DOI: 10.1111/cgf.13772
Eurographics Symposium on Rendering 2019
T. Boubekeur and P. Sen
(Guest Editors)

Volume 38 (2019), Number 4

Real-time Image-based Lighting of Microfacet BRDFs with Varying Iridescence

Tom Kneiphof[1] and Tim Golla[1] and Reinhard Klein[1]

[1]University of Bonn, Germany

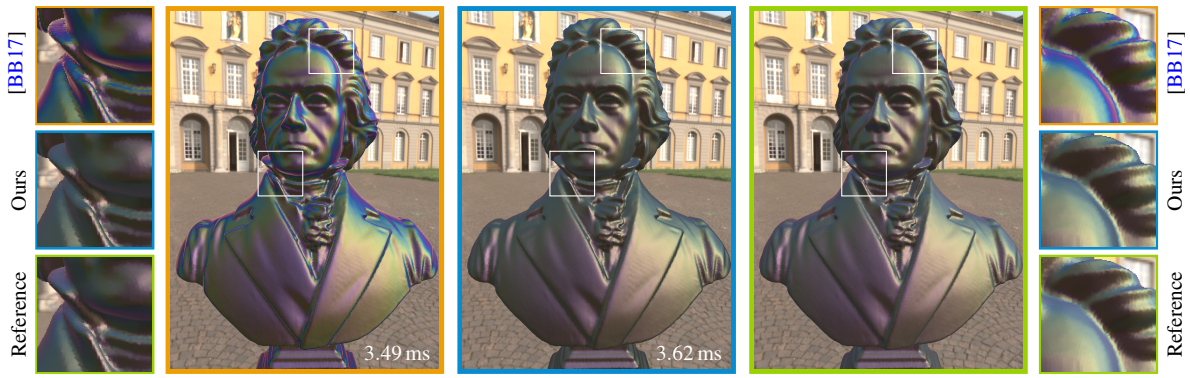

Figure 1: *We present a method for rendering iridescent effects in real-time. In comparison to previous methods, our method shows more realistic results on rough surfaces, which resemble the offline rendered reference more closely. Our method improves upon previous work by considering a distribution of light directions instead of only regarding the direction of perfect reflection. The Beethoven bust model shown here has a surface roughness of* $\alpha = 0.25$. *The thin-film is* $d = 550\,\mathrm{nm}$ *thick and has a refractive index of* $\eta_2 = 1.2$. *The base material has a complex refractive index of* $\eta_3 = 1.27$ *and* $\kappa_3 = 0.44$.

Abstract

Iridescence is a natural phenomenon that is perceived as gradual color changes, depending on the view and illumination direction. Prominent examples are the colors seen in oil films and soap bubbles. Unfortunately, iridescent effects are particularly difficult to recreate in real-time computer graphics. We present a high-quality real-time method for rendering iridescent effects under image-based lighting. Previous methods model dielectric thin-films of varying thickness on top of an arbitrary micro-facet model with a conducting or dielectric base material, and evaluate the resulting reflectance term, responsible for the iridescent effects, only for a single direction when using real-time image-based lighting. This leads to bright halos at grazing angles and over-saturated colors on rough surfaces, which causes an unnatural appearance that is not observed in ground truth data. We address this problem by taking the distribution of light directions, given by the environment map and surface roughness, into account when evaluating the reflectance term. In particular, our approach prefilters the first and second moments of the light direction, which are used to evaluate a filtered version of the reflectance term. We show that the visual quality of our approach is superior to the ones previously achieved, while having only a small negative impact on performance.

CCS Concepts

• ***Computing methodologies*** → ***Computer graphics; Rendering; Reflectance modeling;***

1. Introduction

Iridescence is a phenomenon, where a surface changes its color depending on the direction from which it is viewed or illuminated. Iridescent effects are well known from oil leaks and soap bubbles to the vivid colors of some animals such as insects, birds and fishes. Oftentimes, these effects are very subtle and are not immediately noticeable to the untrained eye, but nevertheless contribute to the photorealism of a scene. Iridescence is a consequence of the wave

nature of light. It is caused by the interference of light waves, which have been scattered in a wavelength-dependent way by the surface. There are two primary mechanisms causing iridescence:

- Diffraction, which results from light reflection on a microscopic structure with features on the scale of the visible wavelengths.
- Interference of light waves interacting with thin-films with a thickness on the scale of visible wave lengths.

Belcour and Barla [BB17] introduced an extension to microfacet theory for the rendering of iridescent effects caused by thin-films of varying thickness on top of an arbitrarily rough base layer. If the features of the microfacet surface are assumed to be much larger than the wavelengths of visible light, diffraction effects can be neglected. Until recently, it was neccessary to consider individual wavelengths separately in order to compute iridescent effects accurately. In their work, Belcour and Barla [BB17] introduced a method that can be evaluated in real-time, using only an RGB light representation. However, their method exhibits the following limitation (see Figure 16 in [BB17]): With increasing surface roughness, the iridescent colors become over-saturated in their real-time image-based lighting implementation in comparison to a path-traced reference image. The reason is that the iridescent colors are only computed for a single light direction in their real-time image-based lighting implementation, which is a good approximation on smooth surfaces, but not on rough surfaces. On rough surfaces, light from different directions contributes different iridescent colors, which are mixed together, resulting in a less saturated appearance.

In this paper, we propose an improved method for the real-time rendering of iridescent effects on rough surfaces under image-based lighting, based on the approach of Belcour and Barla [BB17]. Our new approach takes the actual lighting conditions into account when computing the iridescent colors, as defined by the environment map. We show that our approach visually outperforms the previous approach. The key observation is that the light is not only coming from the ideal reflection direction, but many directions, which are contributing different colors. The ideal reflection direction is replaced with a distribution of directions for which the reflectance term is evaluated. Since the iridescent colors are so sensitive to changes in the light direction, we infer that it is necessary to take information on the distribution of light directions into account. Therefore, a simpler filtering of the iridescent colors without the distribution on light directions would not be able to reproduce the ground truth as closely as our method.

The remainder of this paper is structured as follows: Section 2 gives an overview of the related work. Section 3 explains fundamentals from previous work, required to understand our method. In Section 4, we describe our method and Section 5 contains the evaluation of our method, followed by a conclusion in Section 6.

2. Related Work

We give a short overview of the works on rendering diffractive models, rendering thin-film models and research in the field of real-time image-based lighting that is most relevant for our work.

