



HAL
open science

A Fully-correlated Anisotropic Micrograin BSDF Model

Simon Lucas, Mickaël Ribardière, Romain Pacanowski, Pascal Barla

► **To cite this version:**

Simon Lucas, Mickaël Ribardière, Romain Pacanowski, Pascal Barla. A Fully-correlated Anisotropic Micrograin BSDF Model. ACM Transactions on Graphics, 2024, 43 (4), pp.111. 10.1145/3658224 . hal-04567402

HAL Id: hal-04567402

<https://hal.science/hal-04567402>

Submitted on 3 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Fully-correlated Anisotropic Micrograin BSDF Model

SIMON LUCAS, Université de Bordeaux, INRIA, France
MICKAËL RIBARDIÈRE, Université de Poitiers, France
ROMAIN PACANOWSKI, INRIA, France
PASCAL BARLA, INRIA, France




Fig. 1. Several instances of our material model rendered in Mitsuba [Jakob et al. 2022] (base surface materials shown in the bottom right corner): (a) Anisotropic aluminum micrograins on top of an isotropic rough gold conductor base; (b) Grey diffuse micrograins on top of a diffuse wood texture, conveying a dusty appearance; (c) Green diffuse+specular micrograins, mimicking bacterial deposit on top of a rough dielectric surface; (d) Height-correlated colored diffuse micrograins on a rough aluminium conductor base; (e) Gray diffuse micrograins located in small cavities of white plastic weaves to portray a dirty material.

We introduce an improved version of the micrograin BSDF model [Lucas et al. 2023] for the rendering of anisotropic porous layers. Our approach leverages the properties of micrograins to take into account the correlation between their height and normal, as well as the correlation between the light and view directions. This allows us to derive an exact analytical expression for the Geometrical Attenuation Factor (GAF), summarizing shadowing and masking inside the porous layer. This fully-correlated GAF is then used to define appropriate mixing weights to blend the BSDFs of the porous and base layers. Furthermore, by generalizing the micrograins shape to anisotropy, combined with their fully-correlated GAF, our improved BSDF model produces effects specific to porous layers such as retro-reflection visible on dust layers at grazing angles or height and color correlation that can be found on rusty materials. Finally, we demonstrate very close matches

Authors' addresses: Simon Lucas, Université de Bordeaux, INRIA, Bordeaux, France, simon.lucas@inria.fr; Mickaël Ribardière, Université de Poitiers, Poitiers, France, mickaël.ribardière@univ-poitiers.fr; Romain Pacanowski, INRIA, Bordeaux, France, romain.pacanowski@inria.fr; Pascal Barla, INRIA, Bordeaux, France, pascal.barla@inria.fr.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

© 2024 Association for Computing Machinery.
XXXX-XXXX/2024/7-ART111 \$15.00
<https://doi.org/10.1145/3658224>

between our BSDF model and light transport simulations realized with explicit instances of micrograins, thus validating our model.

CCS Concepts: • **Computing methodologies** → **Reflectance modeling**.

Additional Key Words and Phrases: BSDF, porosity, correlation, visibility

ACM Reference Format:

Simon Lucas, Mickaël Ribardière, Romain Pacanowski, and Pascal Barla. 2024. A Fully-correlated Anisotropic Micrograin BSDF Model. 43, 4, Article 111 (July 2024), 14 pages. <https://doi.org/10.1145/3658224>

1 INTRODUCTION

The microfacet theory [Cook and Torrance 1982; Torrance and Sparrow 1967] has been used for more than five decades in the construction of surfacic material models, or Bidirectional Scattering Distribution Functions (BSDF). As recalled in Section 2, a variety of effects has been reproduced, including anisotropy, glints, iridescence, sheen or hazy gloss. However, the vast majority of microfacet-based models relies on a common assumption, first introduced by Smith [1967]: the height of a given microfacet is uncorrelated from its normal. Even though this might be a valid assumption for random micro-surfaces, it cannot accommodate more structured configurations where a surface is covered with repetitions of similar microscopic elements, such as found in dusty or dirty surfaces.

Recent work (e.g., [d'Eon et al. 2023; Lucas et al. 2023]) has extended microfacet theory to deal with the mixing of multiple microfacet Normal Distribution Functions (NDF), which provides for

a way to model structured configurations as commonly found in porous layers. Such layered structures are different from standard layered materials, where several (potentially rough) interfaces are stacked upon each others. Indeed, porous layers come with new visibility effects at grazing angles. Even though it is likely that height-normal correlations should play an important role in such orientation-dependent visibility effects, to the best of our knowledge there is no existing work that breaks free of the Smith assumption.

Our goal is thus to explore the new appearance effects that may be obtained by explicitly handling the height-normal correlation arising from distributions of (potentially anisotropic) similar microscopic elements. To this end, we rely on the micrograin formalism introduced by Lucas et al. [2023] and recalled in Section 3. However, instead of modeling the Geometric Attenuation Factor (GAF) using Smith’s assumption, we leverage the micrograin structure to explicitly consider the bijection between heights and normals.

As explained in Section 4, our method is based on a fundamental property of random distributions of identical micrograins: the visibility of the whole distribution from an arbitrary direction is determined by the projected area of a single micrograin. We rely on this property to derive all the terms of our new BSDF model, making the following contributions:

- We extend the NDF in the model of Lucas et al. [2023] to deal with anisotropic micrograins (half-ellipsoids) using their projected area at normal incidence (Section 4.1);
- We derive an exact height-dependent GAF by relying on the projected area of micrograins at a given height, which we extend to a fully-correlated GAF by considering correlation among incoming and outgoing directions (Section 4.2);
- We rely on the newly introduced GAF to derive accurate mixing weights between the porous and base layers, and a new diffuse micrograin component (Section 4.3).

Our new model is thoroughly validated against reference simulations on explicit micrograin distributions, and compared to existing models for porous layers in Section 5. Validations show an excellent (near-perfect) fit to simulated BSDFs when the latter are restricted to single-scattering light paths. We showcase in Section 6 the new optical effects made possible by our fully-correlated microfacet model, some of which are shown in Figure 1. In particular, we demonstrate grazing-angle changes in visibility, retro-reflection effects and height-color correlation results. We discuss the limitations of our model (notably multiple scattering) in Section 7, and argue for the micrograin formalism as a promising foundation for future research in the development of material models.

2 PREVIOUS WORK

Microfacet theory. Microfacet theory [Torrance and Sparrow 1967] was first introduced to the Computer Graphics community by Cook and Torrance [1982] in order to model rough microspheres. Most microfacet models consider mirror-like microfacets distributed stochastically on the surface, according to their Normal Distribution Function (NDF). Existing models may differ depending on the nature of the NDF (e.g., [Dong et al. 2016; Ribardière et al. 2017; Walter et al. 2007]) and whether they handle its anisotropy (e.g., [Heitz 2014]). The theory has also been applied to diffuse (Lambertian) microfacets

(e.g., [Heitz and Dupuy 2015; Oren and Nayar 1994]), has been coupled with wave optics theories to render complex color effects such as iridescence [Belcour and Barla 2017] or diffraction [Holzschuch and Pacanowski 2017], and has been used empirically to reproduce sheen effects as found in fabric (e.g. [Estevez and Kulla 2017; Zeltner et al. 2022]). Closer to our approach, specific regular microstructures have been used to define microfacet distributions in previous work (e.g., cloth fibers [Sadeghi et al. 2013], feathers [Huang et al. 2022]), and explicit micro-geometries have been used to instantiate BRDF models in a bi-scale fashion [Wu et al. 2011]. Multiple scattering among microfacets has been addressed in different ways (e.g., [Bitterli and d’Eon 2022; Heitz et al. 2016a]).

Shadowing and masking. In microfacet theory, single-scattering refers to the set of light paths that intersect a single microfacet. From a statistical point of view, this represents the probability that a microfacet is visible from both the view (no masking) and light directions (no shadowing). This is modeled through the so-called Geometrical Attenuation Factor (GAF), which is an essential ingredient of physically-plausible BRDFs based on microfacet theory [Heitz 2014]. In the work of Cook and Torrance [1982], microfacets are considered to form V-groove cavity profiles [Torrance and Sparrow 1967]. Even if V-cavity masking functions are mathematically well-defined, they do not represent realistic microspheres.

A more realistic treatment is provided by Smith’s model [Smith 1967], which is equivalent to the formulation of Ashikhmin et al. [2000]. Since surfaces are represented with statistical distributions, masking is derived by exploiting two fundamental assumptions of the model: a microfacet orientation is not correlated to its position or height, and microfacets have all the same area [Heitz et al. 2016b]. The Smith GAF has been the subject of several studies in the field of computer graphics. It was first used in a separable form, where masking and shadowing are considered as directionally-independent [Walter et al. 2007]. Ross et al. [2005] partially addressed this limitation and introduced the joint masking-shadowing function that is commonly used today, taking into account the correlation between masking and shadowing due to the height of the microsurface. Another source of correlation lies in the incoming and outgoing directions as they get closer to each other. Existing solutions to handle it [Ashikhmin et al. 2000; Heitz et al. 2013a; van Ginneken et al. 1998] remain limited, as they are either based on empirical terms or limited to isotropic Gaussian statistics.

Material mixtures. Mixing microfacet-based models with other models is commonplace. For instance, laying dielectric microfacets onto a diffuse Lambertian base models plastic materials faithfully. However, it is important to take energy conservation considerations into account when doing so (e.g., [Kelemen and Szirmay-Kalos 2001; Meneveaux et al. 2018]). More complex appearance may be produced by layering several rough interfaces on top of each other, each interface potentially modeled by microfacet theory (e.g., [Guo et al. 2018; Jakob et al. 2014; Weidlich and Wilkie 2007]).

