On the Accuracy of Surface Scattering Theories

Matthew Avolio[1], Eugene D'Eon[2], Shlomi Steinberg[1]

[1]University of Waterloo [2]Nvidia

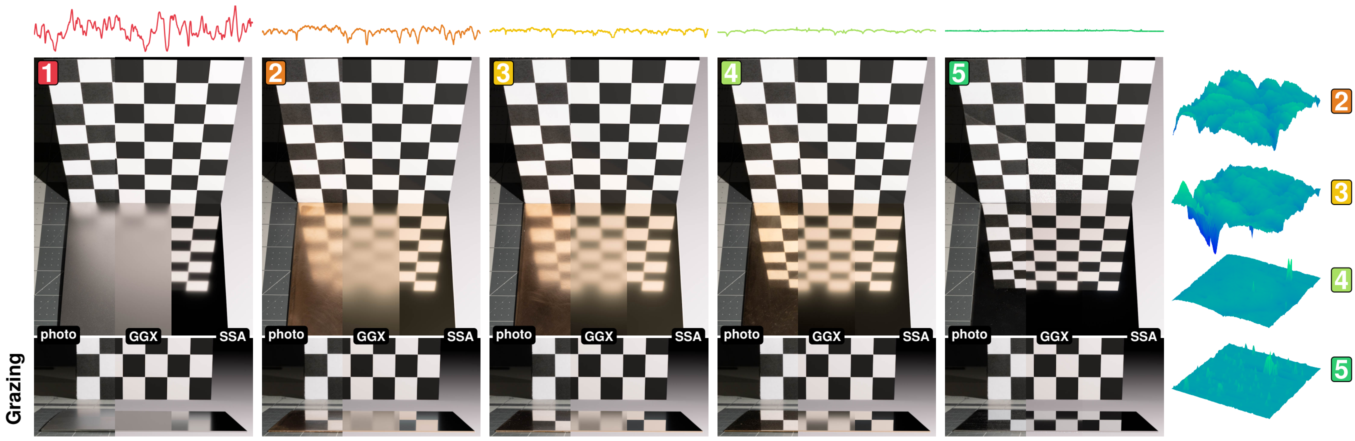

Figure 1: *Photographs compared to rendered microfacet GGX and smooth-surface (SSA) GHS predictions from measured surface geometry. Surface patches measured using atomic force microscopy (AFM) and profilometer scans are shown as on the right and the top, respectively.*

Abstract

Surface scattering models are typically validated by fitting measured reflectance data. However, predictive rendering requires deriving appearance directly from surface geometry. We test the predictive accuracy of Microfacet and Generalized Harvey-Shack (GHS) theories by comparing renderings derived strictly from measured surface profiles (AFM/profilometry) against photographs of metal samples. We demonstrate that no current model succeeds across all angles; specifically, Microfacet theory fails at grazing angles where optical roughness vanishes. Additionally, we analytically link the Trowbridge-Reitz (GGX) distribution to the K-correlation PSD. This suggests that the popular "GGX look" arises from wave-optical effects on Gaussian (i.e. Beckmann NDF) surfaces, rather than suggesting an underlying non-Gaussian geometry.

CCS Concepts

• ***Computing methodologies*** → ***Reflectance modeling;***

1 Introduction

Predicting the optical appearance of surfaces has been a fundamental problem in rendering and optics. Over the years, different surface-scattering theories have emerged that use statistical descriptions of surface profiles to predict material appearance via the bidirectional reflectance distribution function (BRDF). Different applications employ such theories for reproducing and rendering the visual appearance of a modelled surface, or the inverse task where the surface profile is predicted from its visual appearance (i.e., material acquisition). These applications often depend upon the underlying scattering model's accuracy: its ability to connect the appearance—the BRDF—to the surface profile.

Microfacet theory has gained popularity, especially so in computer graphics. This theory understands surfaces as a distribution of oriented independent mirror-like facets, and explains scattering via a geometric-optics process where light rays are reflected by these facets. Popular wave optics-based theories that have emerged in optical literature include the *small surface perturbation* (SPM) models, which are often limited to the 1st order due to computational difficulties, and the related *Harvey-Shack* model, which has also found its way into computer graphics. Many other models exist, but are either computationally intractable—do not admit closed-form expressions for the BRDF—more restricted in the surfaces they are able to simulate, or less accurate.

