Supplemental Document: Designing and Fabricating Color BRDFs with Differentiable Wave Optics

YIXIN ZENG, Zhejiang University, China
KISEOK CHOI, KAIST, South Korea
HADI AMATA, KAUST, Saudi Arabia
KAIZHANG KANG, Zhejiang University, China
WOLFGANG HEIDRICH, KAUST, Saudi Arabia
HONGZHI WU, Zhejiang University, China
MIN H. KIM, KAIST, South Korea

ACM Reference Format:
Yixin Zeng, Kiseok Choi, Hadi Amata, Kaizhang Kang, Wolfgang Heidrich, Hongzhi Wu, and Min H. Kim. 2025. Supplemental Document: Designing and Fabricating Color BRDFs with Differentiable Wave Optics. *ACM Trans. Graph.* 44, 6, Article 240 (December 2025), 5 pages. https://doi.org/10.1145/3763275

A LIST OF SYMBOLS

Table 1 summarizes all mathematical symbols used in the main paper. For the height map $H(s)$, Figure 1 shows a visualization of it.

B LIST OF TRAINING AND RENDERING PARAMETERS

Table 2 lists training and rendering parameters for figures in the main paper. Unless specified, all ablations use their corresponding BRDF's default parameters in this table.

C FABRICATION-AWARE RENDERING

Coherence-aware patch padding. We define each BRDF patch as a 32×32 feature height map, with feature sizes of either 1.5 or 2.0 μm, resulting in physical patch sizes of 48–64 μm. In fabrication, these patches are tiled to create a macroscopic 50 mm × 50 mm wafer for observation. In contrast, prior work [Levin et al. 2013] used larger patches (112 μm), which exceed the typical coherence radius and thus permit independent optimization per tile. However, under realistic lighting conditions–e.g., an incandescent bulb several meters away—the coherence radius can reach 8 μm at a wavelength of 500 nm [Levin et al. 2013], which is comparable to our patch size. This introduces spatial overlap between coherence areas of adjacent patches, violating the independence assumption and causing inter-patch interference. To account for this, we introduce coherence-aware patch padding in our simulation. Specifically, each patch is surrounded by a one-patch border filled with replicated height data, ensuring coherence regions are fully sampled. Rendering is performed only in the center patch to isolate valid BRDF values, as illustrated in Figure 5(c) in the main paper.

Authors' Contact Information: Yixin Zeng, Zhejiang University, Hangzhou, China, 22221238@zju.edu.cn; Kiseok Choi, KAIST, Daejeon, South Korea, kschoi@vclab.kaist.ac.kr; Hadi Amata, KAUST, Thuwal, Saudi Arabia, hadi.amata@kaust.edu.sa; Kaizhang Kang, Zhejiang University, Hangzhou, China, cocoa_kang@zju.edu.cn; Wolfgang Heidrich, KAUST, Thuwal, Saudi Arabia, wolfgang.heidrich@kaust.edu.sa; Hongzhi Wu, Zhejiang University, Hangzhou, China, hongzhi.wu@gmail.com; Min H. Kim, KAIST, Daejeon, South Korea, minhkim@kaist.ac.kr.

ACM 1557-7368/2025/12-ART240
https://doi.org/10.1145/3763275

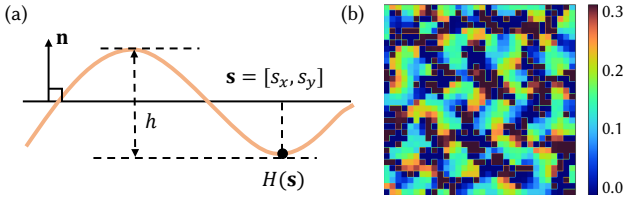

Fig. 1. Example height map (a) and visualization of one optimized height map (b). The surface is represented as a 2D field H, where $H(\mathbf{s})$ donates the height as position $\mathbf{s} = (s_x, s_y)$. We assume the normal vector of the surface is $(0, 0, 1)$, and the total height is h. The optimized height map for the Anti-mirror BRDF is shown in (b).

Table 1. Table of symbols and their descriptions.

