multiple NDFs. Each NDF can be linked to a distinct BRDF, enabling color variations at grazing angles. The second benchmark model is from Lucas et al. [LRPB24]. We use a nested version of the original model, where copper micrograins are deposited atop a bulk surface represented by a gold micrograin model on top of a black bulk surface. Differences are subtle at low filling factors, where the bulk surface appearance dominates (see

the false color insets encoding $\overline{\text{FLIP}}$ error [ANA*20] in Figure 7). However, as the filling factor increases, differences become more pronounced, particularly at grazing angles, where our model preserves the appearance of the golden micrograins. It is important to note that our goal is not to claim superiority over existing models. Instead, our model offers broader parameterization, especially through polydispersity, enabling finer control over retro-reflection effects and color mixing. This is particularly evident on the silhouettes of the objects and at grazing incident and observation angles.

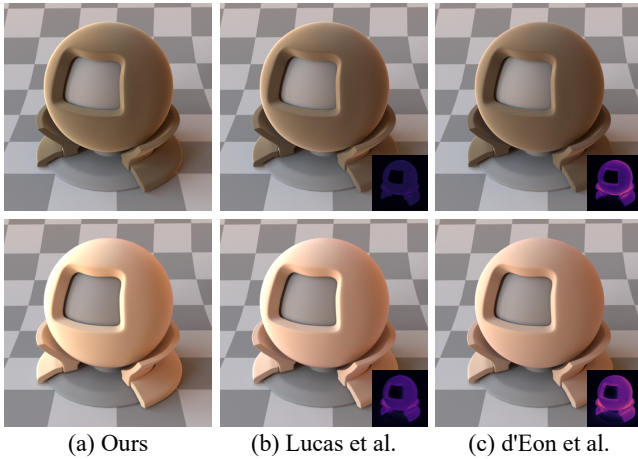


Figure 7: Comparisons between global illumination renderings using our polydisperse micrograin model (a), the model of Lucas et al. [LRPB24] (b) and the one d'Eon et al. [dBWZ23] (c). Polydispersity in our model is achieved with two isotropic micrograins: one made of gold ($r_1 = 0.5$) and the other of copper ($r_2 = 1$). The first row shows results for $\tau_{1,0} = \tau_{2,0} = 0.1$; the second row for $\tau_{1,0} = \tau_{2,0} = 0.5$. For the model of d'Eon et al. [dBWZ23], an equivalent roughness is derived using the same procedure as in [LRPB24]; polydispersity is simulated using a nested version of their approach.

8. Results

We have implemented our model in Mitsuba renderer [JSR*22] and used it to compute all renderings in this paper.

Since variations in radius directly influence the surface filling factors $\tau_{i,0}$ (as defined in Equation 10), which is the key parameter directly governing an object's visual appearance, we adopt an approach similar to the monodisperse model: weights ω_i and density ρ are treated as hidden parameters, while the $\tau_{i,0}$ are exposed as user-adjustable parameters. This approach also provides the added benefit of enabling appearance control through independent adjustments to the radii ratio of different categories. However, this mechanism introduces a limitation: users cannot directly modify the global surface filling factor $\overline{\tau}_0$. Instead, it is computed by Equation 9.

The non-axis-aligned property of our model enables easy control over the orientation of anisotropy. Figure 8 illustrates this capability with three types of conductor micrograins, each with identical shape but different orientations, deposited on a dark gray dif-

fuse bulk. This results in highlights of different colors, elongated in three different directions.




Figure 8: The non-axis-aligned property provides precise control over anisotropy orientation, especially when applied to conductor micrograins. By incorporating polydispersity, complex anisotropic appearance is easily designed. In the bottom-right image, all categories share the same shape and surface filling factor.

Beyond complex anisotropy, our model allows for different micrograin types to exhibit distinct sizes (varying radii r_i), while characterizing the amount of their protrusion above the bulk surface. This feature influences material appearance at grazing angles, as illustrated in Figure 9. By adjusting the size of each micrograin type, users gain additional control over the appearance at grazing angles, with these effects further modulated by variations in the surface filling factors $\tau_{i,0}$. We show these variations of appearance in the isotropic case: we vary radii ratios and surface filling factors. When the filling factor is fixed, increasing the radius ratio enhances the protrusion of micrograins from the surface while simultaneously reducing density to maintain the filling factor (see the diagrams at the top right of Figure 9). The $\overline{\text{FLIP}}$ error between each line shows the visual difference between the upper and lower images. When the height of one micrograin category exceeds that of another, the corresponding material dominates visually at grazing angles. As the filling factor increases, this dominance progressively extends for viewing angles closer to the normal direction. However, once a material reaches its dominant limit at grazing angles for a fixed filling factor, further height increases produce no perceptible visual effect. Additionally, with a fixed filling factor, height variations do not influence the appearance near the normal direction. Figure 10 shows similar manipulations using anisotropic micrograins, on top of opaque and transparent bulks.