In this paper, we analyze the ability of these scattering theories to correctly predict the appearance of the common surfaces that we

DOI: 10.2312/egs.20261024

encounter in real life, and wish to render or capture their appearance. We take accurate, high-resolution scans of several metallic surfaces—varying in roughness from almost mirror-like smooth, to very rough—analyze their surface profile statistics, and study and compare the appearance predicted by scattering theories to their measured real-world appearance. Our primary contribution is showcasing that modern surface scattering theories are only able to accurately predict appearance in very narrow domains (e.g., only for a very limited roughness range), far more limited compared with the typical ranges of surfaces properties that arises in real life and we often wish to model and render.

We also make a novel analytic connection between the K-correlation model and the Trowbridge–Reitz microfacet model. This connection explains why both models predict a similar, perceptually-realistic, optical response; however, they model very different physics, and, as we show, only perform well in opposing roughness domains.

1.1 Related Work Dong et al. [DWMG16] use surface profiles acquired via an optical profilometer scans to study the appearance reproduction of the microfacet BRDF and the Kirchoff scattering theory, but limited their analysis to only very rough surfaces. The Kirchoff integral better predicted the appearance of the surfaces, while the microfacet model overestimated the scattering, consistent with our results. However, the Kirchoff integral admits no closed-form solution, and is not practical for most applications; no comparison to a practical wave-optical scattering model was performed.

Full-wave surface scattering solvers has been proposed [YXW*23] but are only computationally feasible for tiny surface patches and work on *explicit* surface profiles. The Yu et al. [YXW*23] solver was unable to simulate BRDFs for our 93×93 µm AFM scans (shown in Fig. 1) due to memory constraints on a modern workstation computer (they use 24×24 µm patches); and yet our AFM scans are still not large enough to fully capture the samples' surface statistics, hence we complement them with stylus scans. In this work, we analyze statistical surface scattering models that give rise to BRDFs that can be used in practice for rendering.

2 Background

2.1 Scalar Diffraction Theory

The Generalized Harvey Shack (GHS) theory [Kry06] is a modern, non-paraxial formulation of wave scattering from statistical rough surfaces. Scattering is driven by phase variations induced by the optical path differences caused by light reflecting by the surface heightfield—modelled as a random process—and the scattering behaviour is fully determined by the surface's second-order statistics.

Let $h(x)$ be a zero-mean, wide-sense stationary random process describing the surface height at a lateral position $x \in \mathbb{R}$. The height variance is defined as $\sigma_h^2 = \langle h(x)^2 \rangle$, with $\langle \cdot \rangle$ denoting the ensemble average. The random process's autocovariance (ACV), defined as $C(r) = \langle h(x)\, h(x+r) \rangle$, describes how surface heights are correlated as a function of spatial separation r, and the important *power spectral density* (PSD) is the Fourier transform of the ACV:

$$P(\zeta) = \int_{-\infty}^{\infty} C(r) e^{-2\pi i r \zeta} \, \mathrm{d}r \ . \tag{1}$$

Extension to a two-dimensional surface is trivial.

K-Correlation PSD Real surfaces admit a variety of statistics. In practice, many rough surfaces of interest exhibit fractal (self-affine) statistics, with features spanning multiple spatial scales as a result of successive grinding and polishing processes. These kind of surfaces are well described by the popular K-correlation (KC) PSD model [Kry06] (for a 2D surface with isotropic statistics):

$$P_{\mathrm{KC}}\left(\vec{\zeta}\right) = \frac{L_c^2 (\gamma - 1) \sigma_h^2}{2\pi \left[1 + L_c^2 |\vec{\zeta}|^2\right]^{\frac{\gamma+1}{2}}} \ . \tag{2}$$

L_c is the low-frequency roll-off length and γ is the fractal (Hurst) exponent quantifying the fractal PSD's falloff, i.e. its log-log slope.