Symbol	Description
i	Imaginary unit for complex numbers, $i^2 = -1$
λ	Wavelength of light
$\mathbf{n}$	Surface normal (equal to z-axis)
$\mathbf{s}$	2D point
$H(\mathbf{s})$	Height of surface on $\mathbf{s}$
$R(s)$	Surface modulation function
θ	Subtended angle of the light source
σ_c	Size of coherence kernel
σ_b	Size of gaussian blur kernel
σ_n	Size of gaussian noise kernel
S_c	Coherence area
A_c	Area of coherence area
$\boldsymbol{\mu}$	Center of the Gabor kernel
σ_g	Size of the Gabor kernel
$\mathbf{a}$	Frequency vector of the Gabor kernel
$\boldsymbol{\omega}_i$	Direction from which light arrives (3D unit vector)
$\boldsymbol{\omega}_o$	Direction of reflected light (3D unit vector)
$\boldsymbol{\psi}$	$\boldsymbol{\psi} = \boldsymbol{\omega}_i + \boldsymbol{\omega}_o$
$\tilde{\boldsymbol{\psi}}$	2D projection $\boldsymbol{\psi}$ (removing its z-component)
f_r	Bidirectional reflectance distribution function
$\mathcal{F}$	Fourier transform
h	Maximum height of the height map

Table 2. Detailed training and rendering parameters. Here σ_c is the size of coherence area under wavelength of 0.5μm, σ_b is the size of gaussian blur kernel applied to the height map. Unit of feature size, σ_c and σ_b is μm.

Name	Feature Size	σ_c	σ_b	Angle Range (Train)	Angle Range (Render)
Figure 11agh	2	2.6	0.65	$[-30°, 30°]$	$[-17°, 17°]$
Figure 11b	2	2.6	0.65	$[-30°, 30°]$	$[-24°, 24°]$
Figure 11c	1.5	2.6	0.5	$[-30°, 30°]$	$[-17°, 17°]$
Figure 11d	2	3.9	0.65	$[-30°, 30°]$	$[-17°, 17°]$
Figure 11ef	2	5.2	0.65	$[-30°, 30°]$	$[-17°, 17°]$
Figure 11ijkl	1.5	3.9	0.65	$[-30°, 30°]$	$[-17°, 17°]$
Figure 12a	0.8	2.6	0.13	$[-30°, 30°]$	$[-17°, 17°]$
Figure 12b	0.5	2.6	0.13	$[-30°, 30°]$	$[-17°, 17°]$
Figure 12c	1	2.6	0.13	$[-30°, 30°]$	$[-17°, 17°]$
Figure 12d	0.8	2.6	0.13	$[-30°, 30°]$	$[-24°, 24°]$
Figure 12ef	0.5	2.6	0.13	$[-30°, 30°]$	$[-24°, 24°]$
Figure 15	0.8	2.6	0.13	$[-9°, 9°]$	$[-14°, 14°]$
Figure 14a	2	3.9	0.5	$[-30°, 30°]$	$[-17°, 17°]$
Figure 14b	0.8	2.6	0.1	$[-30°, 30°]$	$[-17°, 17°]$
Figure 16	0.8	2.6	0.1	$[-14°, 14°]$	$[-14°, 14°]$
Figure 17	0.8	2.6	0.1	$[-30°, 30°]$	$[-17°, 17°]$

BRDF scale correction. The simulated BRDF is computed as a relative intensity distribution and does not inherently satisfy energy conservation or geometric foreshortening. To ensure alignment with physical measurements, we introduce a global scaling factor that adjusts the rendered BRDF before loss computation. While this scalar could be treated as a trainable parameter, we empirically find that using a fixed value improves optimization stability and yields sharper BRDF reconstructions after fabrication. Figure 8 (left) in the main paper shows how different scalar values affect the spectral distribution of reflected light. Additional comparisons are shown in Figure 5.

Fabrication-aware height normalization. To ensure manufacturability, we constrain the height map to a fixed range $[0, h]$ during optimization. We begin by initializing the surface geometry using a standard Gaussian distribution and allowing the pipeline to explore height variations freely. At each iteration, the result is normalized to the specified fabrication height bound h. Additional comparisons are shown Figure 5. Furthermore, to model the optical diffusion and blurring caused by photoresist behavior, we apply a Gaussian filter with a standard deviation of σ_b to the height map, approximating edge softening during fabrication. To model other noise during fabrication and enhance the robustness of the pipeline, we also incorporate relative Gaussian noise with a standard deviation of σ_n into the height map during training.

