Rendering Specular Microgeometry with Wave Optics

LING-QI YAN, University of California, Berkeley
MILOŠ HAŠAN, Autodesk
BRUCE WALTER, Cornell University
STEVE MARSCHNER, Cornell University
RAVI RAMAMOORTHI, University of California, San Diego

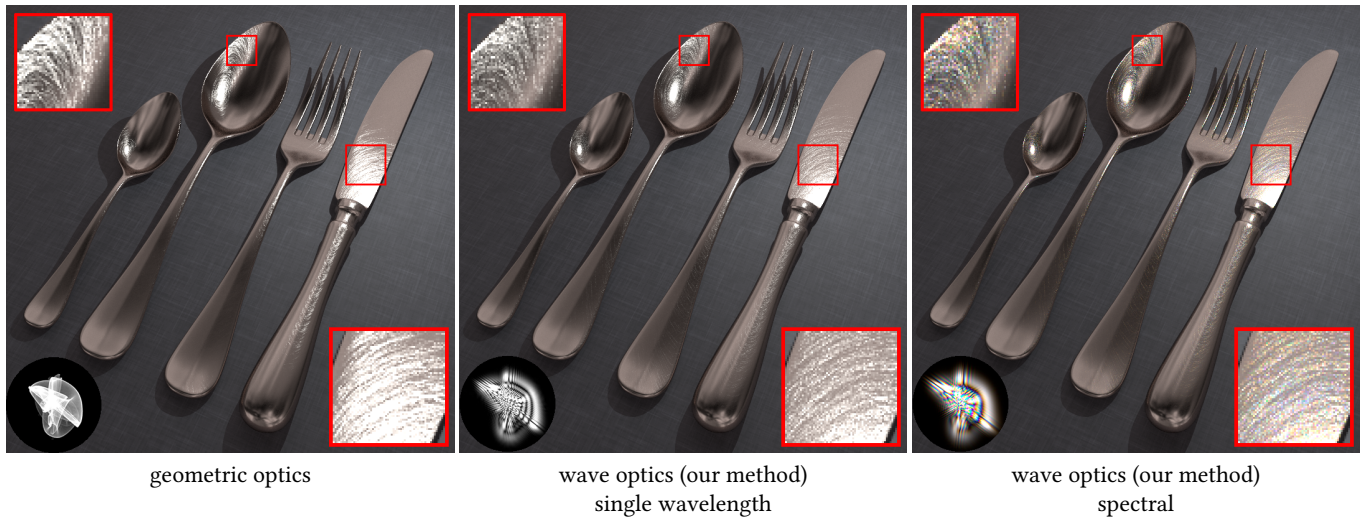| geometric optics | wave optics (our method) single wavelength | wave optics (our method) spectral |

Fig. 1. We present the first practical method for rendering specular reflection from arbitrary high-resolution microstructure (represented as discretized heightfields) using wave optics. Left: Rendering with previous work [Yan et al. 2016], based on the rules of geometric optics. Middle: Using wave optics, even with a single fixed wavelength, our method generates a more natural appearance as compared to geometric optics. Right: A spectral rendering additionally shows subtle but important color glint effects. Insets show enlarged regions and representative BRDFs generated using each method. We encourage readers to zoom in to better see color and detail, and to view the full resolution supplementary images to see the subtle details in all of the figures.

Simulation of light reflection from specular surfaces is a core problem of computer graphics. Existing solutions either make the approximation of providing only a large-area average solution in terms of a fixed BRDF (ignoring spatial detail), or are specialized for specific microgeometry (e.g. 1D scratches), or are based only on geometric optics (which is an approximation to more accurate wave optics). We design the first rendering algorithm based on a wave optics model that is also able to compute spatially-varying specular highlights with high-resolution detail on general surface microgeometry. We compute a wave optics reflection integral over the coherence area; our solution is based on approximating the phase-delay grating representation

of a micron-resolution surface heightfield using Gabor kernels. We found that the appearance difference between the geometric and wave solution is more dramatic when spatial detail is taken into account. The visualizations of the corresponding BRDF lobes differ significantly. Moreover, the wave optics solution varies as a function of wavelength, predicting noticeable color effects in the highlights. Our results show both single-wavelength and spectral solution to reflection from common everyday objects, such as brushed, scratched and bumpy metals.