Instead of mixing materials vertically, some methods have proposed to use lateral mixtures, which are well adapted to the modeling of porous microspheres. Merillou et al. [2000] were the first to propose an empirical modification to microfacet models to handle porosity. They model pores as cylindrical holes with a diffuse

(Lambertian) reflectance, and provide an approximation for multiple scattering inside them. More importantly, they mix the base surface and the pores using an empirical weight factor based on the GAF and on the porosity parameter. Barla et al. [2018] proposed to use a mixture of two microfacet distributions, along with a reparametrization that provides intuitive control over hazy gloss effects. They suggest using a modified GAF to handle the masking and shadowing of the combined distributions. d'Eon et al. [2023] extend this approach to asymmetric mixtures of distributions based on height, exhibiting visibility effects at grazing angles. Interestingly, very similar results are obtained when considering that the microsurface consists of a distribution of elliptical micrograins onto a base surface, as demonstrated by Lucas et al. [2023]. A last approach relies on volumetric descriptions of the material, whereby microfacets are replaced by microflakes (e.g., [Dupuy et al. 2016; Wang et al. 2022]). However, they do not seem to produce grazing-angle visibility effects as is the case of surfacic models.

In contrast to previous methods that all rely on Smith's assumption, we explicitly consider height-normal dependencies at the microscopic level. This is done by extending the micrograin framework of Lucas et al. [2023], which we recall in the next section.

3 BACKGROUND

In the model of Lucas et al. [2023], the microsurface is assumed to consist of a uniform random distribution of identical opaque micrograins, which may potentially intersect (see Figure 2a,b). Micrograins are half-ellipsoids compressed vertically by a factor β .

Lucas et al. model the BSDF of such a microstructure as:

$$f(\mathbf{i}, \mathbf{o}) = w^+(\mathbf{i}, \mathbf{o})f^S(\mathbf{i}, \mathbf{o}) + (1 - w^+(\mathbf{i}, \mathbf{o}))f^B(\mathbf{i}, \mathbf{o}), \quad (1)$$

where \mathbf{i} and \mathbf{o} are the ingoing and outgoing directions respectively, and w^+ is a weight factor used to balance between the contributions of the surfacic (micrograin) material f^S and the bulk material f^B . While f^B is arbitrary, f^S is based on microfacet theory with the addition of a coupled Lambertian component, to accommodate for diffuse, diffuse+specular, or conductor micrograins:

$$f^S(\mathbf{i}, \mathbf{o}) = \frac{K_d}{\pi} \bar{T}(\mathbf{i}, \mathbf{o}) + \frac{D(\mathbf{h})G(\mathbf{i}, \mathbf{o}, \mathbf{h})F(\mathbf{i} \cdot \mathbf{h})}{4|\mathbf{i} \cdot \mathbf{n}||\mathbf{o} \cdot \mathbf{n}|}, \quad (2)$$

with \bar{T} an average transmission term used for coupling, $\mathbf{h} = \frac{\mathbf{i}+\mathbf{o}}{\|\mathbf{i}+\mathbf{o}\|}$ the halfway vector, \mathbf{n} the geometric normal, D and G the micrograin-based Normal Distribution Function (NDF) and Geometric Attenuation Factor (GAF) terms, and F the Fresnel term.

Lucas et al. rely on first-order statistics to derive a probability distribution function (PDF) of closest micrograins, from which they compute the new terms involved in their model. They first show that the corresponding cumulative distribution function (CDF) yields the filling factor τ_0 , which characterizes the relative area of micrograins when projected at normal incidence:

$$\tau_0 = 1 - e^{-\pi\rho r_s^2}, \quad (3)$$

where ρ is the density of micrograins and r_s is their radius. It is extended to model the visible filling factor $\tau_\beta^+(\boldsymbol{\omega})$, which characterizes the relative area of micrograins vertically compressed by a factor

β and projected along an arbitrary direction $\boldsymbol{\omega}$. This is used in the definition of the weight factor of Equation 1:

$$w^+(\mathbf{i}, \mathbf{o}) = 1 - \frac{(1 - \tau_\beta^+(\mathbf{i}))(1 - \tau_\beta^+(\mathbf{o}))}{1 - \tau_0} \quad (4)$$

Lucas et al. derive the NDF of the micrograin surface from the PDF of closest micrograins using a change of variable, yielding:

$$D(\theta_m) = -\frac{\beta^2 \ln(1 - \tau_0)(1 - \tau_0)^{\frac{\tan^2 \theta_m}{\beta^2 + \tan^2 \theta_m}}}{\tau_0(\beta^2 + \tan^2 \theta_m)^2 \cos^4 \theta_m}, \quad (5)$$

where \mathbf{m} is a microfacet normal, and $\theta_m = \arccos(\mathbf{m} \cdot \mathbf{n})$. They derive and invert the following CDF for importance sampling:

$$C(\theta_m) = \frac{1 - (1 - \tau_0)^{\frac{\tan^2 \theta_m}{\beta^2 + \tan^2 \theta_m}}}{\tau_0}. \quad (6)$$

The analytical GAF term G is then derived following Smith's uncorrelated height-normal assumption [Smith 1967].

4 METHOD

Our BSDF model relies on a microstructure similar to the one of Lucas et al.: a random distribution of identical opaque micrograins. We have found that such a microstructure has a fundamental property, which we rely on throughout this section: the visible filling factor τ (from an arbitrary direction) of the *whole* distribution derives from the projected area σ of a *single* micrograin (of arbitrary shape) onto the surface plane. Formally, we write:

$$\tau = 1 - e^{-\rho\sigma}. \quad (7)$$

PROOF. For a random distribution of identical micrograins, as shown in Figure 2, the visible filling factor τ of the microstructure from an arbitrary direction is given by the relative area of projected micrograins. In the finite case where we have a random distribution of N micrograins, the probability of a point on the projected surface to be *outside* of one projected micrograin is $p_{out} = 1 - \frac{\rho\sigma}{N}$, with ρ the distribution density and σ the projected area of a single micrograin. The visible filling factor in the finite case τ_N is then equal to the probability of a surface point to be inside at least one micrograin, which happens when that point is *not* outside of all N projected micrograins: $\tau_N = 1 - p_{out}^N$. The visible filling factor is finally obtained by taking the limit of τ_N as N tends toward infinity:

$$\tau = \lim_{N \rightarrow \infty} 1 - \left(1 - \frac{\rho\sigma}{N}\right)^N = 1 - e^{-\rho\sigma}. \quad \square$$

For instance, the filling factor τ_0 in the model of Lucas et al. (Equation 3) is obtained by using the projected area of a disc at normal incidence $\sigma_0 = \pi r_s^2$ in Equation 7.

We first use this finding in Section 4.1 to handle anisotropic micrograins, and derive an anisotropic NDF by demonstrating that Equation 5 is shape-invariant. We then consider projected areas of micrograins at oblique incidences in Section 4.2, which permits the derivation of a height- and direction-correlated GAF, exhibiting retro-reflection effects. We finally tie all components of our BSDF model together in Section 4.3, where we use the GAF to derive both the exact diffuse micrograin component and pore visibility.




Fig. 2. The surface microstructure is assumed to consist of a distribution of micrograins, as shown in a) for the isotropic case using a top view, where the relative area of the union of micrograin bases corresponds to the filling factor τ_0 . In the work of Lucas et al. [2023], b) micrograins are half-ellipsoids compressed vertically by a factor β . In contrast, c) we use half-ellipsoids stretched horizontally in the lateral directions by a pair of factors $\frac{1}{\beta_x}$ and $\frac{1}{\beta_y}$.

4.1 Anisotropic NDF

A first difference with the model of Lucas et al. [2023] lies in the way we model elliptical micrograins: instead of compressing spherical micrograins vertically, we stretch them laterally (see Figure 2c). This allows for different amounts of stretching $\frac{1}{\beta_x}$ and $\frac{1}{\beta_y}$ in different directions (with $\beta_x \in \mathbb{R}^+$ and $\beta_y \in \mathbb{R}^+$).

The formula of the filling factor τ_0 needs to be modified to account for this anisotropic micrograin geometry. Using Equation 7 in conjunction with the projected area of elliptical micrograins at normal incidence $\sigma_0 = \frac{\pi r_s^2}{\beta_x \beta_y}$, we obtain:

$$\tau_0 = 1 - e^{-\rho \frac{\pi r_s^2}{\beta_x \beta_y}}. \quad (8)$$

Compared to Equation 3, the filling factor is now dependent on the stretching factors $\frac{1}{\beta_x}$ and $\frac{1}{\beta_y}$. Intuitively, when stretching increases, the micrograin density and/or radius must decrease to leave the filling factor unchanged. Formally, inverting Equation 8 yields:

$$\rho r_s^2 = -\frac{\ln(1 - \tau_0) \beta_x \beta_y}{\pi}. \quad (9)$$

Since the density ρ and squared radius r_s^2 are always multiplied together, we consider without loss of generality that $r_s = 1$ in the rest of the paper, which simplifies mathematical derivations.

In the isotropic case (i.e., $\beta_x = \beta_y = \beta$), the NDF of stretched micrograins is identical to that of Lucas et al. who use compressed micrograins (due to NDF normalization). Using Equation 9, we explicitly re-write Equation 5 in a shape-invariant form [Heitz 2014]:

$$D(\mathbf{m}) = \frac{1}{\cos^4 \theta_m} \frac{1}{\beta^2} f(t(\mathbf{m})), \quad (10)$$

$$f(t(\mathbf{m})) = -\frac{\ln(1 - \tau_0)(1 - \tau_0)^{\frac{t(\mathbf{m})^2}{1+t(\mathbf{m})^2}}}{\tau_0 \pi (1 + t(\mathbf{m})^2)^2}, \quad (11)$$

where $t(\mathbf{m}) = \frac{\tan \theta_m}{\beta}$ in the isotropic case. The extension to an anisotropic NDF (see Supplemental material) is obtained by setting

$$t(\mathbf{m}) = \tan \theta_m \sqrt{\frac{\cos^2 \phi_m}{\beta_x^2} + \frac{\sin^2 \phi_m}{\beta_y^2}} \text{ in Equation 10.}$$

As in the model of Lucas et al., the resulting NDF may be expressed in terms of the anisotropic Trowbridge-Reitz (GGX) NDF:

$$D(\mathbf{m}) = -\frac{\ln(1 - \tau_0)(1 - \tau_0)^{\frac{t(\mathbf{m})^2}{1+t(\mathbf{m})^2}}}{\tau_0} D_{ggx}^{\text{aniso}}(\mathbf{m}). \quad (12)$$

Importance sampling. We first rewrite the CDF in the isotropic case (Equation 6, where $\beta_x = \beta_y = \beta$) using $t(\mathbf{m}) = \frac{\tan \theta_m}{\beta}$:

$$C(\mathbf{m}) = \frac{1 - (1 - \tau_0)^{\frac{t(\mathbf{m})^2}{1+t(\mathbf{m})^2}}}{\tau_0}. \quad (13)$$

As before, this is trivially extended to the anisotropic case by using $t^2(\mathbf{m}) = \tan^2 \theta_m g(\phi_m)$ where $g(\phi_m) = \frac{\cos^2 \phi_m}{\beta_x^2} + \frac{\sin^2 \phi_m}{\beta_y^2}$. The angles (θ_m, ϕ_m) of the microsurface normal \mathbf{m} are then given by:

$$\theta_m = \arctan \left(\sqrt{\frac{1}{g(\phi_m)} \frac{\ln(1 - \xi_\theta \tau_0)}{\ln(1 - \xi_\theta \tau_0) - \ln(1 - \tau_0)}} \right), \quad (14)$$

$$\phi_m = \arctan \left(\frac{\beta_y \sin(2\pi \xi_\phi)}{\beta_x \cos(2\pi \xi_\phi)} \right), \quad (15)$$

where ξ_θ and ξ_ϕ are univariate random variables. Here, ϕ_m is obtained by inverting the CDF corresponding to the marginal density function as commonly done for GGX¹, whereas θ_m is obtained by inverting the CDF corresponding to the conditional density function (i.e., by inverting Equation 13).

