



HAL
open science

Anisotropic Specular Image-Based Lighting Based on BRDF Major Axis Sampling

Giovanni Cocco, Cédric Zanni, Xavier Chermain

► **To cite this version:**

Giovanni Cocco, Cédric Zanni, Xavier Chermain. Anisotropic Specular Image-Based Lighting Based on BRDF Major Axis Sampling. Computer Graphics Forum, In press, 10.1111/cgf.15233 . hal-04766640

HAL Id: hal-04766640

<https://inria.hal.science/hal-04766640v1>

Submitted on 5 Nov 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.

Copyright

Anisotropic Specular Image-Based Lighting Based on BRDF Major Axis Sampling

Giovanni Cocco^{1,2} , Cédric Zanni² , and Xavier Chermain² 

¹University of Milan, Italy

²Université de Lorraine, CNRS, Inria, LORIA, France




Figure 1: A circularly brushed metallic Utah teapot is rendered. The anisotropic specular material is illuminated through image-based lighting. From left to right: our method, the bent normal technique [Rev11, LG18], and the reference image. Our technique uses 16 BRDF samples positioned along the major axis of the BRDF lobe, which improves the rendering quality compared to the bent normal technique that uses only one sample. Rendering times at a resolution of 1920×1080 with $8 \times$ MSA are provided below the images.

Abstract

Anisotropic specular appearances are ubiquitous in the environment: brushed stainless steel pans, kettles, elevator walls, fur, or scratched plastics. Real-time rendering of these materials with image-based lighting is challenging due to the complex shape of the bidirectional reflectance distribution function (BRDF). We propose an anisotropic specular image-based lighting method that can serve as a drop-in replacement for the standard bent normal technique [Rev11]. Our method yields more realistic results with a 50% increase in computation time of the previous technique, using the same high dynamic range (HDR) pre-integrated environment image. We use several environment samples positioned along the major axis of the specular microfacet BRDF. We derive an analytic formula to determine the two closest and two farthest points from the reflected direction on an approximation of the BRDF confidence region boundary. The two farthest points define the BRDF major axis, while the two closest points are used to approximate the BRDF width. The environment level of detail is derived from the BRDF width and the distance between the samples. We extensively compare our method with the bent normal technique and the ground truth using the GGX specular BRDF.

CCS Concepts

• Computing methodologies → Reflectance modeling;

1. Introduction

Anisotropic appearances refer to materials that exhibit different visual appearances when viewed from various directions. This directional dependence on appearance is caused by the underlying anisotropic microstructure of the material. Natural fibers such as hair, fur, and woven fabrics like silk exhibit anisotropic appearances

due to their oriented threads or fibers. Engineering materials such as polished wood, brushed metals like aluminum or steel, and 3D prints also display noticeable anisotropy.

Image-based lighting (IBL) renders these materials in real time using an omnidirectional representation of lighting represented as an image. Typically, the image, commonly referred to as the envi-

ronment map, is projected onto a geometry that encompasses the scene, such as a cube or a sphere. A highly detailed environment map provides greater realism, recalling the complex lighting environments we experience in real life. Our objective is to achieve real-time rendering with this type of lighting.

Environment map integration. Computing the light scattering equation requires integrating the environment map while considering the bidirectional reflectance distribution function (BRDF) shape. In other words, the BRDF is a convolution filter over the environment map. This article focuses on specular appearances modeled by an anisotropic microfacet BRDF, using the Smith assumption [Hei14]. This choice is motivated by the widespread adoption of this model in the industry [KC17, LG18]. The integrand involves at least seven parameters: two polar angles and two azimuthal angles for the view and light directions, as well as three parameters to represent the anisotropic roughness of the surface. The precomputation of the integral for all the parameters results in excessively large lookup tables. Additionally, the high dimensionality makes the calculation of light scattering challenging for real-time rendering.

Split-integral. The split-integral approximation [Got12, Laz13, Kar13] reduces the problem’s dimensionality. This technique targets specular isotropic materials and allows for real-time rendering. It can also handle anisotropic appearances using a method that bends the normal toward the anisotropy direction to create a convincing and plausible brushed effect [Rev11, LG18]. These techniques are efficient as they only use one BRDF sample to evaluate a pre-filtered version of the environment map. Consequently, they are used in real-time renderers [McA15, LG18, Goo23]. The bent normal technique [Rev11] offers commendable computational efficiency, but it may not always produce the most realistic results (Figure 1).

Our method. Our work aims to provide a drop-in replacement for the bent normal technique, improving accuracy, while maintaining high-performance rendering. Our method is based on the split-integral approximation. To improve quality, we use several filtered environment map samples – typically sixteen – instead of one. This strategy is inspired by anisotropic texture filtering, where several samples along the anisotropy direction are averaged.

We place the samples used for the environment evaluations on what we define as the BRDF major axis (Figure 2). Doing so ensures that the samples are spatially coherent, which in turn reduces cache misses, thus enabling real-time rendering. In addition, our method allows the rendered images to be noise and artifact-free. We provide an analytical formula to compute the BRDF major axis. Its derivation is based on a geometric approach, using the distance between the reflection direction and the points on a BRDF confidence region boundary approximation. We use the two closest and two farthest points on the confidence region boundary to select the pre-filtered environment map’s level of detail (LOD) and sample it.

We compare our method with the bent normal technique and a converged reference image. Our results are generally closer to the reference image. Notably, the computation time requires approximately 1.5 times the rendering time of the previous technique.




Figure 2: Plots of the anisotropic specular BRDF and its major axis (orange) for three different sets of parameters. Using our method, the major axis is computed and sampled on the fly during rendering.

To summarize, our contributions are:

- An analytical formula for approximating the major axis of an anisotropic specular BRDF.
- A new anisotropic specular image-based lighting technique.

We provide a glsl reference implementation at <https://www.shadertoy.com/view/4fV3W3>. This implementation also demonstrates how our method can replace the bent normal technique.

Our integral estimator is biased because it relies on the split integral technique and samples the BRDF only along its major axis. Consequently, it should be used primarily for high-performance rendering, such as video games, where convergence to the ground truth is not essential. In such cases, we prioritize low rendering times over unbiased rendering.

The remainder of the article is organized as follows. In Section 2, we present related work on specular image-based lighting. We review the split-integral technique in Section 3. We review anisotropic roughness parameterization and the bent normal technique in Section 4. Section 5 describes our new specular anisotropic IBL method. Section 6 compares our results with the previous method and the ground truth. Finally, we discuss the limitations and future work in Section 7.

2. Related work

In computer graphics, anisotropy is significant in several areas, including texture filtering [Hec89], image-based rendering [YCC*22], and geometry processing [CK11]. This section provides background on specular image-based lighting. For a thorough review, refer to Akenine-Möller et al. [AMHH*18].

Greene [Gre86] used the concept of blurring environment maps to simulate the appearance of rough metals. Although this approach yields plausible results, the actual shape of the BRDF should be used as a filtering kernel. Isotropic BRDFs have at least five parameters: the view and light directions (which include two polar and two azimuthal angles) and roughness. This high dimensionality hinders the use of tabulated data for real-time rendering. Consequently, the preintegration of the environment map employs a subset of the BRDF formula to reduce the number of dimensions.