BRDF Under the GHS formulation [Kry06], the BRDF is defined as the Fourier transform of the *surface transfer function*:

$$\mathrm{BRDF}_{\mathrm{GHS}}\left(\hat{\mathbf{i}}, \hat{\mathbf{o}}\right) = R \cdot \mathcal{F}\left\{ \exp\left\{ g[C(\vec{\mathbf{x}}, \vec{\mathbf{y}})/\sigma_h^2 - 1] \right\} \right\} , \tag{3}$$

$$\text{with} \quad g = [k\sigma_h(\gamma_i + \gamma_o)]^2 \quad \text{and} \quad k = 2\pi/\lambda \ . \tag{4}$$

R is the reflectance; $\vec{\mathbf{x}}, \vec{\mathbf{y}}$ are surface spatial coordinates, with respect to which the Fourier transform $\mathcal{F}$ is taken; $\hat{\mathbf{i}}$ and $\hat{\mathbf{o}}$ are the incident and outgoing directions; and, given a (macro) surface normal $\hat{\mathbf{n}}$, $\gamma_i = \hat{\mathbf{i}} \cdot \hat{\mathbf{n}}$ and $\gamma_o = \hat{\mathbf{o}} \cdot \hat{\mathbf{n}}$ are the directional cosines. The variable g plays the role of a roughness indicator, and depends on the surface properties, and the radiation and observation setting. k and λ are the light's wavenumber and wavelength, respectively. Similar to Stam [Sta99], the above can be rewritten as an exponential series:

$$\mathrm{BRDF}_{\mathrm{GHS}}\left(\hat{\mathbf{i}}, \hat{\mathbf{o}}\right) = R \sum_{m=0}^{\infty} \frac{e^{-g} g^m}{m!} \mathcal{F}\left\{ \tfrac{1}{\sigma_h^2} C(\vec{\mathbf{x}}, \vec{\mathbf{y}})^m \right\} \left(\vec{\zeta}^{\perp}\right) , \tag{5}$$

where $\vec{\zeta} = k(\hat{\mathbf{o}} + \hat{\mathbf{i}})$ can be understood as the wave-optics analogue of the geometric half vector, and $\vec{\zeta}^{\perp}$ denotes its projection onto the surface plane. Applying the convolution theorem and simplifying yields a Poisson-weighted self-convolutions of the PSD:

$$\mathrm{BRDF}_{\mathrm{GHS}}\left(\hat{\mathbf{i}}, \hat{\mathbf{o}}\right) = R \sum_{m=0}^{\infty} \frac{e^{-g} g^m}{\sigma_h^2 m!} P\left(\vec{\zeta}^{\perp}\right)^{\star m} . \tag{6}$$

The notation $P^{\star m}$ is the m-fold self convolutions of the PSD P. The $m = 0$ term yields a Dirac delta, i.e. the specular reflection lobe.

Smooth Surface Approximation For sufficiently smooth surfaces, where σ_h is much smaller than wavelength and thus $g \ll 1$, the $m > 1$ terms become negligible, giving rise to the more common form of the GHS BRDF, its *smooth surface approximation* (SSA):

$$\mathrm{BRDF}_{\mathrm{SSA}}\left(\hat{\mathbf{i}}, \hat{\mathbf{o}}\right) = RA\, \delta(\hat{\mathbf{s}} - \hat{\mathbf{o}}) + R \tfrac{1-A}{\sigma_h^2} P\left(\vec{\zeta}^{\perp}\right) . \tag{7}$$

This is a two lobe BRDF with specular (delta) and scattered components weighted by $A = e^{-g}$, with $\hat{\mathbf{s}}$ being the specular direction.

2.2 Microfacet Theory

Recall the well known isotropic GGX (Trowbridge-Reitz) model [WMLT07], with its *normal distribution function* (NDF), and corresponding microfacet BRDF defined as:

$$D_{\mathrm{GGX}}(\theta_h) = \frac{\alpha^2}{\pi\left[(\alpha^2 - 1)\cos^2\theta_h + 1\right]^2} , \tag{8}$$

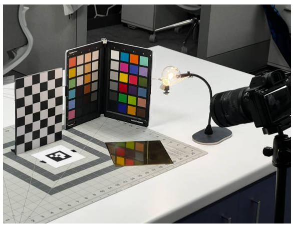

Figure 2: *Photo setup.*

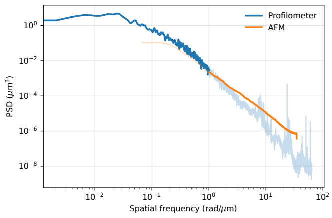

Figure 3: *Sample 2 PSD.*

$$\mathrm{BRDF}_{\mathrm{GGX}}\left(\hat{\mathbf{i}},\hat{\mathbf{o}}\right) = \frac{F(\hat{\mathbf{i}},\hat{\mathbf{h}})\,G(\hat{\mathbf{i}},\hat{\mathbf{o}})\,D_{\mathrm{GGX}}(\theta_h)}{4\,|\hat{\mathbf{i}}\cdot\hat{\mathbf{n}}|\,|\hat{\mathbf{o}}\cdot\hat{\mathbf{n}}|}\,, \tag{9}$$

respectively, where the roughness parameter α controls the width of the distribution, $\theta_h \geq 0$ denotes the angle between the geometric half vector $\hat{\mathbf{h}}$ and the macrosurface normal $\hat{\mathbf{n}}$, and F and G are the Fresnel and Smith shadowing-masking terms.