Our model provides multiple adjustable parameters for each mi-




Figure 9: Variation of appearance by varying radii ratio and surface filling factor. An additional isotropic gold micrograin with $\beta = 2$ is added on the monodisperse case (top-middle image), which is a mixture of a purple diffuse reflectance bulk and silver micrograins ($\tau_0 = 0.3$ and $\beta = 0.2$). β corresponds to parameters a and d in our stretching matrix M as in [LRPB24] for isotropic grains. The diagrams at top right illustrate the effect of increasing radii ratio. The Δ LIP errors between each line show the visual difference between upper and lower images.

micrograin type deposited onto the microgeometry. This flexibility enables the representation of complex material appearance effects, such as retro-reflection from mixed materials on the silhouette of objects, color variations dependent on lighting and observation angles and multiple anisotropy direction, all illustrated in Figure 1 on the car bodywork. The properties of each micrograin category can also be controlled using textures to modulate their respective parameters. As demonstrated in Figure 11, the appearance of the car bodywork varies across its surface, where two textures are applied to adjust the surface filling factor of two micrograins types. This approach effectively simulates rusty and dirty appearance, layered on top of a clean, polished paint finish, represented by blue isotropic and golden flake micrograins deposited on an aluminium bulk.

Performance We evaluated our model on an NVIDIA GeForce RTX 3060 GPU. Table 1 reports the rendering times, which approximately follow a function of the number N of micrograin types: $T(N) \approx 6.03 \times N^{1.92}$. The results indicate an $\mathcal{O}(N^2)$ increase in rendering time with respect to the number of micrograin types N . This behavior is expected, as the evaluation of our model involves $\mathcal{O}(N^2)$ operations arising from the summation in Equations 16 and 17, as well as the pairwise products in Equation 11 and 12. Since the number of micrograin types used in practical applications is typically small (e.g., $N \leq 8$), the computational overhead remains acceptable.

# of micrograin types	2	4	6	8
Render time (s)	24.03	76.92	188.04	340.56

Table 1: Rendering times (in seconds) of the simple scene shown in Figure 9, with a varying number of plastic micrograin types. Rendered image size is 512×512 pixels with 2048 samples per pixel.

9. Discussion

We have introduced a novel polydisperse anisotropic BSDF model that generalizes and extends the capabilities of the micrograin framework [LRPB24]. We derive the exact NDF for microspheres composed of multiple micrograin types, each characterized by distinct shapes, orientations and sizes. Additionally, we extend the corresponding fully correlated GAF and an importance sampling procedure, enabling seamless integration into path-tracing renderers. Our model is validated through comparisons against reference simulations. We demonstrate how material appearance can be controlled by adjusting various parameters that govern distinct visual effects including retro-reflection from mixed materials, color mixtures with respect to lighting and observation angles, and multiple directions of anisotropy.

Limitations and future work. Our model inherits the current limitations of the micrograin formalism. Most importantly, it is limited to single scattering, and exhibits losses of energy as shown in the white furnace tests of Figure 12. Taking into account multiple scattering is more challenging in the case of polydispersity and using the fully-correlated GAF model. A recent solution has been introduced in the monodisperse case with the Smith




Figure 10: By modulating the filling factor of the larger micrograins, the visibility of the smaller grains can be controlled, depending on both incident and observations angles. First row : a diffuse bulk combined with an axis-aligned anisotropic silver micrograin (Lucas et al. [LRPB24]). A non-axis-aligned gold micrograin (polydisperse model) is subsequently added. Second row : a rough dielectric bulk incorporates greenish plastic micrograins (Lucas et al.) followed by the addition of the same gold micrograins. In both case, the filling factor is varied while maintaining a fixed height ratio of $\frac{r_2}{r_1} = 4.0$.

GAF [ZSL*25]. An alternative would be to extend monodisperse vNDF sampling [LPB25] to our polydisperse model and introduce it in a stochastic approach [HHdD16, Bd22].

An interesting avenue of future work would be to further generalize polydispersity by deriving the model using a continuous representation of micrograin shapes, sizes and orientations. As a first step, it might be possible to introduce continuous perturbations of shape, size and reflectance within a given micrograin type. We believe this would be necessary to represent the appearance of real-world materials, which are usually characterized by microscopic irregularities. In particular, we would like to explore the possibility of fitting micrograin distribution parameters directly to measured microstructures, with an additional segmentation step to identify the different types of micrograins that could be used to reproduce the appearance of a real-world surface.

Acknowledgments

This research was funded, in whole or in part, by the French National Research Agency (ANR) under the REFLECTIVITY project: ANR-22-CE22-0006.

References

- [AKDW22] ATANASOV A., KOYLAZOV V., DIMOV R., WILKIE A.: Microsurface transformations. *Computer Graphics Forum* 41, 4 (2022), 105–116. 11
- [ANA*20] ANDERSSON P., NILSSON J., AKENINE-MÖLLER T., OSKARSSON M., ÅSTRÖM K., FAIRCHILD M. D.: FLIP: A Difference Evaluator for Alternating Images. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 3, 2 (2020), 15:1–15:23. 7