4.2 Correlated GAF

We now turn to the derivation of the GAF, first restricted to the case of unit spherical micrograins (i.e., $\beta_x = \beta_y = 1$ and $r_s = 1$). A second difference with the model of Lucas et al. is that we do not rely on Smith's uncorrelated height-normal assumption, but instead couple heights and microfacet normals through micrograins.

Unidirectional GAF. We start by deriving the unidirectional GAF for a direction \mathbf{i} , using a classical separation between local and distant terms [Heitz 2014]: $G_1(\mathbf{i}, \mathbf{m}) = G_1^{\text{local}}(\mathbf{i}, \mathbf{m}) G_1^{\text{dist}}(\mathbf{i}, \mathbf{m})$. The local term accounts for visibility masking by a single, local micrograin: $G_1^{\text{local}}(\mathbf{i}, \mathbf{m}) = \chi^+(\mathbf{i} \cdot \mathbf{m})$. Once accounted for, the local micrograin may be considered as "removed" from the distribution: we are thus

¹E.g., see https://agraphicsguynotes.com/posts/sample_anisotropic_microfacet_brdf/




Fig. 3. The unidirectional GAF G_1 (here for spherical micrograins) is composed of a local and a distant term. Top: once the local term has been accounted for, the local micrograin (dotted) may be removed and we are left with a point of normal \mathbf{m} inside a pore at height $h = \mathbf{m} \cdot \mathbf{n}$. The distant term G_1^{dist} characterizes the proportion of such points that are visible from a direction \mathbf{i} (i.e., points outside shadows). Middle: to characterize G_1^{dist} , we intersect the micrograin and the cylinder tangent to it and oriented along \mathbf{i} with \mathcal{H} , the plane at height h , yielding a circle C_h and an ellipse $\mathcal{E}_h(\mathbf{i})$ (of center e) respectively. Bottom: the shadow lies in the plane \mathcal{H} . Its area is given by $\sigma_s = \sigma_e - \sigma_c - \sigma_t$, with the component areas shown in red (outlined), purple and yellow respectively. As detailed in Appendix A, their computation relies on the points $q_{i\pm}$, where C_h is tangent to $\mathcal{E}_h(\mathbf{i})$.

left with a point at height $h = \mathbf{n} \cdot \mathbf{m}$, which makes the coupling between height and microfacet normals explicit. By construction, this point lies inside a pore at height h (see Figure 3-top).

Yet, throwing a ray from that point in direction \mathbf{i} might still intersect another micrograin. The distant term G_1^{dist} characterizes the probability of such a ray leaving out the surface. Formally, we define it as the relative area of pores at height h visible from the direction \mathbf{i} , with respect to pores at height h (i.e., visible at normal incidence). Since porosity is the complement of the filling factor, we rely on Equation 7 to derive the height-dependent distant term:

$$G_1^{\text{dist}}(\mathbf{i}, h) = \frac{e^{-\rho\sigma(\mathbf{i}, h)}}{e^{-\rho\sigma(\mathbf{n}, h)}} = e^{-\rho\sigma_s(\mathbf{i}, h)}, \quad (16)$$

where $\sigma(\mathbf{i}, h)$ is the projected area in direction \mathbf{i} of a spherical cap onto the plane at height h , and $\sigma_s(\mathbf{i}, h) = \sigma(\mathbf{i}, h) - \sigma(\mathbf{n}, h)$ is the

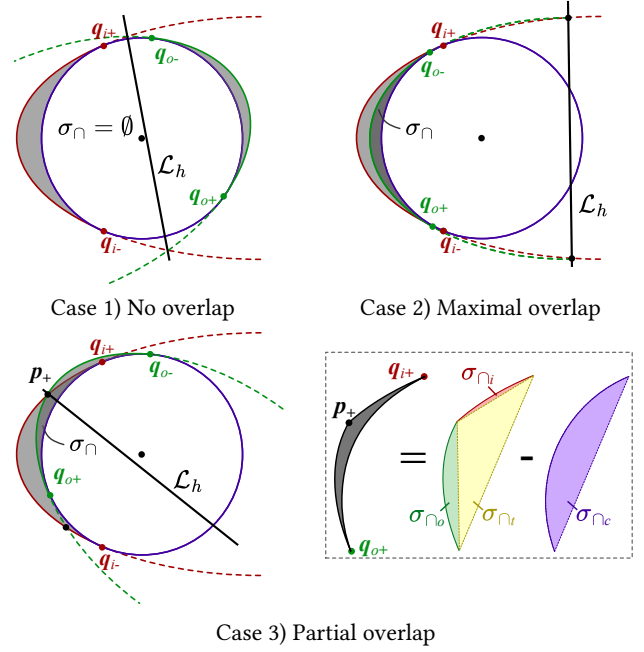


Fig. 4. Directional correlation in the bidirectional GAF requires to compute the area σ_\cap of the overlap (dark gray) between shadows from direction \mathbf{i} (light gray, outlined in red) and direction \mathbf{o} (light gray, outlined in green). Three cases must be distinguished, depending on the position of silhouette points $q_{i\pm}$ and $q_{o\pm}$ with respect to the line \mathcal{L}_h , as detailed in Appendix B. Case 1): $q_{i\pm}$ and $q_{o\pm}$ are on different sides of \mathcal{L}_h and there is no overlap. Case 2): $q_{i\pm}$ and $q_{o\pm}$ are on the same side of \mathcal{L}_h and the overlap is maximal. Case 3): \mathcal{L}_h separates q_{i+} from q_{i-} and q_{o+} from q_{o-} , yielding a partial overlap with shadow boundaries intersecting at p_+ . The shadow overlap area is then given by $\sigma_\cap = \sigma_{\cap i} + \sigma_{\cap o} - \sigma_{\cap c}$, with the component areas shown in yellow, red, green and purple respectively. As detailed in Appendix B, their computation relies on the points p_+ , q_{i+} and q_{o+} .

area of the spherical cap shadow. As illustrated in Figure 3-bottom and detailed in Appendix A, σ_s is obtained through the geometric decomposition: $\sigma_s = \sigma_e - \sigma_c - \sigma_t$, where σ_e is the area of an ellipse sector (outlined in red), σ_c is the area of a circle sector (in purple), and σ_t is the area of a triangle pair (in yellow). When $\mathbf{i} = \mathbf{n}$, or $h = 1$ (i.e., $\mathbf{m} = \mathbf{n}$), there is no shadow ($\sigma_s = 0$); hence $G_1(\mathbf{n}, h) = G_1(\mathbf{i}, 1) = 1$.

Bidirectional GAF. The bidirectional GAF is also separated into local and distant terms: $G(\mathbf{i}, \mathbf{o}, \mathbf{m}) = G^{\text{local}}(\mathbf{i}, \mathbf{o}, \mathbf{m})G^{\text{dist}}(\mathbf{i}, \mathbf{o}, \mathbf{m})$, with the local term given by $G^{\text{local}}(\mathbf{i}, \mathbf{o}, \mathbf{m}) = \chi^+(\mathbf{i} \cdot \mathbf{m})\chi^+(\mathbf{o} \cdot \mathbf{m})$. One may directly rely on the unidirectional GAF (Equation 16) to define the distant term of the bidirectional GAF:

$$G^{\text{dist}}(\mathbf{i}, \mathbf{o}, h) = G_1^{\text{dist}}(\mathbf{i}, h)G_1^{\text{dist}}(\mathbf{o}, h) = e^{-\rho[\sigma_s(\mathbf{i}, h) + \sigma_s(\mathbf{o}, h)]}. \quad (17)$$

Note that contrary to Smith's model, such a GAF is height-correlated by construction since G_1^{dist} is height-dependent. However, it does not account for potential retro-reflection effects that might occur as \mathbf{i} and \mathbf{o} get in alignment.

In order to take such directional correlation into account in G^{dist} , we must consider the potential overlap between the shadows due

to \mathbf{i} and \mathbf{o} (see Figure 4). In terms of projected areas, we have:

$$G^{\text{dist}}(\mathbf{i}, \mathbf{o}, h) = e^{-\rho[\sigma_s(\mathbf{i}, h) + \sigma_s(\mathbf{o}, h) - \sigma_\cap(\mathbf{i}, \mathbf{o}, h)]}, \quad (18)$$

where $\sigma_\cap(\mathbf{i}, \mathbf{o}, h)$ is the area of the intersection of spherical cap shadows corresponding to directions \mathbf{i} and \mathbf{o} .