Split-integral approximation. Phong [MLH02], Gaussian [GKD07], and GGX [Kar13] lobes were used as kernels to compute pre-filtered environment maps. To account for the

other terms of the BRDF, Gotanda [Got12], Lazarov [Laz13], and Karis [Kar13] independently proposed a split-integral approximation of the rendering equation. This approach aims to pre-filter the environment map and the BRDF separately. The approximation is also known as the split-sum when the integral is rewritten with a Monte Carlo estimator [Kar13]. The split-integral approximation is usually applied to isotropic microfacet BRDFs, so the BRDF part – also referred to as the hemispherical-directional reflectance of the specular term – generally includes the Fresnel term, the masking-shadowing function, the normal distribution function (NDF), and normalization factors [Hei14]. Typically, the Fresnel term used is the Schlick’s approximation [Sch94]. The BRDF part is either precomputed and stored in a lookup table [Got12, Kar13] or approximated with simple analytical formulas [Nar14, PI15]. Kneiphof et al. [KGK19] and Fdez-Agüera [FA19] modified the original split-integral approach to include iridescent and multiple-scattering effects, respectively.

One environment sample. Lagarde and de Rousiers [Ldr14] show that the dominant direction of an isotropic BRDF is the reflection direction shifted slightly toward the geometric normal when the view is grazing and the roughness is high. They recommend using the dominant direction instead of the reflection direction when evaluating the environment map. Anisotropic BRDFs are handled using the bent normal technique [Rev11]. This technique uses one prefiltered environment sample defined by a modified reflected direction. The method bends the normal – used to reflect the view – toward the anisotropic direction. This empirical approach is commonly used in the industry because it is stable, computationally efficient, and provides plausible results, even if they are not close to the ground truth [McA15, LG18].

Several environment samples. Multiple radially symmetric lobes can be used as basis functions to approximate the shape of the BRDF [KM00, PI15]. These methods distribute the lobes uniformly along the hemisphere’s meridian and are dedicated to isotropic BRDFs. Additionally, they precompute the lobes’ data for several view angles and roughness values. Our method also uses several environment samples but is designed for both isotropic and anisotropic microfacet BRDFs. The lobes are distributed according to the shape of the BRDF, which does not follow the meridian for anisotropic NDFs. Furthermore, the positions of the lobes are computed on the fly, reducing memory requirements.

Importance sampling techniques. Real-time image-based lighting of anisotropic specular materials can also use Monte Carlo integration with importance sampling [YYTM19, WB20]. Although this technique is simple, it remains computationally expensive and can lead to noisy results. Křivánek and Colbert [KC08] proposed filtered importance sampling to remove the noise by using pre-filtered environment samples. Their method is simple but is limited to slightly anisotropic reflections. Dupuy et al. [DHI*13] extended filtered importance sampling to handle anisotropic materials. Their uniform sampling scheme ensures noise-free results but still requires a substantial number of samples to accurately reconstruct anisotropic lobes (≥ 256) and is restricted to the Beckmann distribution. Our approach targets the GGX distribution and aims to achieve accurate results with a lower number of samples (~ 16).




Figure 3: Illustration of the BRDF f (left) and the omnidirectional incoming radiance L_i (right).

3. Isotropic specular image-based lighting

This section reviews the split integral technique for isotropic specular image-based lighting, as our method is based on it.

The objective is to compute the outgoing radiance L_o from an outgoing light direction $\omega_o = (o_x, o_y, o_z)^T$ and a surface position, which is omitted here for clarity. At this surface position, the microsurface statistics are represented by a roughness parameter $\sigma \in \mathbb{R}_{>0}$ which denotes the standard deviation of the microsurface slope. The scattering equation defines the outgoing radiance

$$L_o(\omega_o, \sigma) = \int_{S^2} f(\omega_o, \omega_i, \sigma) L_i(\omega_i) (\omega_g \cdot \omega_i) d\omega_i. \quad (1)$$

The directions are defined in the tangent space, whose normal $\omega_g = (0, 0, 1)^T$ is the geometric normal. The integral domain is the unit sphere S^2 . We only consider opaque materials, so the BRDF f (Figure 3, left) and the incoming radiance L_i (Figure 3, right) equals zero for incoming rays ω_i from below the surface ($\omega_g \cdot \omega_i < 0$). We assume that the surface is locally flat at the BRDF evaluation point.

The roughness is the parameter of the specular microfacet BRDF [Hei14]

$$f(\omega_o, \omega_i, \sigma) = \frac{F(\omega_o, \omega_i) G_2(\omega_o, \omega_i, \sigma) D(\omega_h, \sigma)}{4|\omega_g \cdot \omega_i| |\omega_g \cdot \omega_o|},$$

where F is the Fresnel reflection coefficient, G_2 is the Smith height-correlated masking-shadowing function, D is the normal distribution function (NDF), and ω_h is the half-vector $(\omega_o + \omega_i) / \|\omega_o + \omega_i\|$. We assume that the NDF is the ground glass unknown (GGX) distribution in this article, as it is the industry standard [KC17, LG18]. Please see Heitz 2014 [Hei14] for the explicit formulas of each term.

The split-integral approximation of the scattering equation [Got12, Laz13, Kar13] is used for high-performance image-based lighting. The approximation is split into two integrals, the first involving the environment map and the second the BRDF. This rewriting aims to evaluate separately pre-filtered versions of the BRDF \bar{f} and the environment map \bar{L}_i . The split-integral approximation is

$$\begin{aligned} L_o(\omega_o, \sigma) &\approx \int_{S^2} K(\omega_i, \check{\omega}_g, \sigma) L_i(\omega_i) d\omega_i \\ &\quad \int_{S^2} f(\omega_o, \omega_i, \sigma) (\omega_g \cdot \omega_i) d\omega_i \\ &\approx \bar{L}_i(\check{\omega}_g, \sigma) \bar{f}(\omega_o, \sigma), \end{aligned} \quad (2)$$

where $\check{\omega}_g = 2(\omega_g \cdot \omega_o) \omega_g - \omega_o$ is the reflection direction. We now describe the two integrals \bar{L}_i and \bar{f} separately (Figure 4).




Figure 4: Illustration of the split integral. Left: The environment part \bar{L}_i is a convolution of the image-based light L_i with a radially symmetric kernel K . Right: The BRDF part \bar{f} assumes constant incoming light and integrates the BRDF for a given view angle and roughness.

Environment pre-filtering. The first integral \bar{L}_i filters the environment map with a low-pass filter K , defining a distribution on the sphere S^2 , i.e., $\int_{S^2} K(\omega_i, \tilde{\omega}_g, \sigma) d\omega_i = 1$. This property guarantees that the split-integral is exact when a constant lighting environment is used, i.e., $L_i(\omega_i) = 1$.

The typical approach for the filter K is to use a radially symmetric filter whose shape is close to that of the BRDF. As the shape of the BRDF is mainly defined by its NDF, the filter is commonly the GGX distribution: $K(\omega_i, \tilde{\omega}_g, \sigma) = D(T_{\tilde{\omega}_g} \omega_i, \sigma)(\tilde{\omega}_g \cdot \omega_i)$, where $T_{\tilde{\omega}_g}$ is the matrix that transforms tangent space directions into reflection space directions. The reflection space is defined by the orthonormal basis whose normal is the reflection direction $\tilde{\omega}_g$.

In the literature, this change of basis is usually described by setting the view direction and the geometric normal to the reflection vector: $\omega_o := \omega_g := \tilde{\omega}_g$ [Kar13, AMHH*18]. Note that the GGX filter satisfies the unit integral property [Hei14], is radially symmetric, and provides the best approximation when $\omega_o = \omega_i = \omega_g$. For these reasons, we also use this filter.