Figure 8 (middle) in the main paper shows that increasing the height range shifts the reflected spectrum across wavelengths, offering a tunable way to bias BRDF coloration. However, large height ranges also increase fabrication difficulty and phase instability. Figure 8 (right) in the main paper shows the effect of varying the feature size. Larger features (e.g., 1.6 μm) attenuate high-frequency spectral content and reduce peak separation.

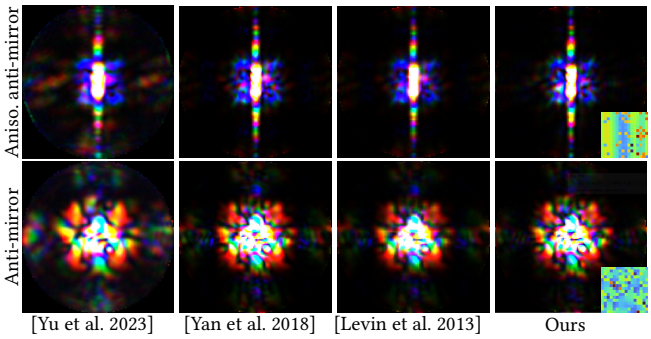

Fig. 2. Comparison on a single feature. We test our BRDF render results on the center feature of two $16\mu m \times 16\mu m$ height maps and compare them with the previous method. The height map is visualized on the right bottom inset. Both the anisotropic and isotropic results on the first and second lines, respectively, validate the accuracy of our simulation framework.

D COMPARISON

To validate our pipeline, we use Anti-mirror and Anisotropic anti-mirror reflections as examples and compare our rendered BRDF for a single feature of the height map with previous works, as shown in Figure 2. The height map size is 16 $\mu m \times 16$ μm, corresponding to the size of the coherence area under light of $\lambda = 0.5$ μm. No blur kernel or gamma correction is applied in this experiment, maintaining consistency with the previous methods.

We compare the average BRDF for the entire height map with previous works in Figure 3. For BEMsim3D [Yu et al. 2023], the Gaussian beam incident fields naturally align with the Gaussian coherence kernel. We decompose the incident illumination into 32×32 Gaussian beams, each corresponding to a distinct coherence area, and compare the average reflectance intensity. For other methods, we uniformly sample 128×128 positions on the height map, render the BRDF at each position, and visualize the average BRDF. The sampling rate of 32×32 is insufficient compared to the 128×128 sampling used in the other three methods, so the rendering result of BEMsim3D appears blurred relative to the others. Similar to Figure 2, no blur kernel or gamma correction is applied.

A key distinction between our method and others lies in the coherence area: ours varies with respect to wavelength, which is more accurate, while others assume a uniform coherence area. This discrepancy can introduce subtle color shifts, particularly noticeable in white patterns. For comparison, we visualize both uniform and nonuniform coherence area rendering results of our method.

E BRDF IMAGE RENDERING

For all our rendering results, the output angular resolution is 64×64 with 8×8 spatial samples per feature, yielding rendering times between 3 and 12 hours depending on the feature size and the size of the coherence area. Higher resolution renders, such as 128×128, require 12 to 48 hours, making this approach computationally prohibitive. As an alternative, we employ trilinear interpolation to upsample the 64×64 BRDF to the target resolution, following a methodology similar to [Yu et al. 2023]. Figure 4 presents a comparison of the original low-resolution BRDF, its interpolated counterpart, and a directly rendered high-resolution BRDF.