We distinguish three cases (see Figure 4): 1) $\sigma_\cap(\mathbf{i}, \mathbf{o}, h) = 0$ when either of the shadows is absent, or if they do not overlap; 2) $\sigma_\cap(\mathbf{i}, \mathbf{o}, h) = \min(\sigma_s(\mathbf{i}, h), \sigma_s(\mathbf{o}, h))$ when one of the shadows fully overlaps the other; 3) $\sigma_\cap = \sigma_{\cap_t} + \sigma_{\cap_i} + \sigma_{\cap_o} - \sigma_{\cap_c}$ when the shadows partially overlap. As illustrated in the inset of Figure 4 and detailed in Appendix B, σ_{\cap_t} is the area of a triangle (in yellow), σ_{\cap_i} and σ_{\cap_o} are areas of elliptic segments (in red and green respectively), and σ_{\cap_c} is the area of a circle segment (in purple). When $\mathbf{i} = \mathbf{o} = \mathbf{n}$, or $h = 1$ (i.e., $\mathbf{m} = \mathbf{n}$), both shadows vanish, hence $G(\mathbf{n}, \mathbf{n}, h) = G(\mathbf{i}, \mathbf{o}, 1) = 1$.

Anisotropic GAF. Up until now we have only considered the case of spherical micrograins. In order to accommodate for elliptical micrograins, we transform the microstructure through a scaling matrix $E = \text{diag}(\beta_x, \beta_y, 1)$ that turns them into spherical micrograins. In such a transformed space, the ingoing and outgoing directions are given by $\mathbf{i}' = \frac{E\mathbf{i}}{\|E\mathbf{i}\|}$ and $\mathbf{o}' = \frac{E\mathbf{o}}{\|E\mathbf{o}\|}$, whereas the transformed microfacet normal is classically given by $\mathbf{m}' = \frac{E^{-1}\mathbf{m}}{\|E^{-1}\mathbf{m}\|}$ with the corresponding height $h' = \mathbf{n} \cdot \mathbf{m}'$. Any projected area computed on this scaled microstructure must then be multiplied by $\det E^{-1} = \frac{1}{\beta_x \beta_y}$.

The distant term of the bidirectional GAF is now given by:

$$G^{\text{dist}}(\mathbf{i}, \mathbf{o}, h') = e^{-\rho[\sigma_s(\mathbf{i}', h') + \sigma_s(\mathbf{o}', h') - \sigma_\cap(\mathbf{i}', \mathbf{o}', h')]} \det E^{-1},$$

Using Equation 9 (with $r_s = 1$) for ρ , we obtain the final formula for the distant term of the anisotropic GAF (G^{local} remains unchanged):

$$G^{\text{dist}}(\mathbf{i}, \mathbf{o}, h') = e^{\frac{\ln(1-\tau_0)}{\pi} [\sigma_s(\mathbf{i}', h') + \sigma_s(\mathbf{o}', h') - \sigma_\cap(\mathbf{i}', \mathbf{o}', h')]} \quad (19)$$

4.3 Our BSDF model

The third and last difference between our model and the one of Lucas et al. lies in the combination of the different BSDF components. We use the following formula, which has a form similar to Equation 1:

$$f(\mathbf{i}, \mathbf{o}) = \tau_0 f^S(\mathbf{i}, \mathbf{o}) + V_p(\mathbf{i}, \mathbf{o}) f^B(\mathbf{i}, \mathbf{o}), \quad (20)$$

where V_p denotes the visibility of pores between micrograins.

We directly use τ_0 to weight the surfacic micrograin component. This is valid for its specular component, since the height- and direction-correlated GAF already exactly accounts for all attenuation effects due to visibility. However, we need to modify the diffuse micrograin component to integrate the visibility of points over micrograins. This is very much similar to the case of rough diffuse microfacets studied by Heitz [2014]. In our case, we obtain:

$$f^{S,d}(\mathbf{i}, \mathbf{o}) = \int_{\Omega} \frac{K_d(\mathbf{m})}{\pi} \frac{|\mathbf{i} \cdot \mathbf{m}| |\mathbf{o} \cdot \mathbf{m}|}{|\mathbf{i} \cdot \mathbf{n}| |\mathbf{o} \cdot \mathbf{n}|} D(\mathbf{m}) G(\mathbf{i}, \mathbf{o}, \mathbf{m}) T(\mathbf{i}, \mathbf{o}, \mathbf{m}) d\mathbf{m}. \quad (21)$$

with $T(\mathbf{i}, \mathbf{o}, \mathbf{m}) = (1 - F(\mathbf{i} \cdot \mathbf{m}))(1 - F(\mathbf{o} \cdot \mathbf{m}))$. Compared to the model of Lucas et al., we explicitly consider the visibility of (anisotropic) micrograins, we use a more accurate transmittance instead of the approximation $\bar{T}(\mathbf{i}, \mathbf{o})$, and a diffuse albedo K_d that may optionally vary as a function of the microfacet normal \mathbf{m} (i.e., the position on the micrograin). Unfortunately, Equation 21 has no closed form,

hence we resort to a stochastic evaluation. This is similar to the work of Heitz and Dupuy [2015], except that we use our D and G terms and sample the NDF instead of the visible NDF (see Section 7).

The surfacic micrograin component is now given by:

$$f^S(\mathbf{i}, \mathbf{o}) = f^{S,d}(\mathbf{i}, \mathbf{o}) + \frac{D(\mathbf{h})G(\mathbf{i}, \mathbf{o}, \mathbf{h})F(\mathbf{i} \cdot \mathbf{h}, \mathbf{h})}{4|\mathbf{i} \cdot \mathbf{n}| |\mathbf{o} \cdot \mathbf{n}|}, \quad (22)$$

where we use our new NDF D and GAF G . As with K_d , we let the refractive index in the Fresnel term F optionally vary as a function of \mathbf{h} , which allows for micrograins with locally varying reflectance.

The bulk material is only visible for light paths that reach pores between micrograins from directions \mathbf{i} and \mathbf{o} , as modeled by $V_p(\mathbf{i}, \mathbf{o})$. At normal incidence ($\mathbf{i} = \mathbf{o} = \mathbf{n}$), it is simply given by $(1 - \tau_0)$. In other configurations, part of the pores may not be visible from either the ingoing or outgoing direction. Such an attenuation is directly given by the distant term of the anisotropic GAF (Equation 19), using $h' = 0$ since pores are located on the surface:

$$V_p(\mathbf{i}, \mathbf{o}) = (1 - \tau_0) G^{\text{dist}}(\mathbf{i}, \mathbf{o}, 0). \quad (23)$$

We show in Appendix C that $G^{\text{dist}}(\mathbf{i}, \mathbf{o}, 0)$ has a simple formula that leads to a slightly more efficient implementation for V_p .

Using the visible filling factor notation of Lucas et al., we observe that in the isotropic case, Equation 16 may be rewritten as:

$$G_1^{\text{dist}}(\mathbf{i}, \mathbf{o}) = \frac{1 - \tau_B^+(\mathbf{i})}{1 - \tau_0}. \quad (24)$$

Hence we may rewrite their bulk weight factor as (see Equation 4):

$$1 - w^+(\mathbf{i}, \mathbf{o}) = (1 - \tau_0) G_1^{\text{dist}}(\mathbf{i}, 0) G_1^{\text{dist}}(\mathbf{o}, 0). \quad (25)$$

In other words, compared to our model, the weighting factor of the bulk component in the model of Lucas et al. (Equation 1) is not direction-correlated and limited to the isotropic case.

Importance Sampling. The sampling of the BSDF of Equation 20 according to an ingoing direction \mathbf{i} first requires the sampling of a BRDF component. We know that the visibility of the bulk component from direction \mathbf{i} corresponds to the complement of the visible filling factor. Using Equation 24, we define the probability $p_B(\mathbf{i})$ of sampling the bulk material by:

$$p_B(\mathbf{i}) = \frac{(1 - \tau_0) G_1^{\text{dist}}(\mathbf{i}, 0)}{\tau_0 + (1 - \tau_0) G_1^{\text{dist}}(\mathbf{i}, 0)}, \quad (26)$$

where $G_1^{\text{dist}}(\mathbf{i}, h') = e^{\frac{\ln(1-\tau_0)}{\pi} \sigma_s(\mathbf{i}', h')}$ is the anisotropic unidirectional GAF (the derivation is similar to that of Equation 19). Note that at normal incidence we have $p_B(\mathbf{n}) = 1 - \tau_0$ as expected.

The surfacic component, which is sampled with a probability $1 - p_B(\mathbf{i})$, is itself made of a pair of diffuse and specular sub-components. We define the probability $p_{S,d}(\mathbf{i})$ of sampling the diffuse micrograin sub-component as:

$$p_{S,d}(\mathbf{i}) = \frac{K_d(\mathbf{n})}{K_d(\mathbf{n}) + F(\mathbf{i} \cdot \mathbf{n}, \mathbf{n})}, \quad (27)$$

where we use the diffuse albedo (resp. refractive index) at the tip of the micrograin (i.e., with normal \mathbf{n}), since this is the part of the micrograin surface that is always visible.




Fig. 5. The azimuthal angle dependence of G_1 (green solid lines) for fixed $\theta_m = 50^\circ$ and $\theta_i = 64.5^\circ$ is only due to G_1^{local} , which equals 1 for a subset of angles ϕ_m (blue simulation samples), and is equal to 0 for most of the other angles (red simulation samples). Outliers (in red) are due to numerical inaccuracies in simulations at points where $\mathbf{i} \cdot \mathbf{m} \approx 0$.

We finally rely on cosine sampling to sample $f^{S,d}$, and on the methodology of Walter et al. [2007] in conjunction with Equations 14-15 to sample the microfacet component.

5 VALIDATIONS AND COMPARISONS

We first provide several validations of our model using light transport simulations on explicit micrograin distributions in Section 5.1. We then compare our BSDF to other methods and highlight the differences brought by explicitly considering height-normal dependencies in Section 5.2. We only present a subset of our comparison results here and refer the interested reader to supplemental material for a more in-depth evaluation.

5.1 Validations against simulations

We build an explicit microstructure by randomly distributing micrograins onto a plane, using ray-tracing to obtain simulated references.

Comparison with simulated GAF. We start by evaluating our unidirectional GAF, which is separated into local and distant terms. This separation is validated in Figure 5 by simulations on spherical micrograins. For each configuration, we inspect the values of G_1 as a function of ϕ_m for fixed elevation angle θ_m (i.e., $h = \cos \theta_m$) and incoming elevation angle θ_i . As expected, G_1 is constant (up to simulation noise) except for azimuthal angles where G_1^{local} is zero.