The integral \bar{L}_i is stored in a three-dimensional lookup table, as computing it on the fly is expensive. A cube map with a fixed number N_c of LODs is typically used. As the LOD l increases, the roughness σ also increases. However, a linear increase in the slope standard deviation σ does not result in a perceptually linear increase in roughness. Appearance designers typically use a linear roughness parameter $\alpha_{\text{lin}} \in (0, 1]$, defined as $\alpha_{\text{lin}} := \sqrt{\alpha} := \sqrt{\sqrt{2}\sigma}$ [Bur12, Hei14]. The α roughness parameter is a common way to parameterize the GGX distribution. To linearly map a LOD $l \in [0, N_c - 1]$ to a linear roughness α_{lin} , we use

$$\alpha_{\text{lin}} = \frac{l}{N_c - 1}. \quad (3)$$

In our implementation, we use Lagarde and de Rousiers's pre-filtered importance sampling algorithm [LdR14] to precompute \bar{L}_i for different roughnesses, but other techniques can also be used [MS16].

BRDF pre-filtering. The second integral, \bar{f} , is the hemispherical-directional reflectance satisfying $\int_{S^2} f(\omega_o, \omega_i, \sigma)(\omega_g \cdot \omega_i) d\omega_i \leq 1$. This integral can be precomputed in a lookup table [Kar13] or approximated by an analytical function [PI15]. In both cases, the

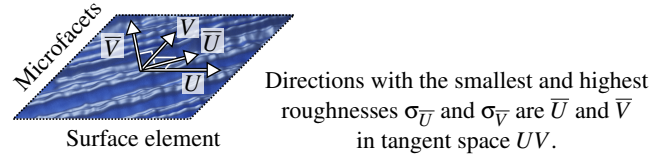


Figure 5: Illustration of an anisotropic heightfield in a surface element. The anisotropic direction \bar{U} may differ from the axes of the UV space.

Schlick formulation [Sch94] is used for the Fresnel term F . In our method, we use the Pesce analytical function [PI15].

4. Anisotropic specular image-based lighting

This section begins by describing the existing parameterizations of anisotropic roughness. Next, we explain the anisotropic scattering equation and how to approximate it using the bent normal technique.

4.1. Anisotropic parameterization

For each surface position, i.e., element, an anisotropic material is modeled by an anisotropic roughness. Each surface element is associated with a UV tangent space whose first and second 2D axes are U and V , respectively. Anisotropic roughness is characterized by two roughness values, $\sigma_{\bar{U}}$ and $\sigma_{\bar{V}}$, defined along two orthogonal axes, \bar{U} and \bar{V} . Note that the roughness axes \bar{U} and \bar{V} can differ from the tangent space axes U and V (Figure 5). Similarly to the isotropic case, the roughness values $\sigma_{\bar{U}}$ and $\sigma_{\bar{V}}$ represent the standard deviations of the microsurface slopes along \bar{U} and \bar{V} , respectively.

Patry [Pat20] proposes using a covariance matrix Σ , specifically a 2×2 symmetric matrix, to encode anisotropy. This choice is motivated by a key property of covariance matrices: their components are linearly filterable. This property is crucial for real-time anti-aliasing of spatially varying anisotropic roughness. Consequently, we adopt this representation for roughness filtering. Note that the eigenvectors and eigenvalues of the covariance matrix Σ correspond to the orthogonal directions \bar{U} and \bar{V} , and the slope variances $\sigma_{\bar{U}}^2$ and $\sigma_{\bar{V}}^2$, respectively.

Setting the components of the covariance matrix Σ is not intuitive, and designers are more accustomed to the previously described perceptual isotropic roughness α_{lin} . To retain this parameter and incorporate anisotropy, Kulla and Conty [KC17] introduced a new parameter $\eta \in (-1, 1)$ to control the amount of anisotropy. We use this parameterization (α_{lin} and η) in our user interface and results section. The mapping between these two parameters and the covariance matrix Σ is

$$\Sigma = [\bar{U} \ \bar{V}] \begin{bmatrix} \sigma_{\bar{U}}^2 & 0 \\ 0 & \sigma_{\bar{V}}^2 \end{bmatrix} \begin{bmatrix} \bar{U}^T \\ \bar{V}^T \end{bmatrix}$$

with

$$\sigma_{\bar{U}} = \frac{\alpha_{\text{lin}}^2(1+\eta)}{\sqrt{2}} \quad \text{and} \quad \sigma_{\bar{V}} = \frac{\alpha_{\text{lin}}^2(1-\eta)}{\sqrt{2}}.$$

4.2. Anisotropic light scattering

The anisotropic scattering equation is

$$L_o(\omega_o, \Sigma) = \int_{S^2} f(\omega_o, \omega_i, \Sigma) L_i(\omega_i)(\omega_g \cdot \omega_i) d\omega_i,$$

where the only change from Equation 1 is the switch from σ to Σ .

4.3. Bent normal technique

The introduction of anisotropy theoretically forbids using the split integral technique for real-time anisotropic IBL. Still, the bent normal technique [Rev11, McA15, LG18] employs it for anisotropic rendering due to its computational efficiency. The idea is to use the pre-filtered terms \bar{L}_i and \bar{f} with modified parameters based on the anisotropic roughness Σ :

$$L_o(\omega_o, \Sigma) \approx \bar{L}_i(g(\check{\omega}_g, \omega_g, \Sigma), h_2(\Sigma)) \bar{f}(\omega_o, h_1(\Sigma)). \quad (4)$$

There are two differences compared to Equation 2. For the evaluation of \bar{L}_i , the mapping $g(\check{\omega}_g, \omega_g, \Sigma)$ bends the reflection direction $\check{\omega}_g$ based on the anisotropy Σ . For both pre-filtered terms, the mappings $h_1(\Sigma)$ and $h_2(\Sigma)$ convert anisotropic roughness to isotropic roughness. We use the mappings from Lagarde 2018 [LG18] to implement the bent normal technique. The results give plausible anisotropic appearances but are still far from the ground truth (Figure 1).

5. Our method

When the reflection direction is not bent, using the pre-filtered environment map \bar{L}_i constrains the shape of the highlight to be radially symmetric. The other term, \bar{f} , does not change the shape of the highlight, as it is a scaling factor independent of the lighting environment. The problem is that the shape of the BRDF is not radially symmetric with anisotropic roughness or isotropic roughness when the view is grazing, as observed in related work [LdR14, PI15]. Our method's key idea is to reconstruct the shape of the BRDF with N evaluations of \bar{L}_i , i.e., radially symmetric primitives. To do this, we replace the pre-filtered environment part of Equation 4 with an average:

$$L_o(\omega_o, \Sigma) \approx \frac{\bar{f}(\omega_o, h_1(\Sigma))}{N} \sum_{j=0}^{N-1} \bar{L}_i(\check{\omega}_j, h_3(\Sigma, \omega_o)). \quad (5)$$

We retain the bent normal mapping h_1 for the pre-filtered BRDF part. For the pre-filtered environment part, we introduce a new mapping h_3 , with the view direction ω_o and the anisotropic roughness Σ as inputs. This mapping is described at the end of our method description. The samples $\check{\omega}_j$ should be placed in the high-density area of the BRDF, i.e., its confidence region. As for anisotropic texture filtering, our strategy is to position them along a major axis. In our case, this axis is the major axis of the BRDF confidence region (Figure 2). Unlike texture filtering, the major axis of the BRDF is not a straight line, as the BRDF is defined in the spherical domain. In the following sections, we will define an approximation of the BRDF confidence region (Section 5.2), its major axis (Section 5.4), its width (Section 5.5), and the mapping h_3 needed to select the environment LOD (Section 5.7).




Figure 6: Illustration of the slope, normal, and reflected spaces. The vectors are of unit length.