- [APS00] ASHIKMIN M., PREMOŽE S., SHIRLEY P.: A microfacet-based brdf generator. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (2000), pp. 65–74. 2
- [Bd22] BITTERLI B., D'EON E.: A position-free path integral for homogeneous slabs and multiple scattering on smith microfacets. *Computer Graphics Forum* 41, 4 (2022), 93–104. 2, 9
- [BPV18] BARLA P., PACANOWSKI R., VANGORP P.: A composite brdf model for hazy gloss. *Computer Graphics Forum* 37, 4 (2018), 55–66. 2
- [BRMS20] BRINGIER B., RIBARDIÈRE M., MENEVEAUX D., SIMONOT L.: Design of rough microgeometries for numerical simulation of material appearance. *Appl. Opt.* 59, 16 (Jun 2020), 4856–4864. 6
- [CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. In *ACM SIGGRAPH proceedings* (1982). 2
- [dBWZ23] D'EON E., BITTERLI B., WEIDLICH A., ZELTNER T.: Microfacet theory for non-uniform heightfields. In *SIGGRAPH 2023 Conference Papers* (New York, NY, USA, 2023), Association for Computing Machinery. 2, 6, 7
- [DHd16] DUPUY J., HEITZ E., D'EON E.: Additional progress towards the unification of microfacet and microflake theories. In *EGSR (EI&I)* (2016), pp. 55–63. 2
- [DWMG16] DONG Z., WALTER B., MARSCHNER S., GREENBERG D. P.: Predicting appearance from measured microgeometry of metal surfaces. *ACM Trans. Graph.* 35, 1 (dec 2016). 2
- [EKM01] ERSHOV S., KOLCHIN K., MYSZKOWSKI K.: Rendering pearlescent appearance based on paint-composition modelling. *Computer Graphics Forum* 20, 3 (2001), 227–238. 2
- [GHZ18] GUO Y., HAŠAN M., ZHAO S.: Position-free monte carlo simulation for arbitrary layered bsdfs. *ACM Trans. Graph.* 37, 6 (2018). 2
- [HD15] HEITZ E., DUPUY J.: Implementing a simple anisotropic rough diffuse material with stochastic evaluation, 2015. 2
- [Hei14] HEITZ E.: Understanding the masking-shadowing function in microfacet-based brdfs. *Journal of Computer Graphics Techniques* 3, 2 (June 2014). 2



Figure 11: Micrograins parameters can be modulated using textures. From a clean version, rust and dust are added on the car bodywork with two textures which control the surface filling factor of copper (c) and diffuse (d) anisotropic micrograins, deposited atop a car paint composed of an aluminium bulk colored with blue isotropic particles (a) and anisotropic golden glitters (b).

(a) $\tau_{A,0} = 0.2$	(b) $\tau_{A,0} = 0.4$	(c) $\tau_{A,0} = 0.8$
$\tau_{B,0} = 0.15$	$\tau_{B,0} = 0.3$	$\tau_{B,0} = 0.6$
$\tau_{C,0} = 0.1$	$\tau_{C,0} = 0.2$	$\tau_{C,0} = 0.4$
$\tau_{D,0} = 0.2$	$\tau_{D,0} = 0.4$	$\tau_{D,0} = 0.8$




Figure 12: White furnace test for 3 configurations using 4 micrograin categories (A, B, C, D with same shape as in Figure 5) : A and B are perfect mirror materials, C is white plastic, and D is white diffuse, all deposited on a white diffuse bulk.

- [HHd16] HEITZ E., HANIKA J., D'EON E., DACHSBACHER C.: Multiple-scattering microfacet bsdfs with the smith model. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–14. [2](#), [9](#)
- [JDJM14] JAKOB W., D'EON E., JAKOB O., MARSCHNER S.: A comprehensive framework for rendering layered materials. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 33, 4 (July 2014), 118:1–118:14. [2](#)
- [JSR*22] JAKOB W., SPEIERER S., ROUSSEL N., NIMIER-DAVID M., VICINI D., ZELTNER T., NICOLET B., CRESPO M., LEROY V., ZHANG Z.: Mitsuba 3 renderer. <https://mitsuba-renderer.org/>, 2022. Version 3.7.1. [7](#)
- [KSK01] KELEMEN C., SZIRMAI-KALOS L.: A Microfacet Based Coupled Specular-Matte BRDF Model with Importance Sampling. In *Eurographics 2001 - Short Presentations* (2001), Eurographics Association. [2](#)
- [LJJ*18] LEE J. H., JARABO A., JEON D. S., GUTIERREZ D., KIM M. H.: Practical multiple scattering for rough surfaces. *ACM Trans. Graph.* 37, 6 (dec 2018). [2](#)
- [LPB25] LUCAS S., PACANOWSKI R., BARLA P.: Importance Sampling of the Micrograin Visible NDF. *Computer Graphics Forum* (2025). [2](#), [9](#)
- [LRPB23] LUCAS S., RIBARDIERE M., PACANOWSKI R., BARLA P.: A micrograin bsdf model for the rendering of porous layers. In *SIGGRAPH Asia 2023 Conference Papers* (New York, NY, USA, 2023), SA '23, Association for Computing Machinery. [2](#)
- [LRPB24] LUCAS S., RIBARDIERE M., PACANOWSKI R., BARLA P.: A Fully-correlated Anisotropic Micrograin BSDF Model. *ACM Transactions on Graphics* 43, 4 (July 2024), 111. [2](#), [3](#), [6](#), [7](#), [8](#), [9](#)
- [MBT*18] MENEVEAUX D., BRINGIER B., TAUZIA E., RIBARDIERE M., SIMONOT L.: Rendering rough opaque materials with interfaced lambertian microfacets. *IEEE Transactions on Visualization & Computer Graphics* 24, 03 (mar 2018), 1368–1380. [2](#)
- [MDG02] MÉRILLOU S., DISCHLER J.-M., GHAZANFARPOUR D.: A brdf postprocess to integrate porosity on rendered surfaces. *IEEE Transactions on Visualization and Computer Graphics* 6, 4 (2002), 306–318. [2](#)
- [MPG*16] MÜLLER T., PAPAS M., GROSS M., JAROSZ W., NOVÁK J.: Efficient rendering of heterogeneous polydisperse granular media. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 35, 6 (Dec. 2016), 168:1–168:14. [2](#)
- [MPH*15] MENG J., PAPAS M., HABEL R., DACHSBACHER C., MARSCHNER S., GROSS M., JAROSZ W.: Multi-scale modeling and rendering of granular materials. *ACM Trans. Graph.* 34, 4 (jul 2015). [2](#)
- [ON94] OREN M., NAYAR S. K.: Generalization of lambert's reflectance model. In *ACM SIGGRAPH proceedings* (1994). [2](#)
- [RBMS17] RIBARDIERE M., BRINGIER B., MENEVEAUX D., SIMONOT L.: STD: Student's t-Distribution of Slopes for Microfacet Based BSDFs. *Computer Graphics Forum* (2017). [2](#)
- [RDP05] ROSS V., DION D., POTVIN G.: Detailed analytical approach to the Gaussian surface bidirectional reflectance distribution function specular component applied to the sea surface. *J. Opt. Soc. Am. A* 22, 11 (Nov. 2005), 2442–2453. [2](#)
- [Smi67] SMITH B.: Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation* 15, 5 (September 1967), 668–671. [2](#)
- [TRM*25] TAMISIER E., RIBARDIERE M., MENEVEAUX D., HORNA S., POULIN P.: Visibility evaluation in microfacet theory. *IEEE Transactions on Visualization and Computer Graphics* 31, 2 (2025), 1422–1434. [2](#)
- [TS67] TORRANCE K. E., SPARROW E. M.: Theory for off-specular reflection from roughened surfaces*. *J. Opt. Soc. Am.* 57, 9 (Sep 1967), 1105–1114. [2](#)