Dhillon et al. [DTS*14] measured the height fields of snake skin and rendered the arising iridescent effects using a Taylor series expansion of the Fourier transform of a function of the measured height field. Toisoul and Ghosh [TG17a] acquired and rendered homogeneous diffraction patterns and use out-of-focus bokeh photography to acquire the height field of the diffraction grating. They also provide a real-time image-based lighting solution by pre-convolving the environment map with the diffraction lookup table. In a follow-up publication [TG17b] they used a low-rank factorization of the diffraction lookup table to convolve the environment map in real-time. Holzschuch and Pacanowski [HP17] introduced an extension to the classical microfacet BRDF [CT82], where each microfacet is modeled by an even smaller nanogeometry, responsible for diffractive effects. Werner et al. [WVJH17] focused on diffraction caused by one-dimensional scratches, which are modeled as one-dimensional curves. Their approach delivers high-quality results, but is not suitable for real-time applications. Velinov et al. [VWH18] provided a closed-form solution to this approach, which requires only a single sample per pixel and allows real-time performance. Yan et al. [YHW*18] described an approach for rendering specular geometry with high-frequency spatial surface variations modeled as height field. Their approach delivers high quality results and is suitable for offline rendering.

One of the earliest approaches to rendering thin-film interference was presented by Smits and Meyer [SM92]. They modeled the phase shift of the light traveling through the thin-film, depending on wavelength and film thickness. However, they only considered the first light path through the thin-film back in to the external medium. Icart and Arquès [IA99] modeled thin-films with complex refractive index on top of rough surfaces. Later, they proposed a model for multi-layer materials with uncorrelated interfaces [IA00]. Hirayama et al. [HKY*01] also considered multi-layer films of metallic and semi-conducting materials, restricted to smooth interfaces between the layers. An early approximation that does not rely on a spectral representation of the light but works with an RGB representation to produce iridescent effects due to thin-films, was presented by Granier and Heidrich [GH03]. Sun [Sun06] simulated biological iridescence caused by multi-layer structures in an RGB-based renderer, by converting the RGB colors to a spectral representation and back. A generic approach to render structural colors in real-time was presented by Imura et al. [IOS*09]. They precomputed a structural light texture, parameterized by the optical path difference. Belcour and Barla [BB17] used an Airy summation to take all light paths inside the thin-film into account. Then, they considered the Fourier transformed reflectance in order to compute the response to an arbitrary spectral band. A more detailed explanation of this method, on which ours is based, is given in Section 3.1.

Kautz et al. [KVHS00] accelerated the prefiltering of environment maps with the goal of rendering specular reflections under arbitrary illumination in real-time. Karis [Kar13] used a split-sum approximation to preintegrate the BRDF and environment map separately. Lagarde and De Rousiers [LDR14] gave a SIGGRAPH course on moving the Frostbite engine to physically based rendering. Among other insights, they noticed a shift of the dominant specular direction at increased viewing angles in the split-sum approximation, and discerned a problem with the anisotropic nature of BRDF-lobes, resulting from a microfacet BRDF [CT82].

3. Preliminaries

This section outlines the concepts and results introduced in prior work, that lay the foundation of our method.

3.1. Iridescent BRDF Model

Because our approach is built on the work of Belcour and Barla [BB17], we briefly summarize their method in this section. They model the surface reflectance based on a classical microfacet BRDF [CT82]. Traditionally, the microfacets are assumed to be perfectly smooth mirrors. Here, the microfacets are no longer just perfectly smooth, but are individually coated by a thin-film, which is incorporated in the BRDF model by replacing the Fresnel reflectance term with a more complex reflectance term R, which accounts for all inter-reflections inside the thin-film and wave-interference effects. The resulting BRDF is written as

$$\rho(\omega_i, \omega_o) = \frac{D(\mathbf{h})G(\omega_i, \omega_o)R(\langle \mathbf{h}, \omega_i \rangle)}{4\langle \mathbf{n}, \omega_i \rangle \langle \mathbf{n}, \omega_o \rangle}, \tag{1}$$

where ω_i and ω_o are the light and viewing direction, respectively. $\mathbf{h} = {\omega_i+\omega_o}/{||\omega_i+\omega_o||}$ is the halfway vector and $\mathbf{n}$ is the surface normal. D defines a distribution on the microfacet orientations, G models geometric shadowing and masking effects and R is the reflectance term, describing the reflectance behaviour of an individual microfacet.

The thin-film thickness d is assumed to be constant across a single microfacet, resulting in parallel interfaces of the film. However, it is not required to be constant on a macroscopic scale, allowing for spatial thickness variations across a surface. The thin-film is modeled as a dielectric material with refractive index η_2, while the underlying base of the microfacet is a dielectric or conducting material with complex refractive index $\eta_3 + i\kappa_3$. The refractive index of the surrounding medium is $\eta_1 = 1$ for air. The angle of incidence on a microfacet in the surrounding medium is called θ_1 with $\cos\theta_1 = \langle \mathbf{h}, \omega_i \rangle$. The refracted angle inside the thin-film is θ_2.

The reflectance term is defined for individual light frequencies ν (or wavelengths λ, respectively). However, we are not directly interested in the reflectance for individual light frequencies or wavelengths, but want to know the response to a predefined set of spectral bands. In our real-time rendering setting, we are interested in the three RGB bands. For each spectral band $j \in \{R, G, B\}$, a sensitivity function $S_j(\nu)$ is defined. The corresponding reflectance R_j is computed by taking the integral over all light frequencies ν and weighting them with the sensitivity $S_j(\nu)$:

$$R_j(\langle \mathbf{h}, \omega_i \rangle) = \int R(\langle \mathbf{h}, \omega_i \rangle; \nu) \cdot S_j(\nu) \, \mathrm{d}\nu \tag{2}$$

$$= \int \hat{R}(\langle \mathbf{h}, \omega_i \rangle; \mu) \cdot \overline{\hat{S}_j(\mu)} \, \mathrm{d}\mu, \tag{3}$$

where $\hat{R}$ and $\hat{S}$ are the Fourier transforms with respect to ν of R and S, respectively. Each of these terms can now be written as a series

$$R_j(\langle \mathbf{h}, \omega_i \rangle) = C_0 + \sum_{m=1}^{\infty} C_m \left[e^{im\phi_2} \hat{S}_j(-m\mathcal{D}) + e^{-im\phi_2} \hat{S}_j(m\mathcal{D}) \right], \tag{4}$$

where $C_0 = R_{12} + R_\star$, $C_m = \left(\sqrt{R_{23}R_{21}}\right)^m \left(R_\star - \sqrt{T_{12}T_{21}}\right)$, $R_\star = \frac{T_{12}T_{21}R_{23}}{1-R_{23}R_{21}}$ with R_{ab} and T_{ab} being the Fresnel reflectance and transmittance at the interface going from layer a to b, respectively. $\phi_2 = \phi_{21} + \phi_{23}$, with ϕ_{ab} being the phase shift induced by a reflection at the interface going from layer a to b. The optical path difference is defined as

$$\mathcal{D} := 2\eta_2 d \cos\theta_2. \tag{5}$$

From the Fresnel equations we know that the phase shifts and reflectances at the interface boundaries also depend on the light polarization. Assuming unpolarized incoming light, we compute the final reflectance for parallel and orthogonal polarization, and write $R(\langle \mathbf{h}, \omega_i \rangle) = \frac{1}{2}(R^{\perp}(\langle \mathbf{h}, \omega_i \rangle) + R^{\parallel}(\langle \mathbf{h}, \omega_i \rangle))$. In the remainder of this work, we do not write the polarization superscript and the subscript j for the spectral band. The reflectance derived in Section 4 is computed for each spectral band and light polarization.