5.1. Slope, normal, and reflected spaces

Our definition of the confidence region and the computation of its major axis require defining the slope, normal, and reflected space used in microfacet theory. The NDF models the distribution of numerous microfacets at a surface position (Figure 5). Each microfacet has a normal $\omega = (x_\omega, y_\omega, z_\omega)^T \in S^2$, which can equivalently be defined by a slope $\check{\omega} = (x_{\check{\omega}}, y_{\check{\omega}})^T \in \mathbb{R}^2$ (Figure 6, left). The relations between a slope and a normal are given by:

$$\omega := \frac{1}{\sqrt{x_{\check{\omega}}^2 + y_{\check{\omega}}^2 + 1}} \begin{pmatrix} -x_{\check{\omega}} \\ -y_{\check{\omega}} \\ 1 \end{pmatrix} \quad \text{and} \quad \check{\omega} := \begin{pmatrix} -x_\omega/z_\omega \\ -y_\omega/z_\omega \end{pmatrix}.$$

As we target glossy materials, all the microfacets are perfectly specular, i.e., they act as mirrors. We define the direction

$$\check{\omega} := 2(\omega \cdot \omega_o)\omega - \omega_o$$

as the microfacet's reflection of the view direction ω_o . We say that the reflected direction $\check{\omega} \in S^2$ is in the reflected space (Figure 6, right).

5.2. Confidence region

We approximate the BRDF confidence region by the NDF confidence region in the reflected space. This choice provides closed-form solutions. The GGX slope distribution $P^{22}(\check{\omega}, \Sigma)$, i.e., the NDF in the slope space, is a Student's t -distribution [RBSM19] whose confidence region boundary is the ellipse (Figure 7(a)):

$$\tilde{\mathcal{E}} := \left\{ \check{\omega}_t := [\bar{U} \quad \bar{V}] \begin{pmatrix} k\sigma_{\bar{U}} \cos(t) \\ k\sigma_{\bar{V}} \sin(t) \end{pmatrix} \middle| t \in [0, 2\pi) \right\}, \quad (6)$$

where the ellipse radii are $k\sigma_{\bar{U}}$ and $k\sigma_{\bar{V}}$. The scale factor k controls the slope density percentage included in the ellipse. The rotation matrix $[\bar{U} \quad \bar{V}]$ represents the orientation of the ellipse (Figure 7(a)). The boundary of the confidence region in normal space is $\mathcal{E} := \{\omega_t | t \in [0, 2\pi)\}$ (Figure 7(b)), and in reflected space is $\tilde{\mathcal{E}} := \{\check{\omega}_t | t \in [0, 2\pi)\}$ (Figure 7(c)). Note that the shapes defined on the sphere S^2 are more complex than the 2D ellipse, which motivated us to start with Equation 6.

5.3. Confidence region extrema

The objective now is to define the major axis of the confidence region in the reflected space. Figure 7 suggests that the major axis of

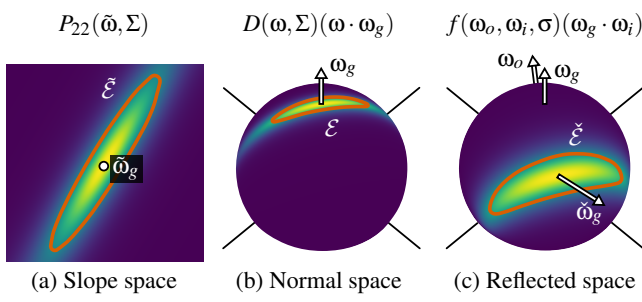


Figure 7: The approximation of the BRDF confidence region boundary (orange) is presented in slope, normal, and reflected space. (a) In slope space, the slope distribution function P^{22} , its confidence region boundary $\tilde{\mathcal{E}}$, and the slope $\tilde{\omega}_g = (0,0)^T$ (white point) are illustrated. (b) In normal space, the cosine-weighted NDF, the confidence region boundary \mathcal{E} , and the geometric normal ω_g are depicted. (c) In reflected space, the cosine-weighted BRDF, its confidence region boundary $\tilde{\mathcal{E}}$, and the reflection direction $\tilde{\omega}_g$ are represented.

the BRDF corresponds to the major axis of the 2D ellipse. However, as shown in Figure 8, this is not always the case. A straightforward example is when the roughness is isotropic. In this situation, the ellipse becomes a circle without a major axis. Nevertheless, even in this case, the confidence region in the reflected space has a major axis, as illustrated in the third row of Figure 8. A microfacet BRDF with isotropic roughness adopts an anisotropic shape when the view direction ω_o is grazing [LdR14].

These observations motivate us to derive the major axis of the confidence region $\tilde{\mathcal{E}}$ in the reflected space. We are interested in identifying the points $\tilde{\omega}_t$ on the boundary of the confidence region $\tilde{\mathcal{E}}$ that are the farthest from and closest to the reflection direction $\tilde{\omega}_g$. We refer to these points as the extrema. Although spherical distance could be used to define them, we use a different distance measure that leads to closed-form solutions. Specifically, we consider the alignment between the center $\tilde{\omega}_g$ and a point $\tilde{\omega}_t$ on the boundary of the confidence region, as defined by the dot product:

$$d(t) := \tilde{\omega}_t \cdot \tilde{\omega}_g.$$

Using the dot product implies that the maximum values identify the closest points and vice versa. We define the closest and farthest points from the center $\tilde{\omega}_g$ as follows:

$$\tilde{\omega}_{\perp 0}, \tilde{\omega}_{\perp 1} = \underset{\tilde{\omega}_t}{\operatorname{argmax}} d(t) \quad \text{and} \quad \tilde{\omega}_{\top 0}, \tilde{\omega}_{\top 1} = \underset{\tilde{\omega}_t}{\operatorname{argmin}} d(t), \quad (7)$$

respectively. These extrema are illustrated in Figure 8. There are always two farthest points, $\tilde{\omega}_{\top 0}$ and $\tilde{\omega}_{\top 1}$, and two closest points, $\tilde{\omega}_{\perp 0}$, $\tilde{\omega}_{\perp 1}$, unless all points are equidistant from the center. Please refer to the supplemental material for the proof and detailed derivation. Closed-form formulas are provided in Appendix A.

5.4. BRDF major axis

Intuitively, we want the major axis of the BRDF confidence region to be the curve that starts at the first maximum extrema, $\tilde{\omega}_{\top 0}$,




Figure 8: The second and third rows show that the ellipse's major axis is not always aligned with the BRDF's major axis (orange dotted line). The two farthest points, $\tilde{\omega}_{\top 0}$ and $\tilde{\omega}_{\top 1}$, and the two closest points, $\tilde{\omega}_{\perp 0}$ and $\tilde{\omega}_{\perp 1}$, from the reflection direction $\tilde{\omega}_g$ (white point) in the reflected space are illustrated in the slope, normal, and reflected spaces. We define the BRDF major axis \mathcal{M} in normal space as the arc of the great circle going from $\omega_{\top 0}$ to $\omega_{\top 1}$.

passes through the center $\tilde{\omega}_g$, and ends at the second maximum extrema, $\tilde{\omega}_{\top 1}$. The BRDF samples, $\tilde{\omega}_j$, should be placed along this curve to maximize the area of the BRDF confidence region covered by the samples.

Figure 8 shows that this curve is a straight line in slope space and a great circle arc passing through the geometric normal ω_g in normal space. Due to the reflection operation, the major axis may take on a more complex shape in the reflected space, as shown in Figures 2 and 8, first row.

These observations motivate us to define the major axis \mathcal{M} in the normal space, which can be achieved by interpolating between the two maximum extrema $\omega_{\top 0}$ and $\omega_{\top 1}$. The interpolated points, i.e., the major axis \mathcal{M} in the normal space, are then reflected to yield the major axis $\tilde{\mathcal{M}}$ in the reflected space. We use normalized vector interpolation (nlerp) as a fast and reasonable approximation of spherical interpolation:

$$\mathcal{M} := \left\{ \omega_\tau := (x_\tau, y_\tau, z_\tau)^T = \operatorname{nlerp}(\omega_{\top 0}, \omega_{\top 1}, \tau) \mid \tau \in [0, 1] \right\},$$

where

$$\operatorname{nlerp}(\omega_0, \omega_1, t) = \frac{(1-t)\omega_0 + t\omega_1}{\|(1-t)\omega_0 + t\omega_1\|}.$$




Figure 9: The maximum extremum, ω_{T0} , is clamped to ω'_{T0} when part of the major axis is below the light-scattering hemisphere. The clamped major axis has a center, $\check{\omega}_c$, that differs from the center, $\check{\omega}_g$, of the unclamped major axis.