2.3 The Two-Scale Model

Holzschuch and Pacanowski's [HP17] two-scale model combines a geometric microfacet model with the diffractive GHS SSA by overlaying a KC SSA BRDF over each microfacet:

$$\mathrm{BRDF}_{\mathrm{2scale}} = RA\,\delta(\hat{\mathbf{s}}-\hat{\mathbf{o}}) + RG(\hat{\mathbf{i}},\hat{\mathbf{o}})(1-A)(P_{\mathrm{KC}}\star D_{\mathrm{exp}})(\theta_h) \tag{10}$$

with A as in Eq. (7), D_{exp} being the exponential micorfacet NDF [HP17], and $P_{\mathrm{KC}}\star D_{\mathrm{exp}}$ its spherical convolution with the KC PSD.

3 Methodology

We aim to analyze the appearance prediction accuracy of the microfacet theory, scalar diffraction theory (GHS), and the two-scale model. To that end we fabricate several metal surfaces, capture their surface profiles, and study their appearance.

Samples We fabricated 150×150 mm metal sheets at varying roughnesses and root-mean square (RMS) heights. A representative subset is used for our analysis, see Fig. 1: samples 1 and 5 are stainless steel treated with industry standard glass bead-blasting and mirror-finish, respectively; samples 2,3,4 are brass, progressively treated by hand with finer abrasives. The RMS roughness of the samples are: 808, 280, 138, 91, and 27 nm respectively.

Surface Metrology Two different surface characterization metrology instruments were used to capture data at different scales:

1. Highly accurate 93x93 μm *atomic force microscope* (AFM) scans of the surface at a 1024x1024 pixels resolution.
2. 6 mm scans using a stylus profilometer at a 50 nm step size.

The stylus complements the AFM scans with low frequency surface data captured over distances greater than possible with the AFM.

To fit a PSD from the scanned data, we compute a 1D PSD for each AFM or stylus scan line using Welch's method [Wel67] and averaged to reduce noise. The final surface PSD is a piecewise function of both the AFM and stylus data, combined at a frequency cutoff of 1 rad/μm, see Fig. 3. This cutoff was chosen ad hoc, and works well for all our samples. Sample 1 was too rough for AFM measurement, so its PSD relies only on profilometer data.

The microfacet NDFs are captured from the profilometer data by making a histogram of normals for each segment weighted by the projected area and angular Jacobian.

Rendering Rendering was done using Mitsuba [JSR*22]. For the microfacet model renderings, we used GGX NDFs fitted to the measured NDFs, smoothed using the heuristic Gaussian filtering method from Dong et al. [DWMG16].

To render the full GHS model, we introduce a new importance sampling strategy where a diffraction order m is sampled from the Poisson PMF with mean g. Then we sample the m-fold self-convolution $P(\vec{\zeta}^{\perp})^{\star m}$ using a KC PSD fitted to the corresponding convolution, following the same importance-sampling strategy as Holzschuch et al. [HP17]. Rendering the smooth surface approximation, Eq. (7), is similar. See our supplemental material for more information on the rendering process and accompanying code.

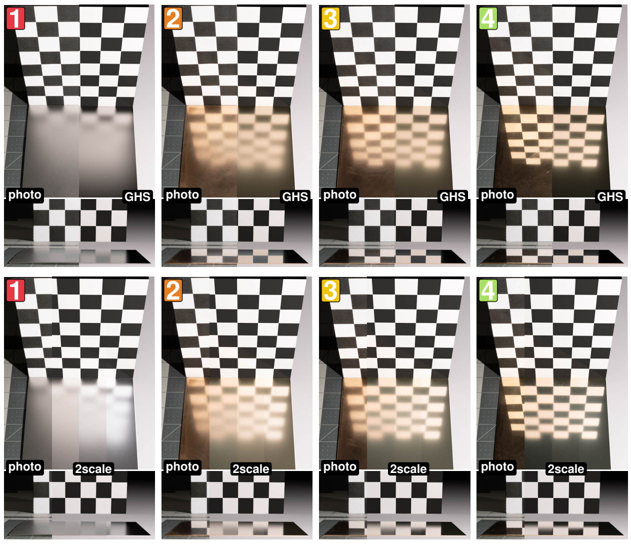

Figure 4: *Photos compared to full GHS and two-scale model predictions. The two-scale splits show the effect of varying the macro/nano-roughness cutoff at periods of 50λ, 100λ, and 150λ.*