3.2. Real-time Image-based Lighting

In order to produce plausible rendering of an object in a certain context, it does not suffice to illuminate the object only by a limited number of analytical light sources. In the most general situation, the incident light is described by an environment map, which defines the radiance reaching an object or scene from each direction. Illuminating an object under such lighting, requires to solve the integral over all incoming directions ω_i in the rendering equation [Kaj86], which we will call the rendering integral:

$$L_o(\omega_o) = \int_{\Omega_i} \rho(\omega_i, \omega_o) \langle \mathbf{n}, \omega_i \rangle L_i(\omega_i) \, \mathrm{d}\omega_i, \tag{6}$$

where $L_o(\omega_o)$ is the radiance leaving a surface in direction ω_o, and $L_i(\omega_i)$ is the incident radiance to the surface from direction ω_i, as defined by the environment map. An arbitrarily good solution can be found by performing Monte-Carlo integration, and therefore drawing many samples from the environment map to approximate the integral. However, this is not feasible in real-time applications, where we wish to achieve a good approximation with as few samples as possible.

Integrating over the product of multiple functions in Eq. 6 is rather problematic. To simplify the situation, we use the following factorization: Given two functions $f, g : U \to \mathbb{R}$, we can write the integral over the product of f and g as

$$\int f(x) \cdot g(x) \, \mathrm{d}x = \int f(x) \, \mathrm{d}x \cdot \int \frac{f(x)}{\int f(x') \, \mathrm{d}x'} \cdot g(x) \, \mathrm{d}x, \tag{7}$$

where $\frac{f(x)}{\int f(x') \, \mathrm{d}x'}$ defines a probability distribution on x.

By applying this to Eq. 6, we get the product of two integrals

$$L_o(\omega_o) = \int_{\Omega_i} \rho(\omega_i, \omega_o) \langle \mathbf{n}, \omega_i \rangle L_i(\omega_i) \, \mathrm{d}\omega_i \tag{8}$$

$$= \int_{\Omega_i} \rho(\omega_i, \omega_o) \langle \mathbf{n}, \omega_i \rangle \, \mathrm{d}\omega_i \cdot \int_{\Omega_i} p(\omega_i) L_i(\omega_i) \, \mathrm{d}\omega_i, \tag{9}$$

where $p(\omega_i) \propto \rho(\omega_i, \omega_o) \langle \mathbf{n}, \omega_i \rangle$ is a probability distribution on the light directions ω_i. The first integral is independent of the incoming light $L_i(\omega_i)$ and is known as the prefiltered BRDF. When Schlick's approximation [Sch94] is used for the Fresnel reflectance term of the BRDF ρ, it can be precomputed for all combinations of

$\langle \mathbf{n}, \omega_o \rangle$ and roughness α, which is then stored in a discretized two-dimensional table [Kar13, LDR14]. However, $p(\omega_i)$ still depends on both, ω_o and ω_i, as well as the surface normal $\mathbf{n}$. Let ω_r be the ideal reflection direction of the view direction ω_o at the surface normal $\mathbf{n}$. Approximating $p(\omega_i)$, by choosing $\omega_o = \mathbf{n} = \omega_r$, is equivalent to the "split sum approximation" introduced by Karis [Kar13]. With this approximation, the second integral only depends on the light direction ω_i and is known as the prefiltered environment map. The approximated $p(\omega_i)$ defines the filter kernel for the prefiltered environment map, parameterized by the reflection direction ω_r and surface roughness α [Kar13, LDR14].

4. Real-time Image-based Lighting with Iridescence

In this section, we present an improved method for evaluating image-based lighting for rough and iridescent microfacet BRDFs in real-time. Belcour and Barla [BB17] proposed to evaluate the reflectance term for the ideal-reflection direction only, and evaluate the image-based lighting without the reflectance term. This method yields good results for smooth surfaces. On rough surfaces however, where light is contributing from various directions, it is no longer viable to assume that all light is coming from a single direction for the evaluation of the reflectance term. Light from different directions causes different iridescent colors, and mixing them together results in a less saturated appearance. When evaluating the reflectance term for a single direction only, this desaturation effect does not occur.

4.1. Factorization of the Rendering Integral

We propose an extension to the method of Belcour and Barla [BB17], which takes the desaturation effect of the iridescent colors on rough surfaces into account. We also separate the evaluation of the reflectance term from the rest of the BRDF. Therefore, we define the BRDF without reflectance term as

$$\tilde{\rho}(\omega_i, \omega_o) := \frac{D(\mathbf{h})G(\omega_i, \omega_o)}{4\langle \mathbf{n}, \omega_i \rangle \langle \mathbf{n}, \omega_o \rangle} = \frac{\rho(\omega_i, \omega_o)}{R(\langle \mathbf{h}, \omega_i \rangle)}. \tag{10}$$

Having this definition at hand, we apply the factorization in Eq. 7 twice to Eq. 6, which yields

$$L_o(\omega_o) = \int_{\Omega_i} \tilde{\rho}(\omega_i, \omega_o) R(\langle \mathbf{h}, \omega_i \rangle) \langle \mathbf{n}, \omega_i \rangle L_i(\omega_i) \, \mathrm{d}\omega_i \tag{11}$$

$$= \int_{\Omega_i} \tilde{\rho}(\omega_i, \omega_o) \langle \mathbf{n}, \omega_i \rangle \, \mathrm{d}\omega_i \cdot \int_{\Omega_i} p_{L_i}(\omega_i) \cdot L_i(\omega_i) \, \mathrm{d}\omega_i \cdot \int_{\Omega_i} p_R(\omega_i) \cdot R(\langle \mathbf{h}, \omega_i \rangle) \, \mathrm{d}\omega_i, \tag{12}$$

where $p_{L_i}(\omega_i) \propto \tilde{\rho}(\omega_i, \omega_o) \langle \mathbf{n}, \omega_i \rangle$ and $p_R(\omega_i) \propto \tilde{\rho}(\omega_i, \omega_o) \langle \mathbf{n}, \omega_i \rangle L_i(\omega_i)$ are probability distributions on ω_i. The first integral over $\tilde{\rho}(\omega_i, \omega_o)$ is the prefiltered BRDF from Eq. 9 with a constant reflectance term. This can be solved exactly for any combination of $\langle \mathbf{n}, \omega_o \rangle$ and surface roughness α, since the reflectance term R is not involved. The second integral over $L_i(\omega_i)$ is the prefiltered environment map from Eq. 9, where $p_{L_i}(\omega_i)$ is approximated by choosing $\omega_o = \mathbf{n} = \omega_r$ as described in Section 3.2. What remains is the third integral of the reflectance term, depending on the coefficients C_m, the phase ϕ_2 and the optical path difference $\mathcal{D}$. In the following, we describe how we approximate the solution of this integral at runtime.