We next compare the distant term of the unidirectional GAF G_1^{dist} to a simulation where the contributions of all visible micrograin surface points at a given height (blue samples in Figure 5) are averaged to reduce variance. This is shown in Figure 6 for several filling factors τ_0 and micrograin elevation angles θ_m , with G_1^{dist} plotted as a function of the incident angle θ_i . Our model exactly matches simulations, validating our approach based on projected areas of spherical caps (or micrograin shadow areas).

Comparison with simulated BRDF. Lastly, we validate our full material model against simulations, obtained by forward path tracing on the explicit micrograin surface followed by collection of outgoing rays in a distant hemisphere. In these comparisons, we restrict simulations to single-scattering light paths, in order to precisely evaluate the potential errors of our model, which neglects multiple-scattering (but see Section 7 for comparisons with full simulations).




Fig. 6. Validation of G_1^{dist} (solid curves) against simulations (dots) as a function of the incident angle θ_i for several (τ_0, θ_0) configurations.





Fig. 7. Validation of BRDF slices ($\theta_i = 86^\circ$) for *conductor* micrograins (of index $\eta + i\kappa = 1.0152 + i6.6273$) distributed onto an absorbing base surface. Our direction-correlated model (right columns) exhibits very high accuracy w.r.t the reference (center). Without direction correlation (left columns), a retro-reflection peak is absent. The green isoline separates angular configurations where the BRDF is below 0.00001, in which case simulation noise prevails and the relative error is not relevant anymore.

We show in Figures 7 and 8 two different configurations using conductor and diffuse micrograins respectively, both on a fully absorbing bulk surface and for different configurations of material parameters $(\tau_0, \beta_x, \beta_y)$. In each case, we compare our model with or without direction correlation (Equations 17 and 18 respectively) at grazing angles to better exhibit retro-reflection effects. Many more such comparisons are given in supplemental materials; they all show near-perfect matches of our full model with the reference.

Each image represents a BRDF slice parametrized according to zenithal and azimuthal outgoing angles as shown in the inset image. The relative errors are computed by taking the difference between the model and the reference, and normalizing by the reference. A gray color (relative error of 0) means a perfect match.



We have not found necessary to separately validate the bidirectional GAF. Indeed, our NDF in the isotropic case is identical to the one of Lucas et al. [2023], which has already been validated in their




Fig. 8. Validation of BRDF slices ($\theta_i = 86^\circ$) for *diffuse* micrograins (of albedo $K_d = 0.5$) distributed onto an absorbing base surface. Once again, our direction-correlated model (right columns) exhibits very high accuracy with respect to the reference (center). The retro-reflection peak due to direction correlation (absent from the two left columns), remains thin even with diffuse micrograin configurations as it is mainly due to visibility.

work; hence the only potential source of error for our BRDF model when considering an absorbing bulk lies in the bidirectional GAF. Since we have obtained excellent comparisons of our full model compared to the reference, it seems reasonable to assume that this is due to a correct bidirectional GAF. The numerous additional validations presented in supplemental material confirm this assumption.

5.2 Comparisons between models

We provide comparisons to a few other BRDF models that handle the mixing of multiple microfacet distributions: the empirical model of Merillou et al. [2000] that assumes cylindrical micropores “drilled” in a rough surface; the model of d’Eon et al. [2023] based on different microfacet distributions combined based on depth; and the micrograin model of Lucas et al. [2023]. As discussed in the latter, the parameters of all three models may be made to coincide to obtain similar results at normal incidence. We compare all four models (including ours) to a reference simulation in the single-scattering case, with conductor micrograins and an absorbing bulk surface, on various combinations of material parameters (τ_0, β_x, β_y). We believe these result in fair comparisons since all models provide analytic formula in such configurations.

One of those comparisons is shown in Figure 9, where it is clear that only our model accurately reproduces the reference simulation. This is confirmed by all the other comparisons provided in supplemental material. In this case, the other three models seem to produce similar results relative to each others, but this is not the case in general, in particular for the model of Merillou et al. We do not take these comparisons to suggest that our model is better, since the other models assume different microstructures. Instead it shows that an explicit modeling of height-normal dependencies yields different BRDF slices compared to previous work. Differences in terms of appearance are shown in Figure 10, where we omit the model of Lucas et al. [2023] since its results are very similar to those of the model of d’Eon et al. [2023]. The model of Merillou et al. [2000] markedly differs from ours, while the model of d’Eon et al. [2023]




Fig. 9. Comparison of different BSDF models with parameters reproducing a distribution of spherical conductor micrograins ($\tau_0 = 0.5$, $\beta_x = \beta_y = 1$, $\eta + i\kappa = 1.0152 + i6.6273$) on an absorbing base.




Fig. 10. Comparisons between global illumination renderings using our specular micrograin model (a), the model of Merillou et al. [2000] (b) and the one d’Eon et al [2023] (c), for a pair of filling factors (rows). We use a gold porous layer on top of an absorbing base. We set $\beta = 1$ and compute an equivalent roughness for (b) and (c) using Equation 28. FLIP images show visual differences with our model.

is visually closer, with remaining differences mostly visible toward grazing angles due to our GAF.

In Figure 11, we further compare our model (Equations 17 and 18) and the one of Lucas et al. [2023] in a retro-reflection configuration. This shows the influence of height-normal dependencies at grazing angles. The effect becomes more noticeable at higher β values, as shown by FLIP differences [Andersson et al. 2020].




Fig. 11. Comparisons in a retro-reflection configuration (i.e., $\mathbf{i} = \mathbf{o}$) between our model (using Equation 17 or 18) and the model of Lucas et al. [2023]. We use a distribution of conductor micrograins ($\tau_0 = 0.2$) with increasing values for β on a Lambertian base surface (of albedo 0.5). The renderings (top row) differ more for increasing values of β when visibility effects become prominent, as shown by FLIP differences with our full model (bottom row).

6 RESULTS

We have implemented our model in Mitsuba [Jakob et al. 2022] to generate results with global illumination. All conductor refractive indices are taken from the IOR database used in Mitsuba.

Our preliminary implementation in GLSL inside BRDFExplorer shows that our model introduces a 3.7% overhead when using the version without directional correlation and an overhead of 7.6% when using the fully-correlated GAF, compared to a standard anisotropic GGX BRDF model implemented with a height-correlated Smith GAF. These timings were performed on a PC equipped with a Nvidia RTX 2080 graphics card.

Material parameters. One key parameter that controls the appearance of our BSDF model is the filling factor τ_0 , which acts as a mixture weight between the appearance of micrograins and the appearance of the bulk surface (see Equation 20). This is illustrated in Figure 12 where we show two different types of micrograins (dielectrics and conductors) on top of a Lambertian base material, for several values of τ_0 . Our model thus provides an alternative solution to existing work on material mixtures (e.g., [d’Eon et al. 2023; Kelemen and Szirmay-Kalos 2001; Meneveaux et al. 2018]).

The other parameters of our model are the micrograin compression/stretching factors (β_x, β_y) and micrograin reflectance properties. As already mentioned by Lucas et al. [2023], the compression/stretching factors may be related to equivalent roughness parameters by the following equation, which also applies to anisotropy:

$$(\tilde{\alpha}_x, \tilde{\alpha}_y) = \sqrt{\frac{\tau_0}{\ln(1 - \tau_0)}} (\beta_x, \beta_y) \quad (28)$$

We show in Figure 1 different combinations of all material parameters on a 3D Buddha model. The golden material of Figure 1a is obtained by laying anisotropic aluminum micrograins ($(\tau_0, \beta_x, \beta_y) = (0.2, 0.05, 2)$) onto an isotropic rough gold conductor base ($\alpha = 0.2$). This shows another example of material mixture, this time between conductor material components. The dusty wooden Buddha of Figure 1b is obtained by applying gray diffuse micrograins ($K_d = 0.2$,




Fig. 12. **Material mixtures.** We use the filling factor parameter to control the mixture of micrograins ($\beta = 0.1$) and bulk surface. (a) Dielectric micrograins ($\eta = 1.6, K_d = [0.3, 0.2, 0.1]$) on top of a beige diffuse bulk (same K_d). (b) Gold micrograins on top of a gray diffuse bulk (of albedo 0.4).

$\beta = 1$) on top of a textured diffuse wood bulk material. We vary the filling factor according to the upward component of the surface normal to produce plausible dust deposition. Similarly, a moldy glass Buddha is shown in Figure 1c, this time using green diffuse + dielectric ($K_d = [0.2, 0.5, 0.05], \eta = 1.5$) micrograins onto a rough dielectric surface ($\alpha = 0.2$ with an index of 1.5), illustrating the ability of our model to use a transparent material for the base surface. A subtle example of color correlation is given in Figure 1d: we use diffuse micrograins ($(\tau_0, \beta) = (0.4, 3)$) with albedos varying from green to red as a function of height, distributed on top of a rough aluminum conductor base ($\alpha = 0.2$). The effect may be observed at extremely grazing angles where the micrograin color takes on a reddish tint. Lastly, in Figure 1e, we use a diffuse white plastic weave texture for the base surface, which is completely transparent in some locations. We apply a distribution of grayish diffuse micrograins ($\beta = 1, K_d = [0.1, 0.05, 0]$) on top of it, with a filling factor that varies spatially in order to mimic the accumulation of dirt around locations where the weaves intersect.

In the supplemental video, we show an interactive manipulation of material parameters using our BRDFExplorer implementation. In




Fig. 13. **Diffuse anisotropy.** We compare the appearance of isotropic (top) and anisotropic (bottom) diffuse materials, with filling factors varying per object: from back to front, $\tau_0 = \{0.05, 0.15, 0.3\}$. As shown in zoom insets, grazing-angle visibility effects vary in intensity with the filling factor, and in orientation according to anisotropy (see also the FLIP difference inset).

particular, we demonstrate normal-correlated complex refractive indices, using Gulbrandsen’s remapping [Gulbrandsen 2014] for simpler control as described in supplemental material.

Anisotropic materials. Our anisotropic model produces interesting view-dependent effects. We first show the case of anisotropic diffuse micrograins on a diffuse base in Figure 13. For all objects in the scene, we use a red diffuse albedo ($K_d = [0.6, 0, 0]$) for the micrograin and a gray diffuse albedo for the bulk component ($K_d = 0.4$). However, we vary the filling factors per object, and we compare isotropic and anisotropic configurations. The reddish grazing angle visibility effects differ in terms of distribution in each case (see zoom insets).