Clamped major axis. At grazing angles, one of the farthest extrema occur in the lower hemisphere, i.e., $z_\tau < 0$ (Figure 9, right). We restrict the parameterization domain of \mathcal{M} by confining τ to the interval $[\tau_0, \tau_1]$, where τ_0 and τ_1 are the solutions to $z_\tau = 0$. The closed-form expressions for τ_0 and τ_1 are provided in Appendix B. By applying this restriction, we obtain two new clamped extrema, ω'_{T0} and ω'_{T1} , that prevent the sampling of the BRDF in the lower hemisphere:

$$\begin{aligned} \omega'_{T0} &= \text{nlerp}(\omega_{T0}, \omega_{T1}, \tau'_0), & \tau'_0 &= \text{clamp}(\tau_0, 0, 1), \\ \omega'_{T1} &= \text{nlerp}(\omega_{T0}, \omega_{T1}, \tau'_1), & \tau'_1 &= \text{clamp}(\tau_1, 0, 1). \end{aligned} \quad (8)$$

Figure 9 illustrates how the maximum extremum, ω_{T0} , is clamped to ω'_{T0} in the normal space (Equation 8).

5.5. BRDF width

The BRDF sample radius must be estimated to select the appropriate LOD of the pre-filtered environment map \bar{L}_i (Equation 5). This radius depends on the width of the BRDF. The closest extrema, $\check{\omega}_{\perp 0}$ and $\check{\omega}_{\perp 1}$, provide a good approximation of the BRDF width (Figure 10, left) in most cases. We define the BRDF width as the distance between these extrema:

$$\Delta := \|\check{\omega}_{\perp 0} - \check{\omega}_{\perp 1}\|.$$

With a perfectly grazing view, the closest extrema collapse into a singular point: the reflection direction $\check{\omega}_g$. Consequently, with grazing views, the width Δ approaches zero. To address this issue, we empirically move the closest extrema toward the farthest extremum on the upper hemisphere, following the boundary of the confidence region (Figure 10, right). The displacement is proportional to the clamping of the major axis (see Appendix C). This results in two new points, $\check{\omega}'_{\perp 0}$ and $\check{\omega}'_{\perp 1}$, used to define the BRDF width

$$\Delta' := \|\check{\omega}'_{\perp 0} - \check{\omega}'_{\perp 1}\|.$$

5.6. Major axis sampling

Each BRDF sample ω_j has an associated area, which is illustrated in Figure 11 as a circle because the environment map is




Figure 10: Left: The BRDF width can be computed using the distance Δ between the two closest extrema $\check{\omega}_{\perp 0}$ and $\check{\omega}_{\perp 1}$. Right: The more the view is grazing, the more the distance Δ between $\check{\omega}_{\perp 0}$ and $\check{\omega}_{\perp 1}$ is closed to zero. Consequently, the BRDF width is not well approximated. To solve this issue, we heuristically move these points toward the farthest extremum that remains in the upper hemisphere, following the boundary of the confidence region.




Figure 11: The farthest sample (left) is shifted along the BRDF major axis toward the center ω_c (right) to reduce its associated area outside the confidence region. The displacement amount is proportional to the BRDF width Δ' .

pre-filtered with a radially symmetric filter. We offset the two farthest samples $\check{\omega}'_{T0}$ and $\check{\omega}'_{T1}$ from the clamped major axis center $\omega_c = \frac{\omega'_{T0} + \omega'_{T1}}{\|\omega'_{T0} + \omega'_{T1}\|}$ (Figure 9, right) towards this center to reduce the samples' area outside the confidence region (Figure 11). When the major axis is not clamped, the center corresponds to the reflection direction ($\omega_c = \omega_g$).

We define the two farthest samples from the clamped major axis as

$$\omega_{\wedge i} := \text{nlerp}(\omega_c, \omega'_{Ti}, \zeta), \quad i \in \{0, 1\}$$

where the interpolant ζ is computed based on the BRDF's shape anisotropy. The more isotropic the BRDF lobe, the more the maximum extrema are shifted towards the center $\check{\omega}_c$. See Appendix D for the definition of ζ .

Finally, the BRDF sample positions are computed by interpolation (Figure 12):

$$\omega_j = \text{nlerp}\left(\omega_{\wedge 0}, \omega_{\wedge 1}, \frac{j}{N-1}\right) \text{ where } j \in [0, N-1]. \quad (9)$$

5.7. Pre-filtered environment LOD selection

We must compute the appropriate LOD given the BRDF width Δ' to access the pre-filtered environment map \bar{L}_i . We assume each sample $\check{\omega}_j$ has a radius r_s , as illustrated by the orange circle in Figure 11.




Figure 12: The BRDF sample positions $\check{\omega}_j$ in reflected space ($N = 8$). They are defined by Equation 9 in normal space. These samples $\check{\omega}_j$ follow the BRDF major axis and are used to evaluate the environment map.




Figure 13: Rendering of a sphere with $N = 16$, $k = 1.0$, $\eta = 0.95$ and $\alpha_{lin} = 0.8$. We compare different values of β with the reference. Small values of β lead to ghosting artifacts, while high values over-blur the reflection.

For highly anisotropic BRDFs, the number of samples N might be insufficient to cover the entire confidence area with a sample radius r_s proportional to the BRDF width Δ' , which can result in ghosting artifacts (Figure 13). Therefore, we ensure that the radius is large enough to cover the entire BRDF confidence region by taking into account the distance between samples:

$$r_s = \max \left\{ \frac{\Delta'}{2k}, \frac{\beta \|\check{\omega}'_{T0} - \check{\omega}'_{T1}\|}{2N} \right\}. \quad (10)$$

The first term is divided by the ellipse scaling factor k (Equation 6), as increasing the confidence region area should not affect the BRDF shape. The second term is based on the distance between two adjacent samples $\check{\omega}_j$ (Figure 12), which is approximated by $\frac{\|\check{\omega}'_{T0} - \check{\omega}'_{T1}\|}{N}$. The meta-parameter β scales the width of the GGX filter, and its value was determined empirically. We found that $\beta = 2.3$ yields the best results for any number of samples N (Figure 13).

By rewriting Equation 3,

$$l = \alpha_{lin}(N_c - 1),$$

the LOD l should be computed from the linear perceptual roughness α_{lin} . Our method determines the environment LOD based on the sample radius r_s . We map r_s to the linear perceptual roughness α_{lin} using $\alpha_{lin} = \sqrt{r_s}$. We tested various exponents and found that the square root best matches the reference for isotropic roughnesses. The environment LOD l is defined as

$$l = \sqrt{r_s}(N_c - 1).$$

To complete the mathematical definition of our method (Equation 5), we now define the function h_3 , which maps the view direc-




Figure 14: Rendering of a sphere with $N = 16$, $\eta = 0.8$ and $\alpha_{lin} = 1.0$. We compare different values of k with a reference. A small value of k fails to provide an anisotropic appearance, while large values may produce artefacts.

tion ω_o and the anisotropic roughness Σ to the isotropic roughness σ needed for pre-filtered environment LOD selection. By rewriting the previous equality $\alpha_{lin} = \sqrt{r_s}$, we obtain

$$h_3(\Sigma, \omega_o) = r_s / \sqrt{2}.$$

Note that the sample radius r_s (Equation 10) depends on the view direction ω_o and the anisotropic roughness Σ .