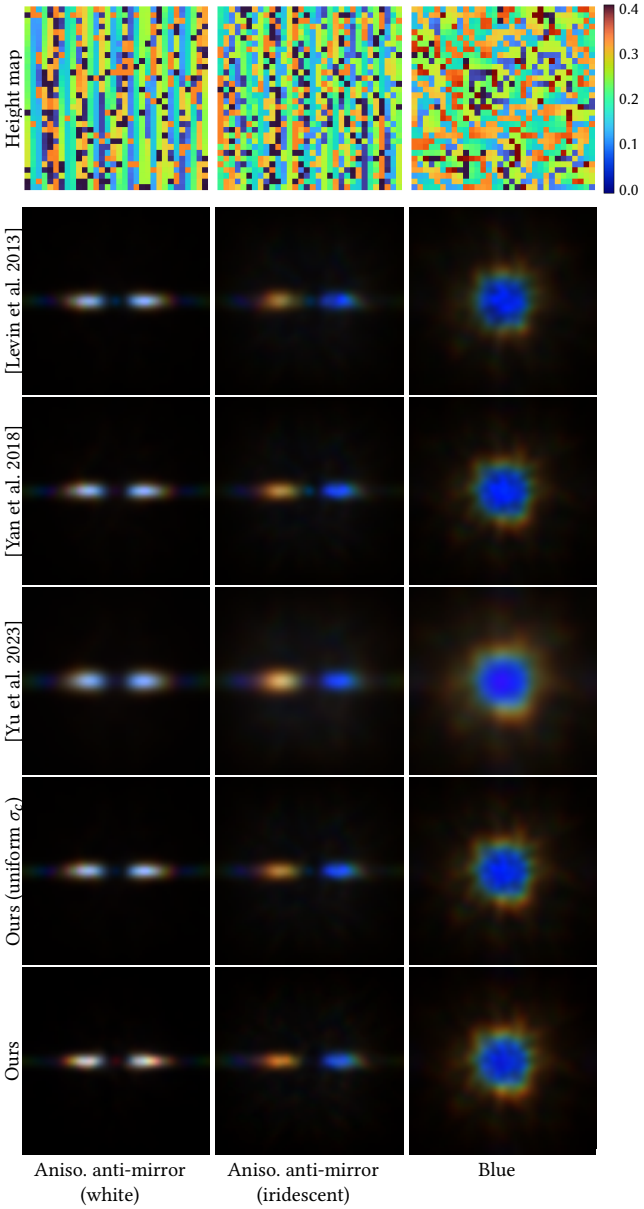

Fig. 3. Comparison on the full height map. We sample multiple spatial positions on the height map and compare the average rendered BRDF of our method against three alternative approaches. For our method, we additionally present rendering results using a uniform σ_c across all wavelengths. The results demonstrate that our optimized height map successfully achieves the target customized BRDF properties.

F ABLATION

In this section, we discuss the ablations for maximum height, intensity, and the ability of our pipeline to create isotropic and anisotropic BRDFs, as well as the angular range, patch size, and color, to justify our real-fabricated experiment parameter selection and further demonstrate the flexibility of our method.

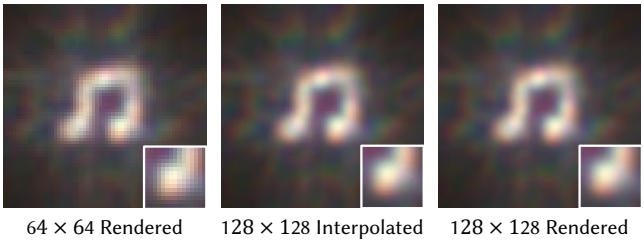

Fig. 4. Comparison of low-resolution rendered BRDF, its interpolated counterpart, and a directly rendered high-resolution BRDF. Resolution of each BRDF is listed on the second line. An inset (bottom right) shows enlarged detail.

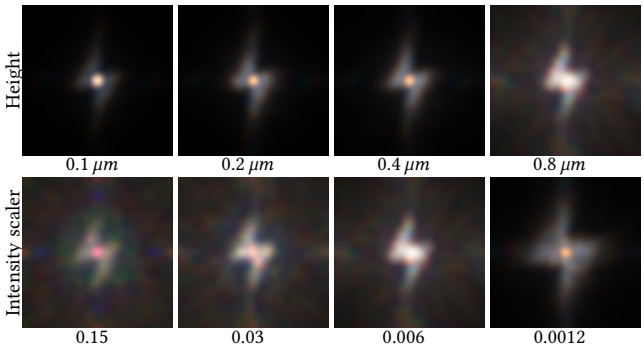

Fig. 5. Ablation of maximum height and intensity scaler on Lightning BRDF. When the maximum height is small, a bright white spot appears at the center due to strong mirror reflection (zeroth-order diffraction). As the maximum height increases, the optimization gains greater freedom, effectively suppressing the mirror reflection and improving overall BRDF quality. For intensity, with both too large or too low intensity, the quality of simulated BRDF is degraded.

Maximum height & Intensity. We tested the influence of maximum height and intensity on the Lightning BRDF, as shown in Figure 5. With a smaller maximum height, the quality of the BRDF degraded. However, our wave optics rendering pipeline does not account for shadowing and masking effects, which creates a trade-off between the quality and the impact of unaccounted-for shadowing and masking effects. For intensity ablation, with both larger and smaller intensities, the quality of the rendered BRDF decreases. For all experiments in the paper, we carefully test different intensity scalars and maximum heights and choose the best one to display and fabricate.