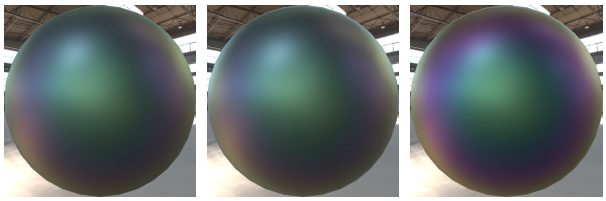

(a) *Constant C_m, ϕ_2* **(b)** *Reference* **(c)** *Constant $\mathcal{D}$*

Figure 2: *Comparison of using the same C_m, ϕ_2 and $\mathcal{D}$ for all ω_i, fixed to the mean $\langle \mathbf{h}, \omega_o \rangle$, and using the ground truth values, respectively. In (a), the same C_m and ϕ_2 have been used for all light directions in the reflectance term (Eq. 4), and in (c) the same optical path difference $\mathcal{D}$ was used instead. The surface roughness is $\alpha = 0.25$. The thin-film is $d = 600\,\mathrm{nm}$ thick, has a refractive index of $\eta_2 = 1.3$ and the base material has a refractive index of $\eta_3 = 2$ and $\kappa_3 = 0.5$.*

4.2. Integration of the Weighted Reflectance

In the reflectance term R, the coefficients C_m, the phase ϕ_2 and the optical path difference $\mathcal{D}$ all depend on the light direction ω_i. We found that the color over-saturation is primarily due to the spectral sensitivity $\hat{S}$, parameterized by $\mathcal{D}$, and not due to the C_m and ϕ_2, as illustrated in Figure 2. Therefore, we assume C_m and ϕ_2 to be identical for all participating light directions ω_i, otherwise $p_R(\omega_i) = 0$. Under this assumption, the integral over the reflectance term R in Eq. 12 is written as

$$\begin{aligned} &\int_{\Omega_i} p_R(\omega_i) \cdot R(\langle \mathbf{h}, \omega_i \rangle) \, \mathrm{d}\omega_i \\ \approx &C_0 + \sum_{m=1}^{\infty} C_m \left[e^{im\phi_2} \cdot \int_{\Omega_i} p_R(\omega_i) \cdot \hat{S}(-m\mathcal{D}) \, \mathrm{d}\omega_i \right. \\ &\left. + e^{-im\phi_2} \cdot \int_{\Omega_i} p_R(\omega_i) \cdot \hat{S}(m\mathcal{D}) \, \mathrm{d}\omega_i \right], \end{aligned} \tag{13}$$

reducing the problem to integrating the Fourier transformed spectral sensitivity $\hat{S}$.

In the following, we write ω_i in terms of $\cos\theta_i$ and ϕ_i around ω_o, where $\cos\theta_i = \langle \omega_i, \omega_o \rangle$ and ϕ_i is the azimuth angle of ω_i around ω_o. Notice that the optical path difference $\mathcal{D}$ indirectly depends on $\cos\theta_i$ in Eq. 5 (via Eq. 24), and not at all on ϕ_i. By performing a change of variables on ω_i in Eq. 13, followed by a change of variables on $\cos\theta_i$, we get

$$\int_{\Omega_i} p_R(\omega_i) \cdot \hat{S}(m\mathcal{D}) \, \mathrm{d}\omega_i \tag{14}$$

$$= \int_{-1}^{1} \int_{0}^{2\pi} p_R(\omega_i) \cdot \hat{S}(m\mathcal{D}) \, \mathrm{d}\phi_i \, \mathrm{d}\cos\theta_i \tag{15}$$

$$= \int_{\mathbb{R}} p_{\mathcal{D}}(\mathcal{D}) \cdot \hat{S}(-m\mathcal{D}) \, \mathrm{d}\mathcal{D}, \tag{16}$$

with

$$p_{\mathcal{D}}(\mathcal{D}) = \left| \frac{\partial \mathcal{D}}{\partial \cos\theta_i} \right|^{-1} \cdot \int_{0}^{2\pi} p_R(\omega_i) \, \mathrm{d}\phi_i. \tag{17}$$

We assume the optical path difference $\mathcal{D}$ to be normally dis-

tributed, i.e. $\mathcal{D} \propto \mathcal{N}\left(\mu_\mathcal{D}, \sigma^2_\mathcal{D}\right)$. This allows us to prefilter the Fourier transformed spectral sensitivity $\hat{S}$ for different means $\mu_\mathcal{D}$ and variances $\sigma^2_\mathcal{D}$ of the optical path difference:

$$\hat{\mathcal{S}}(\mu_\mathcal{D}, \sigma^2_\mathcal{D}) := \int_\mathbb{R} p_\mathcal{D}(\mathcal{D}) \hat{S}(\mathcal{D}) \, \mathrm{d}\mathcal{D}, \tag{18}$$

which is discretized into a two-dimensional lookup table. The same lookup table is used by Belcour and Barla [BB17] for anti-aliasing rapidly varying thin-film thicknesses. Using the prefiltered $\hat{\mathcal{S}}$ yields the following approximation to Eq. 13:

$$\begin{aligned} &\int_{\Omega_i} p_R(\omega_i) \cdot R(\langle \mathbf{h}, \omega_i \rangle) \, \mathrm{d}\omega_i \\ \approx &C_0 + \sum_{m=1}^{\infty} C_m \left[e^{im\phi_2} \cdot \hat{\mathcal{S}}\left(-m\mu_\mathcal{D}, m^2\sigma^2_\mathcal{D}\right) \right. \\ &\left. + e^{-im\phi_2} \cdot \hat{\mathcal{S}}\left(m\mu_\mathcal{D}, m^2\sigma^2_\mathcal{D}\right) \right]. \end{aligned} \tag{19}$$