4 Discussion and Conclusion

In Fig. 1 and 4 we compare the surface scattering appearance predicted by the microfacet theory, smooth-surface approximation (SSA) of GHS, full-order GHS and the two-scale model to the physical samples, and show that no current statistical scattering model accurately reproduces surface appearance across a wide roughness range and incidence and scattering angles. Microfacet theory consistently predicts BRDFs that are far too rough; with one exception being the very rough sample 1, for which the assumptions of geometric optics are more appropriate. Even so, it still diverges at grazing angles. SSA GHS vastly under predicts scattering for all but the smoothest samples, as expected from its explicit smooth-surface approximation. Full GHS with all diffraction orders yields noticeable improvements across roughnesses and viewing angles, but still falls short of reproducing the appearance, and is more complicated to evaluate and sample.

In addition, note that the geometric optics-based microfacet

Proceedings published by Eurographics - The European Association for Computer Graphics.

model drastically fails to predict appearance at grazing angles (Fig. 1). At grazing angles, surfaces appear smoother, as predicted by the GHS model: the parameter g that partitions between the terms in Eq. (6) depends on the light's wavelength and incident/scattering cosines. Thus, *the effective roughness is not an intrinsic property of surface*, but depends on the light and the observer. Although well documented in optics [Ogi91], this phenomenon is not captured in microfacet theory without very particular surfaces (e.g., lacquer on wood, non-uniform heightfields). As expected, the wave optics-based models perform better at grazing angles.

The Physical Origin of the GGX Distribution We identify a novel analytical link between the microfacet GGX distribution and the wave-optical K-correlation (KC) model. From its definition, note that $\vec{\zeta} = k(\hat{\mathbf{o}} + \hat{\mathbf{i}})$ must coincide with the geometric half vector $\hat{\mathbf{h}}$. We wish to study the specular highlight of the BRDF, i.e. where $\hat{\mathbf{o}}^{\perp} \approx -\hat{\mathbf{i}}^{\perp}$. Then $\vec{\zeta} \approx 2k(\hat{\mathbf{i}} \cdot \hat{\mathbf{n}})\hat{\mathbf{h}}$, which implies $\cos^2\theta_h = 1 - \sin^2\theta_h = 1 - \frac{|\vec{\zeta}^{\perp}|^2}{4k^2(\hat{\mathbf{i}}\cdot\hat{\mathbf{n}})^2}$, which together with Eq. (8) yields:

$$D_{\text{GGX}}(\theta_h) = \frac{1}{\pi\alpha^2 \left[1 + \frac{1-\alpha^2}{\alpha^2} \frac{1}{4k^2(\hat{\mathbf{i}}\cdot\hat{\mathbf{n}})^2} |\vec{\zeta}^{\perp}|^2\right]^2} . \qquad (11)$$

The above is analytically identical to the KC PSD (Eq. (2)) with:

$$\gamma = 3 , \qquad L_c^2 = \frac{1-\alpha^2}{\alpha^2} \frac{1}{4k^2(\hat{\mathbf{i}}\cdot\hat{\mathbf{n}})^2} , \qquad \sigma_h^2 = \frac{4k^2(\hat{\mathbf{i}}\cdot\hat{\mathbf{n}})^2}{1-\alpha^2} .$$

The above means that the appearance and shape of the specular highlight is identical between the SSA GHS BRDF (Eq. (7)) with the KC PSD and the GGX microfacet BRDF (Eq. (9)). Indeed, as $\alpha \to 0$ (smooth surfaces), the correlation length $L_c \to \infty$; while as $\alpha \to 1$ (extremely rough surfaces), $L_c \to 0$ and $\sigma_h \to \infty$.