Isotropic and anisotropic. We test our ability to create both isotropic and anisotropic BRDFs, as shown in Figure 6. The height maps exhibit similar isotropic and anisotropic geometric structures as those in geometric optics. At the microscale, a simple anisotropic stripe pattern acts as a diffraction grating, producing rainbow-colored reflectance. To achieve the desired anisotropic BRDF, the strict periodicity must be disrupted by introducing other structural variations, which explains why our optimized height map deviates from a plain stripe pattern.

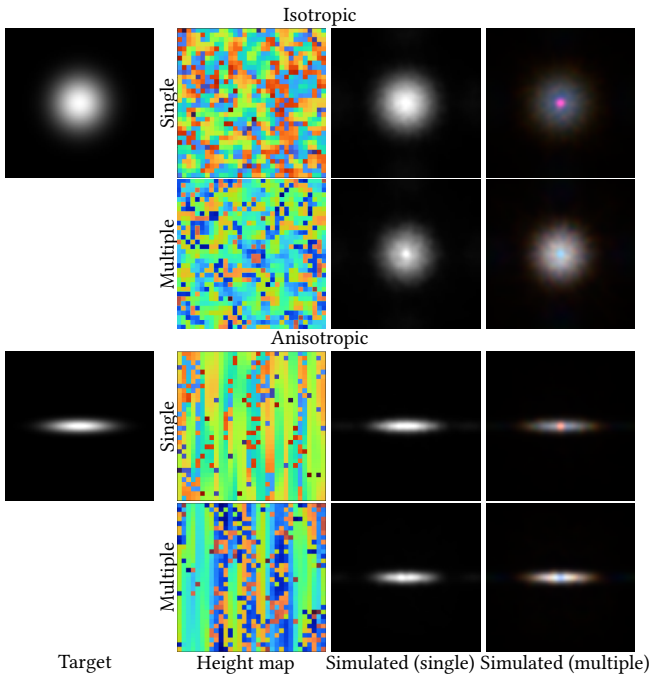

Fig. 6. Height maps and rendering results for isotropic and anisotropic BRDF. We train our pipeline to create simple isotropic and anisotropic BRDF. The height maps are trained under single wavelength ($\lambda = 0.5$, in line 1 and 3) and multiple wavelengths (in line 2 and 4). The BRDFs are rendered under single wavelength ($\lambda = 0.5$, in column 3) and multiple wavelengths (in column 4). No smooth kernel or noise is used during training to maintain consistency with [Levin et al. 2013].

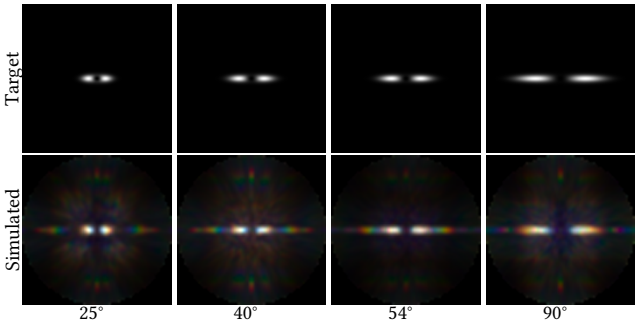

Fig. 7. Ablation of angular range. We evaluate the angular range of our setting using two anti-mirror BRDFs of 25°, 40°, 54° and 90° angle lobe respectively. Theoretically, the maximum angular range achievable with feature size of 0.8 μm is approximately 80°. In our optimization, as the BRDF angle exceeds 80°, the quality of the optimized reflectance deteriorates.

Angular range. Our ability to create BRDF with different angular ranges is shown in Figure 7. We visualize the BRDF with fixed ω_i and ω_o within an angular range of 180°. With a feature size of 0.8 μm, theoretically, our angular range can reach around 80°, according to the discussion of physical step dimensions in [Levin et al. 2013]. The optimized reflectance distribution may degrade at larger angles,