Since the optical path difference depends on the incoming light directions ω_i, we consider the mean μ_{ω_i} and covariance Σ_{ω_i} of ω_i in $\mathbb{R}^3$, distributed according to $p_R(\omega_i)$. $p_R(\omega_i)$ is proportional to the product of $\tilde{\rho}(\omega_i, \omega_o)$ and the incoming light from the environment map $L_i(\omega_i)$, which prohibits us from computing μ_{ω_i} and Σ_{ω_i} exactly. Instead, we use an extended version of the prefiltered environment map, which introduces the $\omega_o = \mathbf{n} = \omega_r$ assumption used to prefilter the environment map (see Section 3.2) in the computation of μ_{ω_i} and Σ_{ω_i}. The environment map is extended by nine additional texture channels, which hold the first and second moments (x, y, z, x^2, y^2, z^2, xy, yz, zx) of ω_i, weighted by $L_i(\omega_i)$. Prefiltering this extended environment map yields a good approximation for the first and second moments of ω_i, distributed proportionally to $\tilde{\rho}(\omega_i, \omega_o)\langle \mathbf{n}, \omega_i \rangle L_i(\omega_i)$. μ_{ω_i} and Σ_{ω_i} are then computed from the prefiltered moments at runtime.

In order to determine the dependency of $\mu_\mathcal{D}$ and $\sigma_\mathcal{D}$ from μ_{ω_i} and Σ_{ω_i}, the distribution of light directions is first projected to the view direction ω_o, which corresponds to the integral over ϕ_i in Eq. 17. For a fixed ω_o, this is just a linear transformation, which we can easily apply to the mean and covariance of a random variable. This defines a distribution on $\langle \omega_i, \omega_o \rangle$ with

$$\mu_{\langle \omega_i, \omega_o \rangle} = \langle \omega_o, \mu_{\omega_i} \rangle, \tag{20}$$

$$\sigma^2_{\langle \omega_i, \omega_o \rangle} = {\omega_o}^t \Sigma_{\omega_i} \omega_o. \tag{21}$$

From here, we have a non-linear relation to the optical path difference. The angle of incidence onto the reflecting microfacet $\langle \mathbf{h}, \omega_i \rangle$ for each light direction is written in terms of $\langle \omega_i, \omega_o \rangle$ as

$$\langle \mathbf{h}, \omega_i \rangle = \sqrt{\frac{1 + \langle \omega_i, \omega_o \rangle}{2}}. \tag{22}$$

Since our frame of reference now changes from the macroscopic surface to individual microfacets, we write $\cos\theta_1 = \langle \mathbf{h}, \omega_i \rangle$ (see Section 3.1). For the refracted angle θ_2 inside the thin-film, Snell's law implies in conjunction with Eq. 22 that

$$\cos\theta_2 = \sqrt{1 - \frac{\eta_1^2}{\eta_2^2}\left(1 - \cos^2\theta_1\right)} \tag{23}$$

$$= \sqrt{1 + \frac{\eta_1^2}{2\eta_2^2}\left(\langle \omega_i, \omega_o \rangle - 1\right)}. \tag{24}$$

This, however, is non-linear in $\langle \omega_i, \omega_o \rangle$ and we cannot get the mean and variance of $\cos\theta_2$ directly from the distribution on $\langle \omega_i, \omega_o \rangle$. To get the mean $\mu_{\cos\theta_2}$ and variance $\sigma^2_{\cos\theta_2}$ of $\cos\theta_2$, that approximates the ground truth reasonably well, we use the first-order Taylor series expansion of $\cos\theta_2$ with respect to $\langle \omega_i, \omega_o \rangle$ around its mean to transform the distribution on $\langle \omega_i, \omega_o \rangle$ to $\cos\theta_2$. This way, the mean is transformed non-linearly, and the standard deviation is scaled by $\frac{\partial \cos\theta_2}{\partial \langle \omega_i, \omega_o \rangle}\Big|_{\mu_{\langle \omega_i, \omega_o \rangle}}$:

$$\mu_{\cos\theta_2} \approx \sqrt{1 + \frac{\eta_1^2}{2\eta_2^2}\left(\mu_{\langle \omega_i, \omega_o \rangle} - 1\right)}, \tag{25}$$

$$\sigma_{\cos\theta_2} \approx \frac{\eta_1^2}{4\eta_2^2} \frac{1}{\sqrt{1 + \frac{\eta_1^2}{2\eta_2^2}\left(\mu_{\langle \omega_i, \omega_o \rangle} - 1\right)}} \cdot \sigma_{\langle \omega_i, \omega_o \rangle}. \tag{26}$$

In a final step, the distribution on $\cos\theta_2$ is linearly transformed to the optical path difference $\mathcal{D}$, as defined in Eq. 5. This yields

$$\mu_\mathcal{D} = 2\eta_2 d \mu_{\cos\theta_2}, \tag{27}$$

$$\sigma_\mathcal{D} = 2\eta_2 d \sigma_{\cos\theta_2}. \tag{28}$$

Given μ_{ω_i} and Σ_{ω_i}, we have derived $\mu_\mathcal{D}$ and $\sigma^2_\mathcal{D}$. Eq. 13 is finally approximated as

$$\begin{aligned} &\int_{\Omega_i} p_R(\omega_i) \cdot R(\langle \mathbf{h}, \omega_i \rangle) \, \mathrm{d}\omega_i \\ \approx &C_0^* + \sum_{m=1}^{M} C_m^* \left[e^{im\phi_2^*} \hat{\mathcal{S}}\left(-m\mu_\mathcal{D}, m^2\sigma^2_\mathcal{D}\right) \right. \\ &\left. + e^{-im\phi_2^*} \hat{\mathcal{S}}\left(m\mu_\mathcal{D}, m^2\sigma^2_\mathcal{D}\right) \right], \end{aligned} \tag{29}$$

where C_j^* and ϕ_2^* are the C_j and ϕ_2 from Eq. 4, evaluated for the mean light direction μ_{ω_i}, for all $j \in \mathbb{N}_0$, and the infinite series over m is truncated after the first $M = 2$ terms. Belcour and Barla [BB17] suggested to use $M \in \{2, 4\}$ for a reasonably good approximation. Since the $\hat{S}$ is a high frequency signal, filtering it according to Eq. 18 leads to a quick attenuation of the signal for higher order terms. Therefore, we opt for $M = 2$ instead of $M = 4$ terms.