[vGSK98] VAN GINNEKEN B., STAVRIDIS M., KOENDERINK J. J.: Diffuse and specular reflectance from rough surfaces. *Appl. Opt.* 37, 1 (Jan 1998), 130–139. 2

[WJHY22] WANG B., JIN W., HASAN M., YAN L.-Q.: Spongecake: A layered microflake surface appearance model. *ACM Trans. Graph.* 42, 1 (sep 2022). 2

[WMLT07] WALTER B., MARSCHNER S. R., LI H., TORRANCE K. E.: Microfacet models for refraction through rough surfaces. In *Computer Graphics Forum, EGSR proceedings* (2007). 2

[WW07] WEIDLICH A., WILKIE A.: Arbitrarily layered micro-facet surfaces. In *GRAPHITE 2007* (Dec. 2007), ACM, pp. 171–178. 2

[XH18] XIE F., HANRAHAN P.: Multiple scattering from distributions of specular v-grooves. *ACM Trans. Graph.* 37, 6 (dec 2018). 2

[ZSL*25] ZHOU F., SHEN H., LI M., ZHAO Y., BI C.: Position-free multiple-scattering computations for micrograin bsdf model. *Graphical Models* 141 (2025), 101288. 2, 9

Appendix A: Monodisperse non-axis-aligned anisotropic NDF

We derive the monodisperse NDF by first re-expressing it in terms of micronormal angles $\theta_{\mathbf{m}}$ and $\phi_{\mathbf{m}}$ and using Bayes' rule:

$$D(\mathbf{m}) = \frac{p(\theta_{\mathbf{m}}, \phi_{\mathbf{m}})}{\sin \theta_{\mathbf{m}} \cos \theta_{\mathbf{m}}} = \frac{p(\theta_{\mathbf{m}}|\phi_{\mathbf{m}})p(\phi_{\mathbf{m}})}{\sin \theta_{\mathbf{m}} \cos \theta_{\mathbf{m}}}, \quad (19)$$

where $p(\phi_{\mathbf{m}})$ is a marginal probability density function and $p(\theta_{\mathbf{m}}|\phi_{\mathbf{m}})$ is a conditional probability density function.

In the isotropic case, the marginal PDF is simply $p(\phi_{\mathbf{m}}) = \frac{1}{2\pi}$. To derive $p(\phi_{\mathbf{m}})$ in the anisotropic case, we consider the unit projected vector $\mathbf{u}_{\phi_{\mathbf{m}}}$ of a micrograin surface normal \mathbf{m} on the XY -plane: $\mathbf{u}_{\phi_{\mathbf{m}}} = \begin{pmatrix} \cos \phi_{\mathbf{m}} \\ \sin \phi_{\mathbf{m}} \\ 0 \end{pmatrix}$. We also define $\mathbf{u}'_{\phi_{\mathbf{m}}} = \begin{pmatrix} u'_x \\ u'_y \\ 0 \end{pmatrix} = M^T \mathbf{u}_{\phi_{\mathbf{m}}}$.

The mapping function $\phi'_{\mathbf{m}}(\phi_{\mathbf{m}})$ of $\phi_{\mathbf{m}}$ (in stretched space) to $\phi'_{\mathbf{m}}$ (in isotropic space) and its corresponding derivative are given by:

$$\phi'_{\mathbf{m}}(\phi_{\mathbf{m}}) = \arctan 2(u'_y, u'_x), \quad \frac{d\phi'_{\mathbf{m}}}{d\phi_{\mathbf{m}}} = \frac{|\det M|}{\|M^T \mathbf{u}_{\phi_{\mathbf{m}}}\|^2}. \quad (20)$$