This connection resolves a key inconsistency in appearance modelling. Since Walter et al. [WMLT07], GGX has displaced the Beckmann distribution due to its realistic "heavy" tails. In geometric optics, this implies the physical surface must have a Trowbridge-Reitz slope distribution; however, constructing continuous heightfields with such statistics is notoriously difficult, yielding physically implausible spikes. Our metrology confirms that real finished metals are typically Gaussian processes (which implies a Beckmann slope distribution). Equation 11 explains the paradox: the "heavy tailed" look does not require a heavy-tailed surface. It arises naturally from *wave optics acting on a standard Gaussian surface*. Thus, GGX serves as an effective, if unintentional, proxy for diffractive effects that the geometric Beckmann model ignores.

Macro- and Nano-Roughness in the Two-Scale Model All of our surfaces have RMS heights on the order of a wavelength and has a continuous PSD showing no strict break in scale (Fig. 3), making the separation between macro and nano-geometry ill-defined. Two-scale models are most appropriate when surface roughness can be decomposed into well-separated spatial-frequency bands that arise from distinct physical processes, as in ocean wave and ripple models used in remote-sensing [Ogi91].

To obtain the two-scale parameters, we low-pass filter each surface profile with a cutoff spatial period at different multiples of the mean wavelength $\lambda = 550$ nm. Then, fit the resulting large-scale geometry to an exponential NDF D_{exp}, and use the residual PSD to parameterize the nano-scale roughness. The choice of cutoff (see Fig. 4) has a large impact on the predicted appearance, and the best cutoff can differ between samples, however there is no physical basis in determining which cutoff is "correct".

While the two-scale model can be fitted to the reproduce the appearance of surfaces across the MERL database, our results suggest that it does so largely due of its numerous free parameters, rather than because it more accurately represents the underlying surface statistics and scattering physics.

Conclusion The lack of a statistical surface scattering model that is accurately able to reproduce surface appearance—across a sufficiently wide range of surfaces, roughnesses and viewing angles—is a limiting factor for predictive rendering, inverse rendering, and light transport simulations across other regimes in the electromagnetic spectrum (e.g., acoustics, radio or other long-wavelength simulations). Microfacet theory is inherently not fit for such purposes as it neglects the wave-optical nature of light or any well-defined sense of scale. Scalar diffraction theories (like GHS) presents a solid foundation for future work due to its stronger relationship with measured surface statistics and physical-optics formulation. We conclude that predicting surface scattering from statistically rough surfaces remains a difficult problem in computer graphics and optics, and continues to warrant further research.

References

[DWMG16] DONG Z., WALTER B., MARSCHNER S., GREENBERG D. P.: Predicting appearance from measured microgeometry of metal surfaces. *ACM Trans. Graph. 35*, 1 (Dec. 2016). URL: https://doi.org/10.1145/2815618, doi:10.1145/2815618. 2, 3

[HP17] HOLZSCHUCH N., PACANOWSKI R.: A two-scale microfacet reflectance model combining reflection and diffraction. *ACM Trans. Graph. 36*, 4 (July 2017). URL: https://doi.org/10.1145/3072959.3073621, doi:10.1145/3072959.3073621. 3

[JSR*22] JAKOB W., SPEIERER S., ROUSSEL N., NIMIER-DAVID M., VICINI D., ZELTNER T., NICOLET B., CRESPO M., LEROY V., ZHANG Z.: Mitsuba 3 renderer, 2022. https://mitsuba-renderer.org. 3

[Kry06] KRYWONOS A.: Predicting surface scatter using a linear systems formulation of non-paraxial scalar diffraction. URL: https://api.semanticscholar.org/CorpusID:118031406. 2

[Ogi91] OGILVY J. A.: *Theory of wave scattering from random rough surfaces*. A. Hilger, Bristol, England ;, 1991. 4

[Sta99] STAM J.: Diffraction shaders. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques* (USA, 1999), SIGGRAPH '99, ACM Press/Addison-Wesley Publishing Co., p. 101–110. URL: https://doi.org/10.1145/311535.311546, doi:10.1145/311535.311546. 2

[Wel67] WELCH P.: The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics 15*, 2 (1967), 70–73. doi:10.1109/TAU.1967.1161901. 3

[WMLT07] WALTER B., MARSCHNER S. R., LI H., TORRANCE K. E.: Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques* (Goslar, DEU, 2007), EGSR'07, Eurographics Association, p. 195–206. 2, 4

[YXW*23] YU Y., XIA M., WALTER B., MICHIELSSEN E., MARSCHNER S.: A full-wave reference simulator for computing surface reflectance. *ACM Trans. Graph. 42*, 4 (July 2023). URL: https://doi.org/10.1145/3592414, doi:10.1145/3592414. 2