5. Results

We have implemented our method in Unity 2018.3.3f1 based on the High Definition Render Pipeline 4.0.1-preview. Our model is built on top of the pre-existing microfacet BRDF, which uses a GGX normal distribution and Smith-Joint geometry term. We compare our proposed method to reference renderings, as specified by the full BRDF model of Belcour and Barla, and their real-time approximation for image-based lighting, where the reflectance term is only

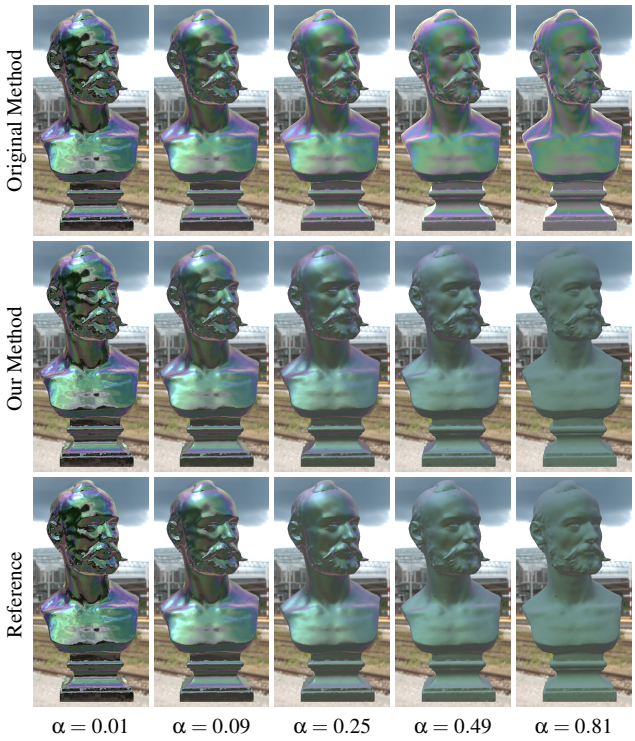

Figure 3: *The Tissot bust model is rendered for different surface roughnesses* α*. The base material has a refractive index of* $\eta_3 = 2$ *and* $\kappa_3 = 0.5$*. The thin-film has a refractive index of* $\eta_2 = 1.3$ *and thickness* $d = 650\,\text{nm}$*.*

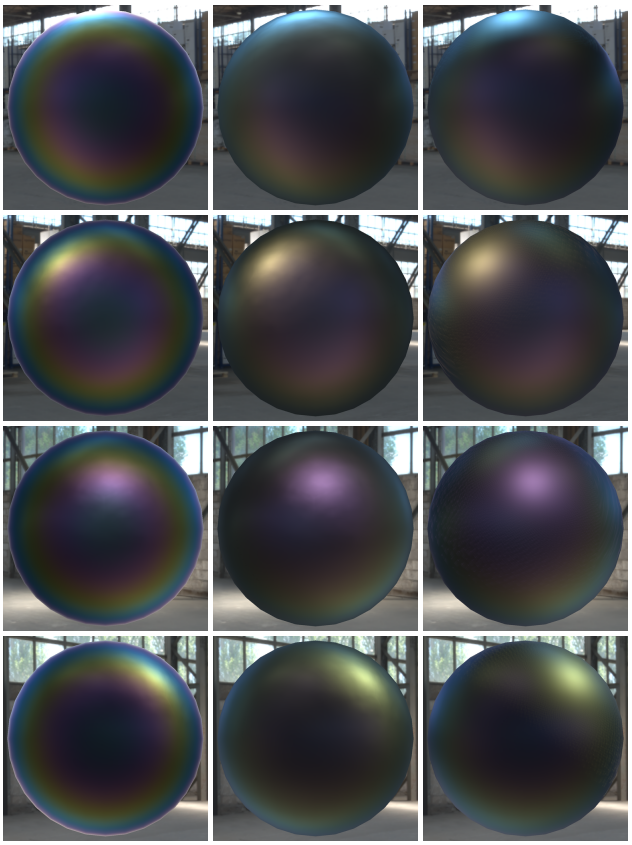

(a) Original Method (b) Our Method (c) Reference

Figure 4: *Rendering of the sphere model from different viewing directions. The surface roughness is* $\alpha = 0.25$*. The thin-film is* $d = 600\,\text{nm}$ *thick and has a refractive index of* $\eta_2 = 1.2$*. The refractive index of the base-material is* $\eta_3 = 2$ *and* $\kappa_3 = 0$*.*

evaluated for the ideal reflection direction [BB17]. In the figures and the rest of this text, we refer to their real-time approximation as the original method. For the reference images, we have drawn 16k samples per pixel from the environment map with importance sampling for the normal distribution function. All renderings have been created using the same red, green and blue spectral bands, where we used the measurements of Stiles and Burch in 1959 [CaU] as our spectral sensitivity functions.

On smooth surfaces, both, the original method and our method, produce results that are very close to the reference, since the light is mostly contributing from the ideal reflection direction, which is exactly the assumption used in the original method. This effect is reproduced in our method: Since the light is almost coming from a single direction, which results in a negligibly small covariance Σ_{ω_i} and the errors introduced by our assumptions are very small. On surfaces with increasing roughness, our proposed method is able to capture the expected desaturation effect, while the original method starts to produce over-saturated color fringes as illustrated in Figures 1 and 3.

Figure 4 shows a sphere, viewed from different directions. The pattern of iridescent colors, produced by the original method, is rather static on the surface and the brightness is modulated by the environment map. This is not the case in the reference rendering, where the iridescent colors depend on the actual light direction. Since our proposed method is able to better account for the real light directions through their mean and covariance, it produces colors very close to the reference. At shallow angles, noticeable distortions of the specular highlights are visible in the original and the propsed method, in comparison to the reference rendering. This artifact is independent of the iridescent reflectance term, and has already been noticed by Lagarde and De Rousiers [LDR14]. It is a consequence of the prefiltered environment map, which approximates the non-isotropic BRDF lobes by isotropic lobes, due to the $\omega_o = \mathbf{n} = \omega_r$ assumption.

The original method by Belcour and Barla [BB17] is intended to be used with varying material parameters. In Figure 5 and 6, we demonstrate this with our new extension, by using a texture map to define the thin-film thickness. Our method produces iridescent colors very close to the reference, while those produced by the original method are over-saturated and shifted, especially at increasing viewing angles.