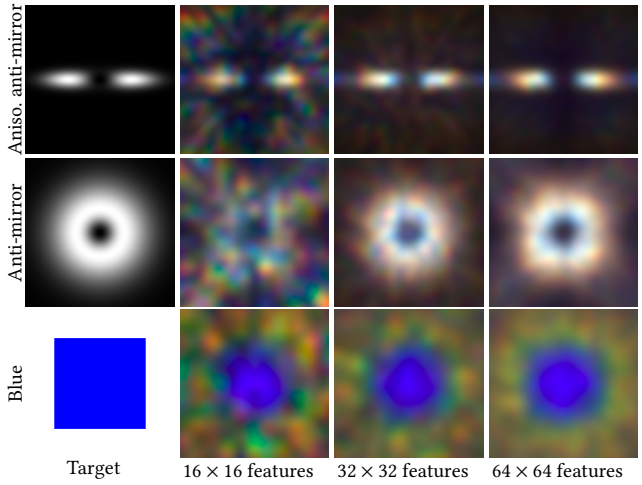

Fig. 8. Ablation study of the patch size. We evaluate the impact of patch size on three representative BRDFs: Aniso. anti-mirror, Anti-mirror, and Blue, using patch sizes of 16×16, 32×32, and 64×64. With a 16×16 features per patch, the quality of the trained BRDFs decreases. With a 64×64 features per patch, the quality improves while the training process need 3 days to converge. As a result, we select 32×32 as the optimal patch size for all experiments.

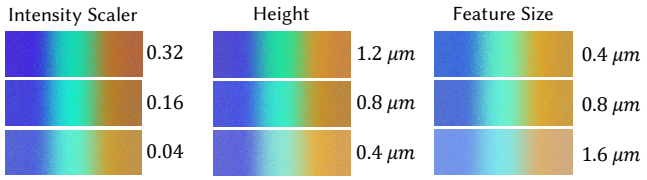

Fig. 9. RGB visualization of spectrum in Figure 8 in the main paper. We use the blue, green, and red colors of the spectrum curve in this figure to create a color bar and visualize them. The color is visualized with gamma correction ($\gamma = 2.2$). The numbers on the right side of each column are the intensity scaler, maximum height, and feature size, respectively.

showing rainbow tails and becoming unable to produce white color at wider angles.

Patch size. We use patch size of 32×32 features in all experiments. Smaller and larger patch sizes are tested in Figure 8. With a smaller patch size, the quality of the BRDF decreases because the degrees of freedom are insufficient to generate diverse BRDFs, especially for color BRDFs. For larger patch sizes, the quality improves, but the training speed becomes too slow, resulting in convergence times of 2–3 days.

Color discussion. Figure 9 illustrate the RGB color of spectrum curve shown in Figure 8 in the main paper.

Coherence Area in Rendering. Figure 10 illustrates how the size of the coherence area influences the rendering result.

REFERENCES

Anat Levin, Daniel Glasner, Ying Xiong, Frédo Durand, William Freeman, Wojciech Matusik, and Todd Zickler. 2013. Fabricating BRDFs at high spatial resolution using

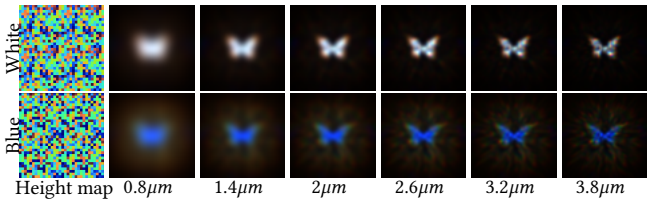

Fig. 10. Comparison of coherence area size in rendering. We investigate the impact of coherence area size on the white and blue Butterfly BRDFs. The figures show rendering results on the left height map with different coherence area sizes. The numbers in the third row represent the coherence area size under light with a wavelength of $\lambda = 0.5\,\mu m$. During training, a coherence area size of 2.6 μm under the wavelength of $\lambda = 0.5\,\mu m$ was used, corresponding to the subtended angle of 1.8°.

wave optics. *ACM Trans. Graph.* 32, 4, Article 144 (jul 2013), 14 pages. https://doi.org/10.1145/2461912.2461981

Ling-Qi Yan, Milos Hasan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering specular microgeometry with wave optics. *ACM Trans. Graph.* 37, 4, Article 75 (jul 2018), 10 pages. https://doi.org/10.1145/3197517.3201351

Yunchen Yu, Mengqi Xia, Bruce Walter, Eric Michielssen, and Steve Marschner. 2023. A Full-Wave Reference Simulator for Computing Surface Reflectance. *ACM Trans. Graph.* 42, 4, Article 109 (jul 2023), 17 pages. https://doi.org/10.1145/3592414