The marginal PDF $p(\phi_{\mathbf{m}}) = p(\phi'_{\mathbf{m}}) \frac{d\phi'_{\mathbf{m}}}{d\phi_{\mathbf{m}}}$ is then given by:

$$p(\phi_{\mathbf{m}}) = \frac{|\det M|}{2\pi \|M^T \mathbf{u}_{\phi_{\mathbf{m}}}\|^2}. \quad (21)$$

For the conditional PDF $p(\theta_{\mathbf{m}}|\phi_{\mathbf{m}})$, we start by defining the related PDF of surface heights $p(h)$. The corresponding CDF $P(h)$ is a function of filling factors: $P(h) = 1 - \frac{\tau(h, \mathbf{n})}{\tau_0}$, with $\tau(h, \mathbf{n})$ the filling factor at height h . $p(h)$ is then obtained by differentiation:

$$p(h) = \frac{dP(h)}{dh} = \frac{2\rho\pi |\det M| h}{\tau_0} \exp\left(-\rho\pi |\det M| (r^2 - h^2)\right). \quad (22)$$

Knowing the normal \mathbf{m} on a micrograin, its corresponding height is given by $h(\mathbf{m}) = \left(\frac{M^T \mathbf{m}}{\|M^T \mathbf{m}\|} \cdot \mathbf{n}\right) r$. In terms of micronormal angles, the height and its derivative according to $\theta_{\mathbf{m}}$ are given by:

$$h(\mathbf{m}) = \frac{r}{\sqrt{t(\mathbf{m})^2 + 1}}, \quad \frac{\partial h(\mathbf{m})}{\partial \theta_{\mathbf{m}}} = \frac{t(\mathbf{m})^2 r}{\cos^2 \theta_{\mathbf{m}} (t(\mathbf{m})^2 + 1)^{\frac{3}{2}}}, \quad (23)$$

where we have used $t(\mathbf{m}) := \|M^T \mathbf{u}_{\phi_{\mathbf{m}}}\| \tan \theta_{\mathbf{m}}$. The conditional

PDF $p(\theta_{\mathbf{m}}|\phi_{\mathbf{m}}) = p(h(\mathbf{m})) \frac{\partial h(\mathbf{m})}{\partial \theta_{\mathbf{m}}}$ is then given by:

$$p(\theta_{\mathbf{m}}|\phi_{\mathbf{m}}) = \frac{2\rho\pi r^2 |\det M| \tan \theta_{\mathbf{m}} \|M^T \mathbf{u}_{\phi_{\mathbf{m}}}\|^2}{\tau_0 \cos^2 \theta_{\mathbf{m}} (t(\mathbf{m})^2 + 1)^2} \times \exp\left(-\rho\pi r^2 |\det M| \frac{t(\mathbf{m})^2}{t(\mathbf{m})^2 + 1}\right). \quad (24)$$

By putting Equations 21 and 24 in Equation 19, followed by a few simplifications (notice that $\|M^T \mathbf{m}\|^2 = (t(\mathbf{m})^2 + 1) \cos^2 \theta_{\mathbf{m}}$), we arrive at the sought-for monodisperse NDF:

$$D(\mathbf{m}) = \frac{|\det M|}{\|M^T \mathbf{m}\|^4} \frac{-\ln(1 - \tau_0)(1 - \tau_0)^{1 - \frac{\cos^2 \theta_{\mathbf{m}}}{\|M^T \mathbf{m}\|^2}}}{\pi \tau_0}. \quad (25)$$

By using the microsurface transformation introduced in [AKDW22], we have $1 - \frac{\cos^2 \theta_{\mathbf{m}}}{\|M^T \mathbf{m}\|^2} = \sin^2 \theta_{\mathbf{m}'}$ (Figure 13). Thus, Equation 25 can be rewritten as Equation 5.




Figure 13: (a) Surface transformation from stretched space (with normal \mathbf{m}) to spherical micrograin space (with normal \mathbf{m}'), using matrix M^{-1} . (b) Illustration of the relationship between normal polar angles in stretched space $\theta_{\mathbf{m}}$ and in spherical micrograin space $\theta_{\mathbf{m}'}$. Thales's theorem yields $\cos \theta_{\mathbf{m}'} = \frac{\cos \theta_{\mathbf{m}}}{\|M^T \mathbf{m}\|}$ [AKDW22].

Appendix B: Monodisperse non-axis-aligned anisotropic NDF importance sampling

We begin with the isotropic spherical NDF $D(\mathbf{m}')$ of Equation 3. The azimuthal marginal density is: $p(\phi_{\mathbf{m}'}) = \frac{1}{2\pi}$. The conditional density of the polar angle is:

$$p(\theta_{\mathbf{m}'}) = \frac{-\ln(1 - \tau_0)(1 - \tau_0)^{\sin^2 \theta_{\mathbf{m}'}}}{\tau_0} 2 \sin \theta_{\mathbf{m}'} \cos \theta_{\mathbf{m}'}. \quad (26)$$

Integrating gives the CDF:

$$P(\theta_{\mathbf{m}'}) = \int_0^{\theta_{\mathbf{m}'}} p(\theta|\phi_{\mathbf{m}'}) d\theta = \frac{1 - (1 - \tau_0)^{\sin^2 \theta_{\mathbf{m}'}}}{\tau_0}. \quad (27)$$

For ξ_ϕ and ξ_θ , two uniform random numbers in $[0, 1)$, the sampling is obtained by inversion:

$$\phi_{\mathbf{m}'} = 2\pi \xi_\phi, \quad \theta_{\mathbf{m}'} = \arcsin \sqrt{\frac{\ln(1 - \tau_0 \xi_\theta)}{\ln(1 - \tau_0)}}. \quad (28)$$