5.8. Choosing the size of the confidence region

The size of the confidence region is controlled by the scale factor k (Equation 6). Our method approximates the BRDF confidence region using the reflected slope distribution confidence region \mathcal{E} . As the scale factor k increases, the slopes of the microfacets at the boundary of the confidence region become steeper. If k is high and the view direction ω_o is grazing, the microfacets at the boundary of the confidence region may become backfacing ($\omega_o \cdot \omega_t < 0$). In such cases, reflections may still be performed even when they should not be. Additionally, our approximation can result in reflected boundary normals pointing below the hemisphere ($\omega_g \cdot \check{\omega}_t < 0$, Figures 9 and 10). For an interactive visualization of the approximated BRDF confidence region boundary, please see the accompanying video. Based on this analysis, we avoid using excessively high values for the scale factor k to prevent backfacing microfacets and reflected normals below the hemisphere. Consequently, we set the scale factor to one ($k = 1$). We ensured there were no artifacts with our renderer (Figure 14) and the confidence region visualizer. This visualizer is available at https://github.com/xavierchermain/brdf_major_axis and can also be used to visualize the major axis and the BRDF samples $\check{\omega}_j$.

6. Results

We compare our technique with the bent normal technique [LG18] and Dupuy et al. 2013 [DHI*13]. We use Monte Carlo integration with 65,536 samples per pixel to provide a ground truth reference. This sampling count is sufficient to generate converged results for the 3D scenes presented in the article.

The implementation uses the C++ programming language and the Vulkan graphics library. We generated the results on a machine with an i7-4770K @ 3.50GHz CPU and an NVIDIA GTX 1080 GPU. The C++/Vulkan source code is available at <https://github.com/iota97/AnisotropyEditor>.




Figure 15: Rendering of a sphere with $N = 16$, $\eta = 0.95$, and $\alpha_{in} = 0.5$. We compare our technique (left) and bent normal (center) with the reference image (right). Frame times are based on a rendering resolution of 1920×1080 with $8 \times$ MSAA.

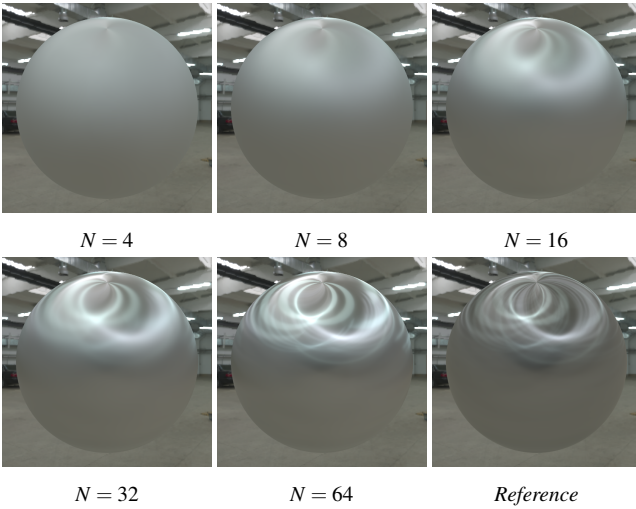


Figure 16: Rendering of a sphere with $\eta = 0.95$ and $\alpha_{in} = 0.7$. We compare our method using different numbers of samples with the reference image.

We use HDR environment cube maps with a resolution of 512×512 . We prefiltered them using the prefiltered importance sampling algorithm by Lagarde and de Rousiers [LdR14], which is based on the work of Křivánek and Colbert 2008 [KC08].

We provide a supplemental video to demonstrate that our method is free of artifacts and noise, and to present a temporal version of Figures 15, 20, and 22.

The reference implementation of our method is available in a Shadertoy code: <https://www.shadertoy.com/view/4fv3W3>. In this Shadertoy, we demonstrate that our technique can be used as a drop-in replacement for the bent normal technique.

Quality analysis for different sample counts. As discussed in Sections 5.7 and 5.8, we set $\beta = 2.3$ and $k = 1$. The only remaining parameter to study is the number of samples, N . Figure 16 visually compares the results for different values of N . Figure 17 shows the rendering time for various N . Figure 18 provides the quantitative

deviations from the reference for our method and the bent normal technique using the FLIP metric [ANA*20].




Figure 17: The average frame time of the video's Garage scene using different numbers of samples is compared with the bent normal baseline. The rendering resolution is 1920×1080 with $8 \times$ MSAA.




Figure 18: The average FLIP mean pixel-wise deviation of the video's Garage scene using different numbers of samples compared with the bent normal baseline.

By considering both qualitative and quantitative comparisons between different sample counts, we find $N = 16$ to be a good compromise between quality and performance, as it is about 50% slower than the bent normal technique. Values of $N < 8$ do not provide sufficient quality, and values of $N > 32$ offer only marginal gains in visual quality while increasing rendering time.

Comparison with the bent normal technique. In Figures 1, 15, 19, 20 and 22, we compare our method with $N = 16$, the bent normal technique, and the ground truth. The results show that the bent normal technique is unable to reproduce the reference image accurately. Our method, while more blurry than the reference image, provides a conservative integration that still closely resembles it.

Figure 19 shows the FLIP pixel-wise deviation as a function of roughness α_{in} and anisotropy η . Our method provides better results than the bent normal technique for $\alpha_{in} < 0.5$. However, our method deviates more from the reference for higher roughness ($\alpha_{in} > 0.5$) because our approximations become less accurate (see Section 5.8).

Comparison with filtered importance sampling. Filtered importance sampling was developed before the introduction of the split-integral approximation for IBL and was utilized by Dupuy

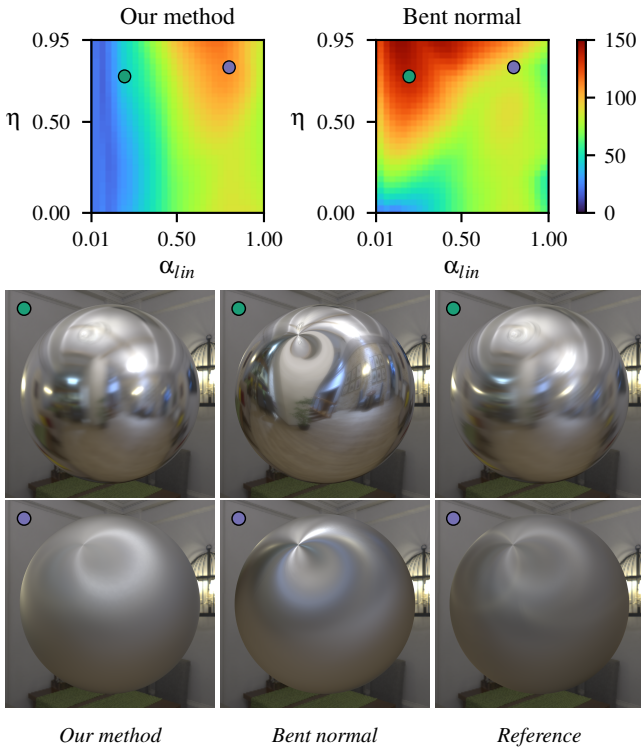


Figure 19: Mean of the ALIP pixel-wise deviations ($\times 10^3$) for different values of roughness α_{lin} and anisotropy η . Rendering images are shown for parameter pairs $(\alpha_{lin} = 0.2, \eta = 0.75)$ and $(\alpha_{lin} = 0.8, \eta = 0.8)$.