Equation 20 may also be used to mix two layers of anisotropic micrograins by using a micrograin layer for the bulk component. This is shown in Figure 14 (top) for anisotropic dielectric micrograins distributed on another layer of anisotropic dielectric micrograins oriented perpendicularly. Both layers use the same refractive index for micrograins, but two different colors for the diffuse albedo. Note that the material is partially transparent in this case. As a result, the material appearance changes with rotations of the camera, as shown in zoom insets. A similar configuration is used in Figure 14 (bottom), except that we use aluminum micrograins and add a third aluminum bulk component in this case to make the material opaque.

Note that even though we have used a thin surface that appears folded on a sphere, our intent is *not* to convey a fabric appearance. Indeed, as discussed in Section 7, the structural assumptions made

by our model are not adapted to such materials. The images in Figure 14 thus resemble more *trompe l’oeil* sculptures².

Retro-reflection effects. We have shown the impact of accounting for directional correlation in Figure 11 in simple retro-reflection configurations. We illustrate this effect on a more complex object in an environment lighting with an additional strong directional light. When the key light comes from the side (Figure 15a), the difference between versions of our model with or without directional correlation (left and right halves respectively) remains subtle (as shown by FLIP differences), especially for conductor micrograins. In contrast, when the light and view directions are aligned (Figure 15b), the difference is much more pronounced: the material with diffuse micrograins conveys a softer look, whereas the one with the conductor micrograins exhibits more intense grazing angle effects.

Correlated color effects. A last feature of our model is the ability to make reflectance vary locally on each micrograin. This is shown in Figure 16, where we use two types of correlation: based on height, and based on microfacet normal orientations, both on dielectric+diffuse and conductor micrograins. This bears some similarities to the filtering of color-mapped textures and surfaces [Heitz et al. 2013b], except that in our case, all the filtering occurs at the microscopic scale and is handled by our fully-correlated GAF. The effect on appearance varies depending on micrograin reflectance. Materials based on dielectric+diffuse micrograins exhibit color variations that are expected from local variations of the diffuse albedo on a micrograin. The appearance of materials based on conductor micrograins where the complex refractive index is made to vary locally (through variations of its imaginary term) is more subtle. In such cases, one should pay particular attention to color fringes *around* highlights (see zoom insets).

Additional results and comparisons. In supplemental material, we further compare our fully correlated GAF against the micrograin Smith GAF from Lucas et al. [2023] (with and without height correlation). We provide an analysis of convergence of our micrograin model for both specular and diffuse micrograins against the model of Lucas et al. [2023].

7 DISCUSSION

We have introduced a new approach to explicitly model height-normal dependencies in structured microspheres, based on the micrograin model of Lucas et al. [2023]. We extend its NDF to deal with anisotropy, and introduce a new fully-correlated GAF that is used in micrograin BSDF components, as well as in the definition of pore visibility. Our model is validated against reference simulations, and we show how its parameters affect material appearance, through retro-reflection, grazing angle visibility and correlated colors.

Limitations. Our model currently only handles single scattering, which is its main limitation. As shown in Figure 17, white furnace tests performed on perfect conductor micrograins on a white diffuse base exhibit losses of energy, as expected. The relevance of the single-scattering approximation depends on all three parameters τ_0 , β_x and β_y , and varies with the incoming direction. This

²See for instance the sculpture “Draped chair” by Marina Karella.




Fig. 14. **Specular anisotropy.** We use a pair of orthogonal anisotropic micrograin layers to create a complex specular anisotropic material appearance. From left to right, we rotate the camera to exhibit view-dependent effects. Top: using dielectric+diffuse micrograins (with $\eta = 1.5$, $(\tau_0, \beta_x, \beta_y) = (0.3, 0.05, 0.8)$ for the top layer and $(\tau_0, \beta_x, \beta_y) = (0.3, 0.7, 0.1)$ for the bottom one) of different albedo colors – respectively red ($K_d = [0.5, 0, 0]$) and green ($K_d = [0, 0.5, 0]$) – conveys the appearance of a corrugated plastic sheet. Bottom: using a pair of aluminum micrograin layers with parameters identical to the dielectric case for $(\tau_0, \beta_x, \beta_y)$ conveys the appearance of a *trompe l'oeil* sculpture.




Fig. 15. **Retro-reflection effects.** Our model with and without directional correlation (left and right halves respectively), in two different orientations (a,b). Top: diffuse micrograins (with $K_d = [0.1, 0.2, 0.4]$ and $(\tau_0, \beta) = (0.4, 1)$) on top of a diffuse base of same albedo. Bottom: gold micrograins (with $(\tau_0, \beta) = (0.3, 2)$) on top of a gray diffuse base (of albedo 0.2). Inset images show FLIP differences, which are much more pronounced in retro-reflection configurations as expected.

is shown in Figure 18 where we compare our model to a full simulation (i.e., including multiple scattering light paths), this time




Fig. 16. **Correlated color effects.** Our model allows for local variations of micrograin reflectances (top left insets, seen from top), for instance based on (a) the micrograin height or (b) the orientation of its normal. Top: using diffuse micrograins (with $\tau_0 = 0.4$, $\beta = 3$) on top of a rough aluminium bulk ($\alpha = 0.2$). Bottom: using conductor micrograins (with $\tau_0 = 0.8$, $\beta = 0.4$ in (a) and $\tau_0 = 0.4$, $\beta = 0.1$ in (b)) onto a diffuse bulk (of albedo 0.01 in (a) and 0.2 in (b)).

using conductor micrograins on a diffuse base of albedo 0.5. We compute directional albedos for both our model and simulation and display ratios between them: the closer the ratio is to 1, the better




Fig. 17. White furnace tests for different values of τ_0 and β , using perfect conductor isotropic micrograins on a white diffuse base.




Fig. 18. Ratios of directional albedos between our model and the full reference simulation, for varying incidence angles θ_i and a fixed $\phi_i = 0$. The contribution of multiple scattering increases with increasing filling factor τ_0 , increasing β and increasing incidence angle θ_i .

the single-scattering approximation. These matrices are not symmetric, since we vary θ_i but use a fixed $\phi_i = 0$. The proportion of multiple-scattering paths increases with higher values of τ_0 and β since incoming rays have less probability of directly being reflected outside of the micro-surface in these cases. However, at extremely grazing angles, the single-scattering approximation yields better results since only micrograin tips become visible.

Extending our model to handle multiple-scattering requires the derivation of the Visible Normal Distribution Function (VNDF), which would also help reduce variance in importance sampling. Note that for simplicity, we have used a direct transmission term $T(\mathbf{i}, \mathbf{o}, \mathbf{m})$ in the definition of the diffuse microgain component $f^{S,d}$ in Equation 21, which neglects inter-reflections *inside* a dielectric+diffuse micrograin. This could be corrected by using a more complex transmission term, such as the one used in the interfaced Lambertian model [Simonot 2009]. Our BSDF model could also be optimized by providing analytical approximations for $f^{S,d}$.

Future work. The shape of micrograins could be made more general than half-ellipsoids. Indeed, as long as we keep considering a

uniform random distribution of *identical* micrograins, Equation 7 will remain valid. The only constraint on micrograin shape in the micrograin framework is the existence of a bijection between the height h and the microfacet normal \mathbf{m} for a point on a micrograin.

Variations in micrograin radii or densities would be more complex to handle inside the BSDF model, but would be particularly interesting for multi-scale rendering. A more challenging direction of future work would be to consider other types of micrograin distributions. For instance, blue noise distributions of micrograins could be used to model the appearance of microscopic Voronoi patterns (e.g., cells in leaves), but this remains an open problem. In its current form, our model should be mainly considered for artistic use. However, the interesting appearance effects it produces might be replicated on real-world examples through fabrication, taking inspiration from previous work (e.g., [Luongo et al. 2019; Rouiller et al. 2013]).

Finally, we believe the micrograin framework could be extended to model a volumetric bulk, as an alternative to methods based on the diffusion approximation (e.g., [Hapke 1986]). This would require the extension of micrograins to full spheres, and an accurate account of the interplay between surfacic and volumetric components.

ACKNOWLEDGMENTS

We would like to thank the reviewers for helping us improve the paper. This work has been funded by the Inria Exploratory Action Ecoptics.

REFERENCES

- Pontus Andersson, Jim Nilsson, Tomas Akenine-Möller, Magnus Oskarsson, Kalle Åström, and Mark D. Fairchild. 2020. FLIP: A Difference Evaluator for Alternating Images. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 3, 2 (2020), 15:1–15:23. <https://doi.org/10.1145/3406183>
- M. Ashikhmin, S. Premoze, and P. Shirley. 2000. A microfacet-based BRDF generator. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 2000, New Orleans, LA, USA, July 23-28, 2000*. ACM, 65–74.
- P. Barla, R. Pacanowski, and P. Vangorp. 2018. A Composite BRDF Model for Hazy Gloss. *Computer Graphics Forum* 37, 4 (2018), 55–66. <https://doi.org/10.1111/cgf.13475> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13475>
- Laurent Belcour and Pascal Barla. 2017. A Practical Extension to Microfacet Theory for the Modeling of Varying Iridescence. *ACM Transactions on Graphics* 36, 4 (July 2017), 65. <https://doi.org/10.1145/3072959.3073620>
- Benedikt Bitterli and Eugene d'Eon. 2022. A Position-Free Path Integral for Homogeneous Slabs and Multiple Scattering on Smith Microfacets. *Computer Graphics Forum* 41, 4 (2022), 93–104. <https://doi.org/10.1111/cgf.14589> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14589>
- R. L. Cook and K. E. Torrance. 1982. A Reflectance Model for Computer Graphics. In *ACM SIGGRAPH proceedings*.
- Eugene d'Eon, Benedikt Bitterli, Andrea Weidlich, and Tizian Zeltner. 2023. Microfacet theory for non-uniform heightfields. In *SIGGRAPH 2023 Conference Papers* (Los Angeles, CA, USA). Association for Computing Machinery, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3588432.3591486>
- Zhao Dong, Bruce Walter, Steve Marschner, and Donald P. Greenberg. 2016. Predicting Appearance from Measured Microgeometry of Metal Surfaces. *ACM Trans. Graph.* 35, 1, Article 9 (dec 2016), 13 pages. <https://doi.org/10.1145/2815618>
- Jonathan Dupuy, Eric Heitz, and Eugene d'Eon. 2016. Additional Progress towards the Unification of Microfacet and Microflake Theories. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations* (Dublin, Ireland) (EGSR '16). Eurographics Association, Goslar, DEU, 55–63.
- Alejandro Conty Estevez and Christopher Kulla. 2017. Production Friendly Microfacet Sheen BRDF. *ACM SIGGRAPH 2017* (2017).
- Ole Gulbrandsen. 2014. Artist Friendly Metallic Fresnel. *Journal of Computer Graphics Techniques (JCGT)* 3, 4 (9 December 2014), 64–72. <http://jcgt.org/published/0003/04/03/>
- Yu Guo, Miloš Hašan, and Shaung Zhao. 2018. Position-Free Monte Carlo Simulation for Arbitrary Layered BSDFs. *ACM Trans. Graph.* 37, 6 (2018).
- Bruce Hapke. 1986. Bidirectional reflectance spectroscopy: 4. The extinction coefficient and the opposition effect. *Icarus* 67, 2 (1986), 264–280. <https://doi.org/10.1016/0019->