In Figure 7, we compare the reflected light from a surface patch under uniform illumination, produced by different combinations of viewing angle θ_o, thin-film thickness d and surface roughness α. When the reflectance term is only evaluated for the ideal re-

(a) *Original Method* **(b)** *Our Method* **(c)** *Reference*

Figure 5: *The Erato model is rendered using a varying thin-film thickness* $d \in [413\,\text{nm}, 604\,\text{nm}]$*, modulated by Perlin noise. The base material has a refractive index of* $\eta_3 = 1.4$ *and* $\kappa_3 = 1.2$*, the thin-film refractive index is* $\eta_2 = 1.36$*, and the surface roughness is* $\alpha = 0.49$*.*

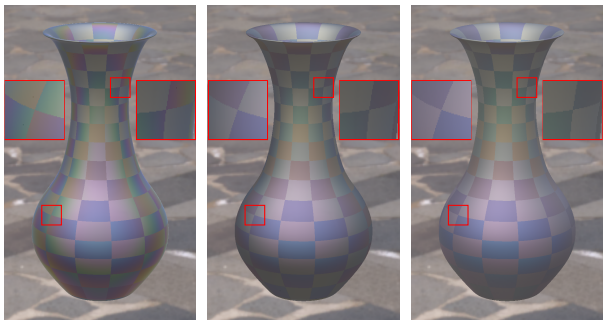

(a) *Original Method* **(b)** *Our Method* **(c)** *Reference*

Figure 6: *The Vase model is rendered using a varying thin-film thickness* $d \in [200\,\text{nm}, 700\,\text{nm}]$*, modulated by two vertical gradients masked by a checkerboard pattern. The base material has a refractive index of* $\eta_3 = 1.4$ *and* $\kappa_3 = 1.9$*, the thin-film refractive index is* $\eta_2 = 1.3$*, and the surface roughness is* $\alpha = 0.49$*.*

flection direction, as done in the original method, the color gradient of the reflected light does not depend on the surface roughness. This introduces noticeable errors, especially at grazing angles. Overall, the color gradients are well approximated by our proposed method. However, our colors are slightly less saturated at large surface roughnesses. Using the ground truth mean and variance on the optical path difference, produces results closer to the reference. Comparing columns (b), (c) and (f), (g) in Figure 7 shows where the linearized transformation from the light directions to the optical path difference introduces additional errors.

The performance of our proposed method is compared to the original method in Table 1 for different scenes. We consider two artificial edge cases for the performance evaluation: When rendering a flat plane, that fills the entire screen, we only have a very

Scene	Original Method (ms)	Our Method (ms)
Flat Plane	3.08 ± 0.14	3.19 ± 0.13
Random Plane	4.04 ± 0.13	5.41 ± 0.16
Beethoven	3.49 ± 0.13	3.62 ± 0.16
Erato	3.31 ± 0.13	3.46 ± 0.13

Table 1: *We measured the runtime of the original and our proposed method in different scenes. The measurements have been taken using the Vulkan backend on Ubuntu 16.04, running an i7-4790k CPU and Nvidia GTX 1080 GPU. The screen resolution was set to* 2560×1440 *pixels. In the Flat Plane scene, a plane fills the entire screen to maximize the number of shaded pixels. The Random Plane scene additionally has random surface normals at each shaded pixel, to provide a worst-case scenario where the texture data cannot be cached efficiently. The Beethoven scene is shown in Figure 1 and the Erato scene is shown in Figure 5. For comparison, an empty scene with just the environment map takes* (2.63 ± 0.26) ms *to render.*

small runtime disadvantage in comparison to the original method. A worst-case situation is simulated, by randomizing the surface normals of this plane. Each pixel now samples the environment map at a different location, and also the prefiltered Fourier transformed spectral sensitivity table is accessed incoherently by neighboring pixels on the screen. With this modification, the GPU is not able to cache the texture data efficiently, which already reduces the performance of the original method. However, our method reads a much larger chunk of data from the environment map, which reduces the performance by a larger amount. As more realistic settings, we use the Beethoven bust model and the Erato model shown in Figures 1 and 5, respectively. In these settings, the runtime differences between our method and the original method are low, while the visual quality is improved.

6. Conclusion

We presented an improved method for real-time image-based lighting of rough microfacet surfaces, covered with a thin-film that causes iridescent effects. Previously, the reflectance term was evaluated for the ideal reflection direction only [BB17], which is a reasonable approximation for smooth surfaces. On rough surfaces however, this no longer the case. Our key observation is, that the incident light is not only contributing from the ideal reflection direction. Depending on the environment map and surface roughness, the light is contributing from various directions, which mixes different iridescent colors.

We factorize the rendering integral into three parts, which are handled individually. The prefiltered BRDF and prefiltered environment map are already well known from previous work [Kar13, LDR14]. Our contribution lies in the newly emerging prefiltered reflectance term. We prefilter the mean and covariance of the distribution of light directions, shaped by the environment map and surface roughness. This information is stored in nine additional texture channels of the environment map. Besides the additional memory requirement, this increases the runtime slightly, due to the additional texture lookups. Yet, our method allows real-time