Figure 20: Rendering of a sphere with $N = 16$, $\eta = 0.8$, and $\alpha_{lin} = 0.4$. We compare our technique (left) and the bent normal technique (center) with a reference image (right). In some situations, the bent normal technique can flip the environment upside down. Frame rendering times refer to a resolution of 1920×1080 with $8 \times$ MSAA.

et al. [DHI*13] for real-time rendering of specular materials. Figure 21 demonstrates that their method produces overly blurry specular reflections, even with 25 environment samples and 2.5 times the rendering time of our method.

7. Limitations and future work

We presented a new anisotropic specular anisotropic IBL technique based on the split-integral approximation. Our method can replace




Figure 21: Rendering of a sphere with $\eta = -0.8$ and $\alpha_{lin} = 0.4$. We compare our technique using $N = 16$ (left) with Dupuy et al. 2013 [DHI*13], which uses 5×5 samples (center), and with a reference image (right). Frame rendering times are for a resolution of 1920×1080 with $8 \times$ MSAA.

the bent normal technique, providing more realistic results within approximately $1.5 \times$ its rendering time budget.

Our method has limitations. We assume that the NDF mainly determines the shape of the BRDF. While this is true in most cases, the masking-shadowing function G_2 significantly impacts the BRDF shape for grazing views and high roughnesses. In these cases, our approximation of the BRDF confidence region deviates from its actual shape. To avoid large deviations, we limited the size of the confidence region.

Currently, all samples have the same weight; an interesting future direction is to investigate a weighting scheme that considers the BRDF density to further improve the rendering quality. Additionally, performance could be improved by dynamically adjusting the sample count based on the anisotropy level. However, this requires special care to avoid seam artifacts where the number of samples changes.

Normal mapping support could be added by expressing the view direction in a rotated space whose normal is defined by the normal map. Another interesting direction for future work is extending the method for non-centered NDFs [DHI*13]. Closed-form formulas might exist when the slope mean is not zero. In that case, our technique could be used with this linear slope filtering method. Additionally, support for multiple-scattering microfacet models could be investigated, as done by Fdez-Agüera [FA19] for isotropic materials.

Acknowledgments

This work was supported by the French National Research Agency (ANR) under grant number ANR-24-CE10-6403.

References

- [AMHH*18] AKENINE-MÖLLER T., HAINES E., HOFFMAN N., PESCE A., IWANICKI M., HILLAIRE S.: *Real-Time Rendering 4th Edition*. A K Peters/CRC Press, Boca Raton, FL, USA, 2018. 2, 4
- [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. doi:10.1145/3406183. 9




Figure 22: Rendering of a plate with $N = 16$, $\eta = 0.95$, and $\alpha_{in} = 0.3$. The anisotropic direction \bar{U} is defined using "The Great Wave off Kanagawa", where its colors are mapped to an angle in $[0, \pi)$. We compare our method (left) and the bent normal technique (center) with the reference image (right). Frame rendering times are for a resolution of 1920×1080 with $8 \times$ MSAA.

- [Bur12] BURLEY B.: Physically-Based Shading at Disney. In *Practical Physically-Based Shading in Film and Game Production* (Los Angeles, California, 2012), pp. 10:1–10:7. 4
- [CK11] CHUANG M., KAZHDAN M.: Interactive and anisotropic geometry processing using the screened Poisson equation. *ACM Trans. Graph. (Proc. SIGGRAPH)* 30, 4 (2011). doi:10.1145/2010324.1964952. 2
- [DHI*13] DUPUY J., HEITZ E., IEHL J.-C., POULIN P., NEYRET F., OSTROMOUKHOV V.: Linear Efficient Antialiased Displacement and Reflectance Mapping. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 32, 6 (2013). 3, 8, 10
- [FA19] FDEZ-AGÜERA C. J.: A Multiple-Scattering Microfacet Model for Real-Time Image Based Lighting. *Journal of Computer Graphics Techniques* 8, 1 (2019), 45–55. 3, 10
- [GKD07] GREEN P., KAUTZ J., DURAND F.: Efficient Reflectance and Visibility Approximations for Environment Map Rendering. *Comput. Graph. Forum (Proc. of Eurographics)* 26, 3 (2007), 495–502. doi:10.1111/j.1467-8659.2007.01072.x. 2
- [Goo23] GOOGLE: Filament, 2023. <https://github.com/google/filament>. 2
- [Got12] GOTANDA Y.: Beyond a Simple Physically Based Blinn-Phong Model in Real-Time. In *Practical physically-based shading in film and game production* (2012), ACM SIGGRAPH Courses. doi:10.1145/2343483.2343493. 2, 3
- [Gre86] GREENE N.: Environment Mapping and Other Applications of World Projections. *IEEE Computer Graphics & Applications* 6, 11 (1986), 21–29. doi:10/b7m3gc. 2
- [Hec89] HECKBERT P. S.: Fundamentals of texture mapping and image warping, 1989. 2
- [Hei14] HEITZ E.: Understanding the Masking-Shadowing Function in Microfacet-Based BRDFs. *Journal of Computer Graphics Techniques* 3, 2 (2014), 48–107. 2, 3, 4
- [Kar13] KARIS B.: Real Shading in Unreal Engine 4. In *Physically based shading in theory and practice* (2013), ACM SIGGRAPH Courses. doi:10.1145/2504435.2504457. 2, 3, 4
- [KC08] KRIVÁNEK J., COLBERT M.: Real-time Shading with Filtered Importance Sampling. *Comput. Graph. Forum (Proc. Eurographics Symposium on Rendering)* 27, 4 (2008), 1147–1154. doi:10.1111/j.1467-8659.2008.01252.x. 3, 9
- [KC17] KULLA C., CONTY A.: Revisiting Physically Based Shading at Imageworks. In *Physically based shading in theory and practice* (2017), ACM SIGGRAPH Courses. doi:10.1145/3084873.3084893. 2, 3, 4
- [KGK19] KNEIPHOF T., GOLLA T., KLEIN R.: Real-time Image-based Lighting of Microfacet BRDFs with Varying Iridescence. *Comput. Graph. Forum (Proc. Eurographics Symposium on Rendering)* 38, 4 (2019), 77–85. doi:10.1111/cgf.13772. 3
- [KM00] KAUTZ J., MCCOOL M.: Approximation of Glossy Reflection with Prefiltered Environment Maps. In *Proc. of the Conference on Graphics Interface* (2000), pp. 119–126. doi:10.20380/GI2000.17. 3
- [Laz13] LAZAROV D.: Getting More Physical in Call of Duty: Black Ops II. In *Physically based shading in theory and practice* (2013), ACM SIGGRAPH Courses. doi:10.1145/2504435.2504457. 2, 3
- [LdR14] LAGARDE S., DE ROUSIERS C.: Moving Frostbite to Physically Based Rendering. In *Physically based shading in theory and practice* (2014), ACM SIGGRAPH Courses. doi:10.1145/2614028.2615431. 3, 4, 5, 6, 9
- [LG18] LAGARDE S., GOLUBEV E.: The Road Toward Unified Rendering with Unity's High Definition Render Pipeline. In *Advances in Real-Time Rendering in Games* (2018), ACM SIGGRAPH Courses. doi:10/gf3tbf. 1, 2, 3, 5, 8
- [McA15] MCAULEY S.: Rendering the World of Far Cry 4. In *Game Developers Conference* (2015). URL: <https://www.gdcvault.com/play/1022235/Rendering-the-World-of-Far-Cry-4>. 2, 3, 5
- [MLH02] MCALLISTER D. K., LASTRA A., HEIDRICH W.: Efficient rendering of spatial bi-directional reflectance distribution functions. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hardware* (2002), pp. 79–88. doi:10.2312/EGGH/EGGH02/079-088. 2
- [MS16] MANSON J., SLOAN P.-P.: Fast filtering of reflection probes. *Comput. Graph. Forum (Proc. Eurographics Symposium on Rendering)* 35, 4 (2016), 119–127. doi:https://doi.org/10.1111/cgf.12955. 4
- [Nar14] NARKOWICZ K.: Analytical DFG term for IBL, 2014. URL: <https://knarkowicz.wordpress.com/2014/12/27/analytical-dfg-term-for-ibl/>. 3
- [Pat20] PATRY J.: Samurai Shading in Ghost of Tsushima. In *Physically based shading in theory and practice* (2020), ACM SIGGRAPH Courses. doi:10.1145/3388769.3407523. 4
- [PI15] PESCE A., IWANICKI M.: Approximate Models For Physically Based Rendering. In *Physically based shading in theory and practice* (2015), ACM SIGGRAPH Courses. doi:10.1145/2776880.2787670. 3, 4, 5