- 1035(86)90108-9
- Eric Heitz. 2014. Understanding the Masking-Shadowing Function in Microfacet-Based BRDFs. *Journal of Computer Graphics Techniques* 3, 2 (June 2014).
- Eric Heitz, Christophe Bourlier, and Nicolas Pinel. 2013a. Correlation effect between transmitter and receiver azimuthal directions on the illumination function from a random rough surface. *Waves in Random and Complex Media* 23, 3 (Aug. 2013), 318–335. <https://doi.org/10.1080/17455030.2013.822611>
- Eric Heitz and Jonathan Dupuy. 2015. *Implementing a Simple Anisotropic Rough Diffuse Material with Stochastic Evaluation*. Technical Report.
- Eric Heitz, Johannes Hanika, Eugene d'Eon, and Carsten Dachsbacher. 2016a. Multiple-Scattering Microfacet BSDFs with the Smith Model. *ACM Trans. Graph.* 35, 4, Article 58 (jul 2016), 14 pages. <https://doi.org/10.1145/2897824.2925943>
- Eric Heitz, Johannes Hanika, Eugene d'Eon, and Carsten Dachsbacher. 2016b. Multiple-Scattering Microfacet BSDFs with the Smith Model. 35, 4, Article 58 (July 2016), 14 pages.
- Eric Heitz, Derek Nowrouzezahrai, Pierre Poulin, and Fabrice Neyret. 2013b. Filtering Color Mapped Textures and Variances. In *13D '13 - ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (13D '13 Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games)*, Stephen N. Spencer (Ed.). ACM, Orlando, United States, 129–136. <https://doi.org/10.1145/2448196.2448217>
- Nicolas Holzschuch and Romain Pacanowski. 2017. A two-scale microfacet reflectance model combining reflection and diffraction. *ACM Trans. Graph.* 36, 4, Article 66 (jul 2017), 12 pages. <https://doi.org/10.1145/3072959.3073621>
- Weizhen Huang, Sebastian Merzbach, Clara Callenberg, Doekele Stavenga, and Matthias Hullin. 2022. Rendering Iridescent Rock Dove Neck Feathers. In *ACM SIGGRAPH 2022 Conference Proceedings (Vancouver, BC, Canada) (SIGGRAPH '22)*. Association for Computing Machinery, New York, NY, USA, Article 43, 8 pages. <https://doi.org/10.1145/3528233.3530749>
- Wenzel Jakob, Eugene D'Eon, Otto Jakob, and Steve Marschner. 2014. A Comprehensive Framework for Rendering Layered Materials. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 33, 4 (July 2014), 118:1–118:14. <https://doi.org/10.1145/2601097.2601139>
- Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. 2022. *Mitsuba 3 renderer*. <https://mitsuba-renderer.org>.
- Csaba Kelemen and Laszlo Szirmay-Kalos. 2001. A Microfacet Based Coupled Specular-Matte BRDF Model with Importance Sampling. In *Eurographics 2001 - Short Presentations*. Eurographics Association. <https://doi.org/10.2312/egs.20011003>
- Simon Lucas, Mickael Ribardiere, Romain Pacanowski, and Pascal Barla. 2023. A Micrograin BSDF Model for the Rendering of Porous Layers. In *SIGGRAPH Asia 2023 Conference Papers (Sydney, NSW, Australia) (SA '23)*. Association for Computing Machinery, New York, NY, USA, Article 40, 10 pages. <https://doi.org/10.1145/3610548.3618241>
- Andrea Luongo, Viggo Falster, Mads Doest, M. Ribo, Er Eiriksson, David Pedersen, and Jeppe Frisvad. 2019. Microstructure Control in 3D Printing with Digital Light Processing. *Computer Graphics Forum* 39 (08 2019). <https://doi.org/10.1111/cgf.13807>
- D. Meneveaux, B. Bringier, E. Tautzia, M. Ribardiere, and L. Simonot. 2018. Rendering Rough Opaque Materials with Interfaced Lambertian Microfacets. *IEEE Transactions on Visualization & Computer Graphics* 24, 03 (mar 2018), 1368–1380. <https://doi.org/10.1109/TVCG.2017.2660490>
- Stéphane Mérrillou, Jean-Michel Dischler, and Djamchid Ghazanfarpour. 2000. A BRDF Postprocess to Integrate Porosity on Rendered Surfaces. 6, 4 (2000). <https://doi.org/10.1109/2945.895876>
- M. Oren and S. K. Nayar. 1994. Generalization of Lambert's Reflectance Model. In *ACM SIGGRAPH proceedings*.
- M. Ribardière, B. Bringier, D. Meneveaux, and L. Simonot. 2017. STD: Student's t-Distribution of Slopes for Microfacet Based BSDFs. *Computer Graphics Forum* (2017).
- Vincent Ross, Denis Dion, and Guy Potvin. 2005. 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. <https://doi.org/10.1364/JOSAA.22.002442>
- Olivier Rouiller, Bernd Bickel, Jan Kautz, Wojciech Matusik, and Marc Alexa. 2013. 3D-printing spatially varying BRDFs. *Computer Graphics and Applications, IEEE* 33 (11 2013), 48–57. <https://doi.org/10.1109/MCG.2013.82>
- Iman Sadeghi, Oleg Bisker, Joachim De Deken, and Henrik Wann Jensen. 2013. A practical microcylinder appearance model for cloth rendering. *ACM Trans. Graph.* 32, 2, Article 14 (apr 2013), 12 pages. <https://doi.org/10.1145/2451236.2451240>
- Lionel Simonot. 2009. Photometric model of diffuse surfaces described as a distribution of interfaced Lambertian facets. *Appl. Opt.* 48, 30 (Oct 2009), 5793–5801. <https://doi.org/10.1364/AO.48.005793>
- B. Smith. 1967. Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation* 15, 5 (September 1967), 668–671.
- K. E. Torrance and E. M. Sparrow. 1967. Theory for Off-Specular Reflection From Roughened Surfaces*. *J. Opt. Soc. Am.* 57, 9 (Sep 1967), 1105–1114. <https://doi.org/10.1364/JOSA.57.001105>
- Bram van Ginneken, Marigo Stavridi, and Jan J. Koenderink. 1998. Diffuse and Specular Reflectance from Rough Surfaces. *Appl. Opt.* 37, 1 (Jan 1998), 130–139. <https://doi.org/10.1364/AO.37.000130>
- B. Walter, S. R. Marschner, H. Li, and K. E. Torrance. 2007. Microfacet Models for Refraction Through Rough Surfaces. In *Computer Graphics Forum, EGSR proceedings*.
- Beibei Wang, Wenhua Jin, Miloš Hašan, and Ling-Qi Yan. 2022. SpongeCake: A Layered Microflake Surface Appearance Model. *ACM Trans. Graph.* 42, 1, Article 8 (sep 2022), 16 pages. <https://doi.org/10.1145/3546940>
- Andrea Weidlich and Alexander Wilkie. 2007. Arbitrarily Layered Micro-Facet Surfaces. In *GRAPHITE 2007 (Perth, Australia)*. ACM, 171–178. https://www.cg.tuwien.ac.at/research/publications/2007/weidlich_2007_almfms/
- Hongzhi Wu, Julie Dorsey, and Holly Rushmeier. 2011. Physically-based interactive bi-scale material design. *ACM Trans. Graph.* 30, 6 (dec 2011), 1–10. <https://doi.org/10.1145/2070781.2024179>
- Tizian Zeltner, Brent Burley, and Matt Jen-Yuan Chiang. 2022. Practical Multiple-Scattering Sheen Using Linearly Transformed Cosines. In *ACM SIGGRAPH 2022 Talks (Vancouver, BC, Canada) (SIGGRAPH '22)*. Association for Computing Machinery, New York, NY, USA, Article 7, 2 pages. <https://doi.org/10.1145/3532836.3536240>

A SPHERICAL CAP SHADOW

As illustrated in Figure 3, we want to compute the area σ_s of the shadow cast by a unit spherical micrograin from a direction $\mathbf{i} = (\sin \theta_i, 0, \cos \theta_i)$ onto a plane \mathcal{H} parallel to the surface and at a height h above it. This is the area of a spherical cap shadow.

The cast shadow (gray region in Figure 3-bottom) lies inside an ellipse $\mathcal{E}_h(\mathbf{i})$ and outside of a circle \mathcal{C}_h . $\mathcal{E}_h(\mathbf{i})$ is obtained by intersecting the cylinder of radius 1 and direction \mathbf{i} tangent to the unit sphere with \mathcal{H} . $\mathcal{E}_h(\mathbf{i})$ has a unit minor axis, a major axis of $\cos^{-1} \theta_i$, and a center $\mathbf{e} = (e_x, 0, h)$ where $e_x = \tan \theta_i \cos \theta_m$ is obtained by intersecting \mathcal{H} with a ray of direction \mathbf{i} starting from the origin (see Figure 3-middle). Intersecting the unit sphere with \mathcal{H} yields the circle \mathcal{C}_h , of radius $\sin \theta_m$ since $h = \cos \theta_m$. The projected area of the cast shadow is then computed as $\sigma_s = \sigma_e - \sigma_c - \sigma_t$, where σ_e , σ_c and σ_t are the projected areas of the ellipse sector, the circle sector and the triangle pair, shown in Figure 3-bottom with a red outline, a purple region and a yellow region respectively.