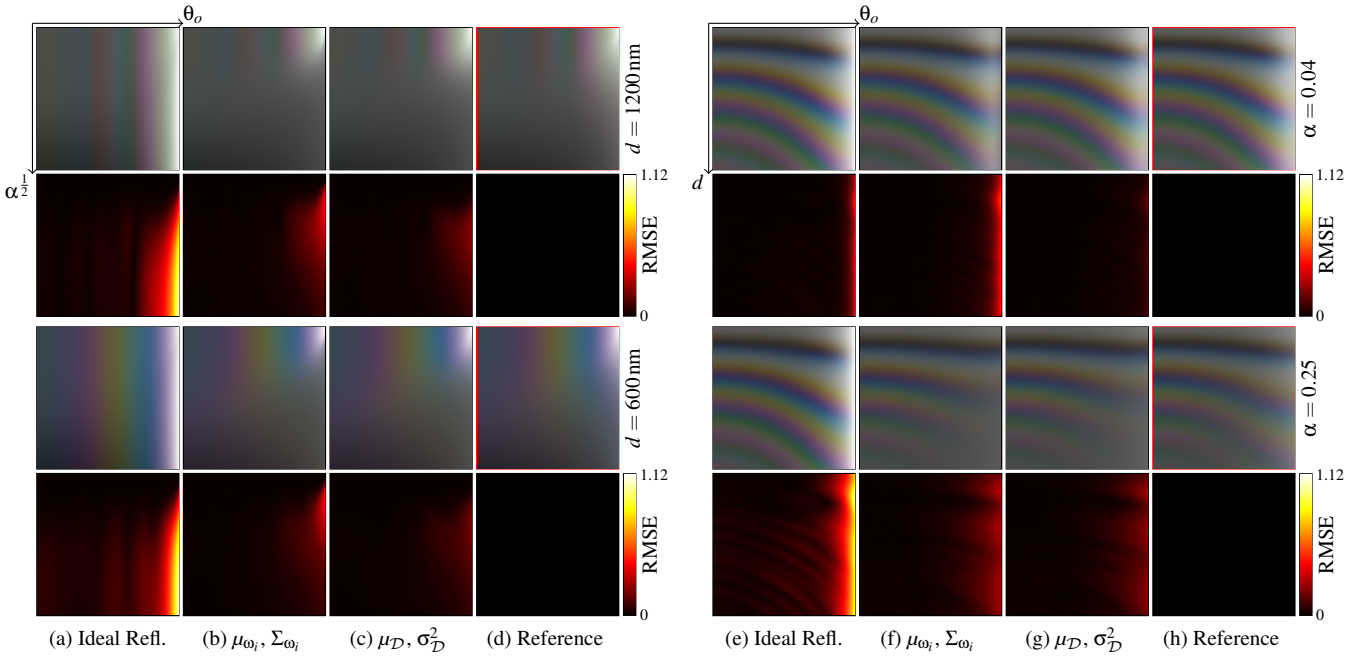

Figure 7: *The reflected light is visualized for varying viewing directions* $\theta_o \in [0, \frac{\pi}{2}]$*, perceptual surface roughness* $\sqrt{\alpha} \in [0,1]$ *in (a) to (d) and thin-film thickness* $d \in [0\,\text{nm}, 1000\,\text{nm}]$ *in (e) to (h) under uniform illumination. In (a) and (e), the reflectance term is evaluated for the ideal reflection direction only, and is therefore independent of* α*. In (b) and (f), our method was used with ground truth* μ_{ω_i} *and* Σ_{ω_i}*. In (c) and (g), the ground-truth mean* $\mu_{\mathcal{D}}$ *and variance* $\sigma^2_{\mathcal{D}}$ *are used with our method, omitting the non-linear transformations. The ground-truth reflectance is shown in (d) and (h).*

framerates. From this mean and covariance, we derive a normal distribution on the optical path difference, that is used to sample the prefiltered Fourier transformed spectral sensitivity. The Fourier transformed sensitivity is replaced by a two-dimensional lookup table, which is parameterized by the mean and variance of the optical path difference instead of a single value of interest. Belcour and Barla [BB17] use the same prefiltered data for anti-aliasing of rapidly varying thin-film thicknesses, which can be easily integrated with our method.

We are able to accurately handle individual discrete embedded light sources in the environment map, like the sun. Under unfortunate lighting conditions however, our method tends to overdo the desaturation of the iridescent colors a little.

In summary, our method provides a higher rendering quality of iridescent effects on rough surfaces than previous real-time approaches, while having only a small negative impact on performance.

Acknowledgements

The authors would like to thank the anonymous reviewers for their comments and suggestions. The Erato and Tissot bust model have been provided by Geoffrey Marchal under CC BY-NC 4.0 license.

References

[BB17] BELCOUR L., BARLA P.: A practical extension to microfacet theory for the modeling of varying iridescence. *ACM Transactions on Graphics (TOG) 36*, 4 (2017), 65. 1, 2, 3, 4, 5, 6, 7, 8

[CaU] COLOUR, AT UCL V. R. L.: Colour matching functions. URL: http://cvrl.ioo.ucl.ac.uk/cmfs.htm. 6

[CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *ACM Transactions on Graphics (TOG) 1*, 1 (1982), 7–24. 2, 3

[DTS*14] DHILLON D. S., TEYSSIER J., SINGLE M., GAPONENKO I., MILINKOVITCH M. C., ZWICKER M.: Interactive diffraction from biological nanostructures. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 177–188. 2

[GH03] GRANIER X., HEIDRICH W.: A simple layered rgb brdf model. *Graphical Models 65*, 4 (2003), 171–184. 2

[HKY*01] HIRAYAMA H., KANEDA K., YAMASHITA H., YAMAJI Y., MONDEN Y.: Visualization of optical phenomena caused by multilayer films based on wave optics. *The Visual Computer 17*, 2 (2001), 106–120. 2

[HP17] HOLZSCHUCH N., PACANOWSKI R.: A two-scale microfacet reflectance model combining reflection and diffraction. *ACM Transactions on Graphics (TOG) 36*, 4 (2017), 66. 2

[IA99] ICART I., ARQUÈS D.: An illumination model for a system of isotropic substrate-isotropic thin film with identical rough boundaries. In *Rendering TechniquesâĂŹ99*. Springer, 1999, pp. 261–272. 2

[IA00] ICART I., ARQUÈS D.: A physically-based brdf model for multilayer systems with uncorrelated rough boundaries. In *Rendering Techniques 2000*. Springer, 2000, pp. 353–364. 2

[IOS*09] IMURA M., OSHIRO O., SAEKI M., MANABE Y., CHIHARA K., YASUMURO Y.: A generic real-time rendering approach for structural colors. In *Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology* (2009), ACM, pp. 95–102. 2

[Kaj86] KAJIYA J. T.: The rendering equation. In *ACM Siggraph Computer Graphics* (1986), vol. 20, ACM, pp. 143–150. 3

[Kar13] KARIS B.: Real shading in unreal engine 4. *PPhysically Based Shading Theory Practice – SIGGRAPH Courses* (2013), 621–635. 2, 4, 7

[KVHS00] KAUTZ J., VÁZQUEZ P.-P., HEIDRICH W., SEIDEL H.-P.: A unified approach to prefiltered environment maps. In *Rendering Techniques 2000*. Springer, 2000, pp. 185–196. 2

[LDR14] LAGARDE S., DE ROUSIERS C.: Moving frostbite to PBR. *Physically Based Shading Theory Practice – SIGGRAPH Courses* (2014). 2, 4, 6, 7

[Sch94] SCHLICK C.: An Inexpensive BRDF Model for Physically-based Rendering. *Computer Graphics Forum 13*, 3 (1994), 233–246. 3

[SM92] SMITS B. E., MEYER G. W.: Newton's colors: simulating interference phenomena in realistic image synthesis. In *Photorealism in Computer Graphics*. Springer, 1992, pp. 185–194. 2

[Sun06] SUN Y.: Rendering biological iridescences with rgb-based renderers. *ACM Transactions on Graphics (TOG) 25*, 1 (2006), 100–129. 2

[TG17a] TOISOUL A., GHOSH A.: Practical acquisition and rendering of diffraction effects in surface reflectance. *ACM Transactions on Graphics (TOG) 36*, 5 (2017), 166. 2

[TG17b] TOISOUL A., GHOSH A.: Real-time rendering of realistic surface diffraction with low rank factorisation. In *Proceedings of the 14th European Conference on Visual Media Production (CVMP 2017)* (2017), ACM, p. 2. 2

[VWH18] VELINOV Z., WERNER S., HULLIN M. B.: Real-time rendering of wave-optical effects on scratched surfaces. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 123–134. 2

[WVJH17] WERNER S., VELINOV Z., JAKOB W., HULLIN M. B.: Scratch iridescence: Wave-optical rendering of diffractive surface structure. *ACM Transactions on Graphics (TOG) 36*, 6 (2017), 207. 2

[YHW*18] YAN L.-Q., HAŠAN M., WALTER B., MARSCHNER S., RAMAMOORTHI R.: Rendering specular microgeometry with wave optics. *ACM Transactions on Graphics (TOG) 37*, 4 (2018), 75. 2