- [RBSM19] RIBARDIÈRE M., BRINGIER B., SIMONOT L., MENEVEAUX D.: Microfacet BSDFs Generated from NDFs and Explicit Microgeometry. *ACM Trans. Graph.* 38, 5 (June 2019). doi:10.1145/3338697.5
- [Rev11] REVIE D.: Implementing fur using deferred shading. In *GPU Pro 2*. A. K. Peters, Ltd., 2011, pp. 57–75. 1, 2, 3, 5
- [Sch94] SCHLICK C.: An Inexpensive BRDF Model for Physically-based Rendering. *Comput. Graph. Forum (Proc. of Eurographics)* 13, 3 (1994), 233–246. doi:10.1111/1467-8659.1330233. 3, 4
- [WB20] WEIER P., BELCOUR L.: Rendering Layered Materials with Anisotropic Interfaces. *Journal of Computer Graphics Techniques* 9, 2 (2020), 37–57. URL: <http://jcgt.org/published/0009/02/03/.3>
- [YCC*22] YU H., CHEN A., CHEN X., XU L., SHAO Z., YU J.: Anisotropic Fourier Features for Neural Image-Based Rendering and Relighting. *Proceedings of the AAAI Conference on Artificial Intelligence* 36, 3 (2022), 3152–3160. doi:10.1609/aaai.v36i3.20223.2
- [YYTM19] YAMAGUCHI T., YATAGAWA T., TOKUYOSHI Y., MORISHIMA S.: Real-time Rendering of Layered Materials with Anisotropic Normal Distributions. In *SIGGRAPH Asia Technical Briefs* (2019), pp. 87–90. doi:10.1145/3355088.3365165.3

Appendix A: Closed-form solution for the extrema

The derivation of formulas and properties are presented in the Supplementary Material. When solving Equation 7, the first thing to note is that the alignment measure $d(t)$ is π -periodic, solution can therefore be expressed based on t_0 and t_1 verifying:

$$\operatorname{argmax}_t d(t) = t_0 \bmod \pi \quad \text{and} \quad \operatorname{argmin}_t d(t) = t_1 \bmod \pi.$$

When directions are more aligned, their associated points on the unit sphere are closer in spherical distance. Therefore the two maximum alignments are achieved when $t = t_0$ and $t = t_0 + \pi$, giving the two closest extrema $\check{\omega}_{\perp 0} = \check{\omega}_{t_0}$ and $\check{\omega}_{\perp 1} = \check{\omega}_{t_0 + \pi}$. Similarly, the two minimum alignments are achieved when $t = t_1$ and $t = t_1 + \pi$, giving the two farthest extrema $\check{\omega}_{\top 0} = \check{\omega}_{t_1}$ and $\check{\omega}_{\top 1} = \check{\omega}_{t_1 + \pi}$.

To simplify the formulas, we rotate the view direction:

$$\begin{pmatrix} \dot{o}_x \\ \dot{o}_y \end{pmatrix} = \begin{bmatrix} \overline{U}^T \\ \overline{V}^T \end{bmatrix} \begin{pmatrix} o_x \\ o_y \end{pmatrix}.$$

The solutions are then provided by:

$$t_0 = \begin{cases} \tan^{-1} \left(\frac{\varphi_1 + \sqrt{\varphi_1^2 - \varphi_3}}{\varphi_2} \right) & \text{if } \varphi_2 \neq 0, \\ \frac{\pi}{2} & \text{if } \varphi_2 = 0 \text{ and } \varphi_1 \geq 0, \\ 0 & \text{if } \varphi_2 = 0 \text{ and } \varphi_1 < 0 \end{cases}$$

$$t_1 = \begin{cases} \tan^{-1} \left(\frac{\varphi_1 - \sqrt{\varphi_1^2 - \varphi_3}}{\varphi_2} \right) & \text{if } \varphi_2 \neq 0, \\ t_0 + \frac{\pi}{2} & \text{if } \varphi_2 = 0, \end{cases}$$

with

$$r_x = k\sigma_{\overline{U}}, \quad r_y = k\sigma_{\overline{V}}, \quad \varphi_0 = 2\dot{o}_x\dot{o}_y r_x r_y,$$

$$\varphi_1 = \dot{o}_x^2 (r_x^2 r_y^2 + r_y^2) - \dot{o}_y^2 (r_x^2 r_y^2 + r_x^2) + r_x^2 - r_y^2,$$

$$\varphi_2 = -\varphi_0 (r_y^2 + 1), \quad \text{and} \quad \varphi_3 = \varphi_0 \varphi_2 (r_x^2 + 1).$$

Appendix B: Closed-form solution for τ_0 and τ_1

A solution can be derived knowing that the extrema are symmetric with respect to the origin in slope space. Consequently, the following is true in normal space:

$$x_{\top 0} = -x_{\top 1}, \quad y_{\top 0} = -y_{\top 1}, \quad z_{\top 0} = z_{\top 1}.$$

Solving the equation $\check{z}_\tau = 0$ with the previous equalities leads to the following formulas:

$$\tau_0 = \frac{1}{2} + \frac{z_{\top 1} \left(p_0 - \sqrt{p_0^2 + o_z p_1} \right)}{2p_1}$$

$$\tau_1 = \frac{1}{2} + \frac{z_{\top 1} \left(p_0 + \sqrt{p_0^2 + o_z p_1} \right)}{2p_1}$$

$$p_0 = x_{\top 1} x_o + y_{\top 1} y_o$$

$$p_1 = z_o (1 - z_{\top 1}^2)$$

Note that $\omega_{\top 0}$ and $\omega_{\top 1}$ are not equivalent in this formula as τ_0 and τ_1 must be used as an interpolation parameter between the two.

Appendix C: Shifted minimum extrema

To take into account the collapse of the closest extrema $\omega_{\perp 0}$ and $\omega_{\perp 1}$, we move these points along the boundary of the confidence region by relying on the initial ellipse parameterization using the following formula:

$$\delta = \begin{cases} \tau'_0 (t_0 - t_1), & \tau'_0 > 0 \\ (\tau'_1 - 1)(t_0 - t_1), & \tau'_0 \leq 0 \end{cases}$$

$$\omega'_{\perp 0} = \omega_{t_0 + \delta}$$

$$\omega'_{\perp 1} = \omega_{t_0 - \delta + \pi},$$

which can be seen as an interpolation of the angles t_0 and t_1 .

Appendix D: Farthest samples interpolant

The interpolant ζ is defined as the ratio

$$\zeta = \frac{\|\check{\omega}'_{\top 0} - \check{\omega}'_{\top 1}\| - \|\check{\omega}'_{\perp 0} - \check{\omega}'_{\perp 1}\|}{\|\check{\omega}'_{\top 0} - \check{\omega}'_{\top 1}\|}.$$

This ratio is computed in the reflected space to account for the BRDF's shape anisotropy.