The sampled normal in the anisotropic domain is $\mathbf{m} = \frac{(M^{-1})^T \mathbf{m}'}{\|(M^{-1})^T \mathbf{m}'\|}$. Its corresponding density is:

$$p(\mathbf{m}) = D(\mathbf{m}) \cos \theta_{\mathbf{m}}. \quad (29)$$

Appendix C: Global filling factor

In the finite case, consider a randomly distributed surface consisting of K discrete elements. These elements are divided into N distinct categories. The number of elements in i -th category is given by $K_i = \omega_i K$ where ω_i represents the fraction of elements belonging to i -th category, i.e., $\sum_{i=1}^N \omega_i = 1$. The corresponding density of i -th category elements is $\rho_i = \omega_i \rho$. Let each element of i -th category has a projected area $\sigma_i(h, \mathbf{d})$ at height h and direction \mathbf{d} . For notation simplicity, we write $\sigma_i = \sigma_i(h, \mathbf{d})$. Then, the probability that a randomly chosen point on the surface lies outside a single element of i -th category is $p_{out,i} = 1 - \frac{\rho_i \sigma_i}{K_i}$. Accordingly, the probability that this point lies outside all K_i elements of i -th category is $(p_{out,i})^{K_i}$. The probability that the point lies inside at least one element—referred to as the finite filling factor $p_{in,K}$ is given by: $\tau_K = 1 - \prod_{i=1}^N (p_{out,i})^{K_i}$. We thus have the global filling factor $\bar{\tau}(h, \mathbf{d}) = \lim_{K \rightarrow \infty} p_{in,K}$.

$$\bar{\tau}(h, \mathbf{d}) = \lim_{K \rightarrow \infty} 1 - \prod_{i=1}^N \left(1 - \frac{\rho_i \sigma_i}{K}\right)^{\omega_i K} = 1 - \exp\left(-\rho \sum_{i=1}^N \omega_i \sigma_i\right) \quad (30)$$

Appendix D: Polydisperse non-axis-aligned anisotropic NDF

As in Appendix A, we derive the polydisperse NDF by first re-expressing it in terms of micronormal angles $\theta_{\mathbf{m}}$ and $\phi_{\mathbf{m}}$. However, in the polydisperse case, we must also account for the different micrograin types when applying Bayes' rule, yielding:

$$D(\mathbf{m}) = \frac{\sum_{i=1}^N \omega_i p(\theta_{\mathbf{m}}|\phi_{\mathbf{m}}, i) p(\phi_{\mathbf{m}}|i)}{\sin \theta_{\mathbf{m}} \cos \theta_{\mathbf{m}}} \quad (31)$$

The PDF $p(\phi_{\mathbf{m}}|i)$ is directly obtained by adapting Equation 21:

$$p(\phi_{\mathbf{m}}|i) = \frac{|\det M_i|}{2\pi \|M_i^T \mathbf{u}_{\phi_{\mathbf{m}}}\|^2} \quad (32)$$

For the PDF $p(\theta_{\mathbf{m}}|\phi_{\mathbf{m}}, i)$, we once again start by defining a PDF of surface heights $p(h|i)$, now restricted to micrograins of type i . Following a similar derivation, we arrive at:

$$p(h|i) = \frac{-\rho}{\bar{\tau}_0} \frac{d\sigma_i(h)}{dh} \exp\left(-\rho \sum_{j=1}^N \omega_j \sigma_j(h)\right), \quad (33)$$

where the area $\sigma_i(h) = \pi[r_i^2 - h^2]^+ |\det M_i|$ uses a clamping operator $[\cdot]^+$ to ensure that the i th micrograin makes no contribution to heights when $h \geq r_i$. Consequently, transforming $p(h|i)$ to $p(\theta_{\mathbf{m}}|\phi_{\mathbf{m}}, i)$ only requires the mapping $h_i = \left(\frac{M_i^T \mathbf{m}}{\|M_i^T \mathbf{m}\|} \cdot \mathbf{n}\right) r_i$:

$$p(\theta_{\mathbf{m}}|\phi_{\mathbf{m}}, i) = \frac{-\rho}{\bar{\tau}_0} \frac{d\sigma_i(h_i(\mathbf{m}))}{dh} \frac{dh_i(\mathbf{m})}{d\theta_{\mathbf{m}}} \exp\left(-\rho \sum_{j=1}^N \omega_j \sigma_j(h_i(\mathbf{m}))\right).$$

Developing h_i , σ_i and their derivatives, we obtain:

$$p(\theta_{\mathbf{m}}|\phi_{\mathbf{m}}, i) = \frac{2\rho\pi r_i^2 |\det M_i| \tan \theta_{\mathbf{m}} \|M_i^T \mathbf{u}_{\phi_{\mathbf{m}}}\|^2}{\bar{\tau}_0 \cos^2 \theta_{\mathbf{m}} (t_i(\mathbf{m}) + 1)^2} \times \exp\left(-\rho\pi \sum_{j=1}^N \omega_j |\det M_j| \left[r_j^2 - \frac{r_i^2}{t_i(\mathbf{m})^2 + 1}\right]^+\right), \quad (34)$$

which is of a form similar to Equation 24 (here with $t_i(\mathbf{m}) := \|M_i^T \mathbf{u}_{\phi_{\mathbf{m}}}\| \tan \theta_{\mathbf{m}}$), with the difference that the second line now accounts for intersections among different types of micrograins.