In order to compute these areas, we first locate the pair of points $\mathbf{q}_{i\pm} = (q_x, \pm q_y, h) = (\sin \theta_m \cos \phi_q, \pm \sin \theta_m \sin \phi_q, \cos \theta_m)$ where \mathcal{C}_h is tangent to \mathcal{E}_h . By construction, we have $\mathbf{i} \cdot \mathbf{q}_{i\pm} = 0$, yielding:

$$\cos \phi_q = -\frac{1}{\tan \theta_m \tan \theta_i}.$$

Note that $\mathbf{q}_{i\pm}$ exists iff $\cos \theta_m \leq \sin \theta_i$, otherwise there is no shadow (i.e., $\sigma_s = 0$) as the spherical cap is fully visible from direction \mathbf{i} .

The areas of the triangle pair and circle sector are then given by:

$$\sigma_t = |e_x q_y| \quad \text{and} \quad \sigma_c = (\pi - \phi_q) \sin^2 \theta_m.$$

To compute σ_e , we first apply a transformation that maps \mathcal{E}_h to the unit circle, which is achieved by a scaling of $\cos \theta_i$ along the X axis. The area of the ellipse sector is then given by the area of a unit circle sector corrected by the Jacobian $\cos^{-1} \theta_i$ of the scaling:

$$\sigma_e = \frac{\arccos(-(q_x - e_x) \cos \theta_i)}{\cos \theta_i}.$$

B OVERLAPPING SPHERICAL CAP SHADOWS

When a micrograin casts two shadows from directions \mathbf{i} and \mathbf{o} at a height h (i.e., when $\sigma_s(\mathbf{i}, h) \neq 0$ and $\sigma_s(\mathbf{o}, h) \neq 0$), there is the possibility that these shadows overlap. We then want to compute the corresponding overlapping area $\sigma_\cap(\mathbf{i}, \mathbf{o}, h)$ to properly account for retro-reflection effects (as done in Equation 18).


(a) Cylinders intersect at $\mathcal{E}_n(\mathbf{i}, \mathbf{o})$ (b) $\mathcal{E}_n(\mathbf{i}, \mathbf{o})$ and \mathcal{L}_h intersect at \mathbf{p}_+

Fig. 19. Geometry of the shadow overlap region. (a) The cylinders oriented along \mathbf{i} and \mathbf{o} (in red and green respectively) and tangent to the unit spherical micrograin (not shown) intersect in an ellipse $\mathcal{E}_n(\mathbf{i}, \mathbf{o})$, which lies in a plane \mathcal{P} . (b) We then compute several intersections with the plane \mathcal{H} at height h : with \mathcal{P} to get the line \mathcal{L}_h , with the unit sphere to get \mathcal{C}_h , and with the cylinders from (a) to get $\mathcal{E}_h(\mathbf{i})$ and $\mathcal{E}_h(\mathbf{o})$. The shadow overlap region (in gray) is then found by using points where $\mathcal{E}_h(\mathbf{i})$ and $\mathcal{E}_h(\mathbf{o})$ are tangent to \mathcal{C}_h , and the point \mathbf{p}_+ at the intersection of \mathcal{L}_h and $\mathcal{E}_n(\mathbf{i}, \mathbf{o})$.

In order to identify whether there is an overlap, we introduce the ellipse $\mathcal{E}_n(\mathbf{i}, \mathbf{o})$, which results from the intersection of the cylinders of radius 1 and direction \mathbf{i} and \mathbf{o} , both tangent to the unit sphere (see Figure 19a). $\mathcal{E}_n(\mathbf{i}, \mathbf{o})$ lies on the plane $\mathcal{P} = \text{span}(\mathbf{h}, \mathbf{b})$ where \mathbf{h} is the halfway vector and $\mathbf{b} = \frac{\mathbf{i} \times \mathbf{o}}{\|\mathbf{i} \times \mathbf{o}\|}$. Intuitively, this leads to three cases: 1) if spherical cap shadows lie on different sides of \mathcal{P} , they do not overlap; 2) if they lie on the same side of \mathcal{P} , one fully overlaps the other; and 3) if they both cross \mathcal{P} , they partially overlap.

Formally, we rely on the line $\mathcal{L}_h = \mathcal{P} \cap \mathcal{H}$ and on silhouette points $\mathbf{q}_{i\pm}$ and $\mathbf{q}_{o\pm}$ (computed as in Appendix A) to test for these three cases, as illustrated in Figure 4. Case 1) occurs when $\mathbf{q}_{i\pm}$ and $\mathbf{q}_{o\pm}$ are on different sides of \mathcal{L}_h , yielding $\sigma_n(\mathbf{i}, \mathbf{o}, h) = 0$. Case 2) happens when $\mathbf{q}_{i\pm}$ and $\mathbf{q}_{o\pm}$ are on the same side of \mathcal{L}_h , resulting in $\sigma_n(\mathbf{i}, \mathbf{o}, h) = \min(\sigma_s(\mathbf{i}, h), \sigma_s(\mathbf{o}, h))$. In case 3), the shadow boundaries intersect at a point $\mathbf{p}_+ \in \mathcal{P}$, which lies at the intersection of \mathcal{L}_h and $\mathcal{E}_n(\mathbf{i}, \mathbf{o})$ by construction (see Figure 19b). Writing $\mathbf{p}_+ = p_h \mathbf{h} + p_b \mathbf{b}$, we have:

$$\begin{aligned} h &= p_h(\mathbf{h} \cdot \mathbf{n}) + p_b(\mathbf{b} \cdot \mathbf{n}), \\ 1 &= p_h^2 + [1 - (\mathbf{h} \cdot \mathbf{o})^2] p_b^2, \end{aligned}$$

where the first line derives from $\mathbf{n} \cdot \mathbf{p}_+ = h$ since $\mathcal{L}_h \subset \mathcal{H}$, and the second line is the equation of $\mathcal{E}_n(\mathbf{i}, \mathbf{o})$ in \mathcal{P} . Solving for both equations yields a pair of intersection points, and we denote by \mathbf{p}_+ the one on the shadows side (i.e., with $p_h < 0$). Similarly, we enforce \mathbf{q}_{i+} and \mathbf{q}_{o+} to be always on the shadow overlap region.

The area of the shadow overlap region is then decomposed as $\sigma_n = \sigma_{n_t} + \sigma_{n_i} + \sigma_{n_o} - \sigma_{n_c}$, with the component areas illustrated in the inset of Figure 4. The area σ_{n_t} of the triangle formed by points $\{\mathbf{p}_+, \mathbf{q}_{o+}, \mathbf{q}_{i+}\}$ (in yellow) is given by:

$$\sigma_{n_t} = \frac{1}{2} \|(\mathbf{q}_{i+} - \mathbf{p}_+) \times (\mathbf{q}_{o+} - \mathbf{p}_+)\|.$$

The area σ_{n_c} of the circle segment (in purple) is obtained by subtracting the area of the triangle formed by points $\{\mathbf{0}, \mathbf{q}_{o+}, \mathbf{q}_{i+}\}$

from the area of the circle sector formed by the same points, yielding:

$$\sigma_{n_c} = \arccos\left(\frac{\mathbf{q}_{i+} \cdot \mathbf{q}_{o+}}{\sin(\theta_m)^2}\right) \sin(\theta_m)^2 - \frac{\|\mathbf{q}_{i+} \times \mathbf{q}_{o+}\|}{2}.$$

The areas σ_{n_i} and σ_{n_o} (in red and green respectively) correspond to ellipse segment areas. To compute them, we first apply a transform matrix M that turns a unit circle into an ellipse with the same axes.

Taking the example of σ_{n_i} , we write $M = \begin{bmatrix} \cos \phi_i & -\sin \phi_i \\ \sin \phi_i & \cos \phi_i \end{bmatrix}$ with $\det M = \frac{1}{\cos(\theta_i)}$, yielding the sought-for area:

$$\sigma_{n_i} = \det M \left(\arccos(\mathbf{q}_M \cdot \mathbf{p}_M) - \frac{\|\mathbf{q}_M \times \mathbf{p}_M\|}{2} \right),$$

where $\mathbf{q}_M = M^{-1}(\mathbf{q}_{i+} - R_{\phi_i} \mathbf{e})$ and $\mathbf{p}_M = M^{-1}(\mathbf{p}_+ - R_{\phi_i} \mathbf{e})$, with R_{ϕ_i} the rotation of angle ϕ_i in the plane \mathcal{H} , and \mathbf{e} the ellipse center.

C PORE VISIBILITY

The definition of the pore visibility term V_p in Equation 23 involves evaluating G^{dist} at $h' = 0$, which itself relies on shadow areas and their overlap (see Equation 19) evaluated at surface level.

The shadow area σ_s at a zero height has a simpler form. Indeed, the ellipse $\mathcal{E}_0(\mathbf{i})$ is centered at the origin and the circle \mathcal{C}_0 has a radius of 1. As a result, $\mathbf{q}_{i\pm} = (0, \pm 1, 0)$, yielding $\sigma_t(\mathbf{i}, 0) = 0$, $\sigma_c(\mathbf{i}, 0) = \frac{\pi}{2}$ and $\sigma_e(\mathbf{i}, 0) = \frac{\pi}{2 \cos \theta_i}$, hence $\sigma_s(\mathbf{i}, 0) = \frac{\pi}{2} \left(\frac{1}{\cos \theta_i} - 1 \right)$. We observe that $\sigma(\mathbf{i}, 0) = \sigma(\mathbf{n}, 0) + \sigma_s(\mathbf{i}, 0) = \frac{\pi}{2} \left(1 + \frac{1}{\cos \theta_i} \right)$ corresponds to the area of an elliptical lune, as defined in the Ellipsoid NDF model of Dong et al. [2016].

The area of shadow overlap σ_n also has a slightly simpler form, since the line \mathcal{L}_0 now goes through the origin.