We arrive at the sought-for polydisperse NDF (Equation 11) by putting Equations 32 and 34 in Equation 31, once again followed by a few straightforward simplifications as $\|M_i^T \mathbf{m}\|^2 = (t_i(\mathbf{m})^2 + 1) \cos^2 \theta_{\mathbf{m}}$ and $\cos \theta_{\mathbf{m}'} = \frac{\cos \theta_{\mathbf{m}}}{\|M_i^T \mathbf{m}\|}$.

Appendix E: Importance sampling of the polydisperse NDF

As pointed out in Section 6, the polydisperse NDF exhibits C^1 discontinuities which hampers importance sampling. We circumvent this issue by decomposing heights in N slices S_k , and reorganize the sampling process in four main steps.

The first step is to sample a slice S_k , which requires a PDF $p(S_k)$. We start from the distribution of heights $p(h) = \sum_i \omega_i p(h|i)$. Using Equation 33, we obtain:

$$p(h) = \frac{-\rho}{\bar{\tau}_0} \left(\sum_{i=1}^N \omega_i \frac{d\sigma_i(h)}{dh}\right) \exp\left(-\rho \sum_{j=1}^N \omega_j \sigma_j(h)\right).$$

If we now restrict heights to the slice S_k and order micrograin types with increasing radii, the sum only has to go from k to N . As a result, the area term does not need to rely on clamping: $\sigma_j(h) = \pi(r_j^2 - h^2) |\det M_j|$. Assuming $h \in [r_{k-1}, r_k]$ in S_k , we obtain:

$$p(h, S_k) = \frac{2\rho\pi h}{\bar{\tau}_0} \left(\sum_{i=k}^N \omega_i |\det M_i|\right) \frac{\exp\left(\rho\pi \sum_{j=k}^N \omega_j |\det M_j| h^2\right)}{\exp\left(\rho\pi \sum_{l=k}^N \omega_l |\det M_l| r_l^2\right)}.$$

The sought-for probability $p(S_k)$ is then obtained by integration:

$$\begin{aligned} p(S_k) &= \int_{r_{k-1}}^{r_k} p(h, S_k) dh \\ &= \frac{\exp(\rho\pi r_k^2 \sum_{j=k}^N \omega_j |\det M_j|) - \exp(\rho\pi r_{k-1}^2 \sum_{j=k}^N \omega_j |\det M_j|)}{\bar{\tau}_0 \exp(\rho\pi \sum_{l=k}^N \omega_l r_l^2 |\det M_l|)} \\ &= \frac{\prod_{j=k}^N (1 - \tau_{j,0})^{1 - \frac{r_k^2}{r_j^2}} - \prod_{j=k}^N (1 - \tau_{j,0})^{1 - \frac{r_{k-1}^2}{r_j^2}}}{\bar{\tau}_0} \end{aligned}$$

where we have used Equation 10 to obtain the third line above.

Using $\lambda_k = \prod_{j=k}^N (1 - \tau_{j,0})^{-\frac{1}{r_j^2}}$ yields a more compact formulation:

$$p(S_k) = \frac{1}{\bar{\tau}_0} \left(\lambda_k^2 - \lambda_{k-1}^2\right) \prod_{j=k}^N (1 - \tau_{j,0}). \quad (35)$$

Observe that if $r_k = r_{k-1}$, then $p(S_k) = 0$, hence the slice vanished between two successive micrograin types of the same radii.

The second step consists in sampling a micrograin type i appearing inside the chosen slice S_k , which requires a probability $p(i|S_k)$. Using Bayes' rule, we have:

$$p(i|S_k) = \frac{p(i, S_k)}{p(S_k)} = \frac{\int_{r_{k-1}}^{r_k} p(i, h', S_k) dh'}{p(S_k)}. \quad (36)$$

Writing $p(i, h, S_k) = \frac{\omega_i |\det M_i|}{\sum_{j=k}^N \omega_j |\det M_j|} p(h, S_k)$ and observing that

$\int_{r_{k-1}}^{r_k} p(h, S_k) dh = P(S_k)$, we obtain (assuming $k \leq i \leq N$):

$$p(i|S_k) = \frac{\omega_i |\det M_i|}{\sum_{l=k}^N \omega_l |\det M_l|} = \frac{1}{\ln(\lambda_k)} \ln \left((1 - \tau_{i,0})^{-\frac{1}{r_i}} \right). \quad (37)$$

The third step is to sample the specular component or diffuse component based on the selected micrograin type i in slice S_k , we define the probability of sampling the diffuse component as:

$$p_i^d(\mathbf{i}) = \frac{K_i^d(\mathbf{n})}{K_i^d(\mathbf{n}) + F_i(\mathbf{i} \cdot \mathbf{n}, \mathbf{n})}. \quad (38)$$

The probability of sampling the specular component is $1 - p_i^d(\mathbf{i})$.

The fourth and last step consists in sampling a micronormal \mathbf{m} on a micrograin of type i in a slice S_k for the specular component, which requires a PDF $p(\mathbf{m}|i, S_k)$. Expressing this PDF in terms of micronormal angles, we have:

$$p(\mathbf{m}|i, S_k) = \frac{p(\theta_{\mathbf{m}}|\phi_{\mathbf{m}}, i, S_k) p(\phi_{\mathbf{m}}|i, S_k)}{\sin \theta_{\mathbf{m}}}. \quad (39)$$

The PDF $p(\phi_{\mathbf{m}}|i, S_k)$ is the same as $p(\phi_{\mathbf{m}}|i)$ (Equation 32) since $\phi_{\mathbf{m}}$ does not depend on height, and is thus independent of the choice of slice S_k . To derive the PDF $p(\theta_{\mathbf{m}}|\phi_{\mathbf{m}}, i, S_k)$ we begin by considering the conditional PDF of heights $p(h|i, S_k) = \frac{p(i, h, S_k)}{p(i, S_k)}$. If we write its denominator $p(i, S_k) = \frac{\omega_i |\det M_i|}{\sum_{l=k}^N \omega_l |\det M_l|} p(S_k)$, we directly obtain $p(h|i, S_k) = \frac{p(h, S_k)}{p(i, S_k)}$. Now using $p(\theta_{\mathbf{m}}|\phi_{\mathbf{m}}, i, S_k) = p(h|i, S_k) \frac{\partial h(\mathbf{m})}{\partial \theta_{\mathbf{m}}}$ in conjunction with Equation 23 and after several straightforward simplifications, we obtain:

$$p(\theta_{\mathbf{m}}|\phi_{\mathbf{m}}, i, S_k) = \frac{2r_i^2 t_i(\mathbf{m})^2}{\cos \theta_{\mathbf{m}} \sin \theta_{\mathbf{m}} (t_i(\mathbf{m})^2 + 1)^2} \frac{\ln(\lambda_k) \lambda_k^{\frac{r_i^2}{t_i(\mathbf{m})^2 + 1}}}{\lambda_k^{r_k^2} - \lambda_k^{r_{k-1}^2}}. \quad (40)$$

Putting it all together, we have:

$$p(\mathbf{m}|i, S_k) = \frac{|\det M_i|}{\|M_i^T \mathbf{m}\|^4} \frac{r_i^2 \ln(\lambda_k) \lambda_k^{r_i^2 \cos^2 \theta_{\mathbf{m}'}}}{\pi (\lambda_k^{r_k^2} - \lambda_k^{r_{k-1}^2})} \cos \theta_{\mathbf{m}} \quad (41)$$

where $\cos \theta_{\mathbf{m}'} = \frac{\cos \theta_{\mathbf{m}}}{\|M_i^T \mathbf{m}\|}$. Also, note that when there is only one type of grain (i.e., $r_i = r_k = r, r_{k-1} = 0, \lambda_k = (1 - \tau_0)^{-\frac{1}{r^2}}$), the PDF above reduces to Equation 29 as expected.

As in the monodisperse NDF sampling procedure (Section B), we first draw sphere-space micro-normals \mathbf{m}' for grain type i , and subsequently apply the stretching transform

$$\mathbf{m} = \frac{(M_i^{-1})^T \mathbf{m}'}{\|(M_i^{-1})^T \mathbf{m}'\|}.$$

We define two arbitrary random values ξ_θ and ξ_ϕ uniformly distributed in $[0, 1)$ to sample $\theta_{\mathbf{m}'}$ and $\phi_{\mathbf{m}'}$. By assuming the stretching matrix M_i to be the identity in Equation 40, we obtain the condi-

tional PDF of the polar angle in sphere space:

$$p(\theta_{\mathbf{m}'}|\phi_{\mathbf{m}'}, i, S_k) = 2r_i^2 \sin \theta_{\mathbf{m}'} \cos \theta_{\mathbf{m}'} \frac{\ln(\lambda_k) \lambda_k^{r_i^2 \cos^2 \theta_{\mathbf{m}'}}}{\lambda_k^{r_k^2} - \lambda_k^{r_{k-1}^2}}. \quad (42)$$

Then we calculate its corresponding CDF by integrating over $\theta_{\mathbf{m}'}$ in the corresponding height range $\theta_{\min} = \arccos\left(\frac{r_k}{r_i}\right)$, $\theta_{\max} = \arccos\left(\frac{r_{k-1}}{r_i}\right)$ in the current slice S_k as:

$$P(\theta_{\mathbf{m}'}|\phi_{\mathbf{m}'}, i, S_k) = \int_{\theta_{\min}}^{\theta_{\max}} p(\theta|\phi_{\mathbf{m}'}, i, S_k) d\theta = \frac{\lambda_k^{r_k^2} - \lambda_k^{r_i^2 \cos^2 \theta_{\mathbf{m}'}}}{\lambda_k^{r_k^2} - \lambda_k^{r_{k-1}^2}}. \quad (43)$$

When we reduce 32 to isotropic case, we obtain $p(\phi_{\mathbf{m}'}|i, S_k) = \frac{1}{2\pi}$. By inverting Equation 43, the sampling formula for $\phi_{\mathbf{m}'}$ and also for $\theta_{\mathbf{m}'}$ is:

$$\phi_{\mathbf{m}'}(i) = 2\pi \xi_\phi$$

$$\theta_{\mathbf{m}'}(i, S_k) = \arccos \left(\sqrt{\frac{\ln \left((1 - \xi_\theta) \lambda_k^{r_k^2} + \xi_\theta \lambda_k^{r_i^2} \right)}{r_i^2 \ln(\lambda_k)}} \right) \quad (44)$$

Note that when there is only one type of grain (i.e., $r_i = r_k = r, r_{k-1} = 0, \lambda_k = (1 - \tau_0)^{-\frac{1}{r^2}}$), the above sampling formulas reduce to Equations 28 as expected.