CCS Concepts: • **Computing methodologies → Rendering**;

Additional Key Words and Phrases: specular surface rendering, glints, material appearance, Harvey-Shack, wave optics

Authors' addresses: Ling-Qi Yan, University of California, Berkeley, lingqi@berkeley.edu; Miloš Hašan, Autodesk, milos.hasan@gmail.com; Bruce Walter, Cornell University, bruce.walter@cornell.edu; Steve Marschner, Cornell University, srm@cs.cornell.edu; Ravi Ramamoorthi, University of California, San Diego, ravir@cs.ucsd.edu.

1 INTRODUCTION

Simulation of material appearance is a core problem of computer graphics, and specular highlight appearance is among the most common effects. One of the most fundamental questions is: given a

Fig. 2. Left: Rendering of a laptop with a point light and environment lighting using our method. Top right: Close up rendering of the corner of the laptop with the same lighting condition. Bottom right: A photograph of a MacBook (around $20\ cm \times 4\ cm$ region) lit by a small LED light in a dark room. Our method is able to produce appearance that is perceptually similar to the photograph, showing colored glints from the underlying noisy microstructure of the aluminum laptop body.

heightfield specifying the microgeometry of a surface, how do we compute the surface reflection according to laws of physics?

Most existing solutions either make the approximation of providing only a large-area average solution in terms of a mean BRDF (ignoring spatial detail), or are based on geometric optics (an approximation to wave optics), or both. For example, the standard microfacet BRDF models commonly used in rendering [Cook and Torrance 1982] assume the surface consists of infinitely small unresolved microfacets, which act as perfect mirrors following the rules of geometric optics. Thus, these models make both of the approximations mentioned (no spatial variation and no wave optics).

Some previous research focused on lifting one or the other of these two limitations. Several wave optics reflection models have been proposed in various areas of physics, including Kirchoff and Harvey-Shack reflection theories [Krywonos 2006], and versions of these have been used in computer graphics, but always assuming large-area averages.

Large-area averages are successful for distant views and smooth illumination, but the high-frequency spatially-varying structure seen in real specular highlights cannot be replicated without modeling discrete, finite microgeometry features. Recent work on rendering glints [Yan et al. 2014, 2016] has moved beyond large-area averages, by presenting solutions for surfaces defined by explicit high-resolution heightfields (or normal maps). However, these methods continue using the rules of geometric optics at scales approaching micron resolution, where they are known to become less accurate. Moreover, recent work has introduced solutions based on wave optics that are capable of handling high-resolution surfaces of a specific kind, defined by a flat surface with randomly oriented 1D scratch curves [Werner et al. 2017].

This leads to the question: can we design a BRDF model based on the more accurate wave optics, but also able to compute spatially-varying solutions with high-resolution detail? To our knowledge,

our paper is the first to consider this question in full generality, modeling the surfaces as arbitrary discretized heightfields.

We present an algorithm that can evaluate a spatially varying BRDF for a given position and incoming/outgoing directions, by computing a wave optics reflection integral over the coherence areas around the position of interest. This requires more computation than in the geometric optics solutions, which makes our method slower, but not prohibitively so; see Table 1. Our solution is based on approximating the micron-resolution surface wave effects using Gabor kernels (products of Gaussians with complex exponentials). We use a reciprocal modification of the Harvey-Shack theory in our results, but our approach also applies to other wave optics models.

We found that the difference between the geometric and wave solution is more dramatic when spatial detail is taken into account. The visualizations of the corresponding BRDF lobes differ dramatically, with the sharp folds typical of geometric normal distribution functions (NDFs) replaced by very different directional patterns more akin to laser speckle (Figure 10). The rendered highlights change appearance, typically with more realistic-looking sharper peaks and longer tails. Moreover, the wave optics solution varies as a function of wavelength, predicting noticeable color effects in the highlights (Figure 1). Our results show both single-wavelength and spectral solutions to reflection from common everyday objects, such as brushed, scratched and bumpy metals; see the result figures and supplementary video.

2 RELATED WORK

We organize the related work into four areas, based on whether they use geometric or wave optics, and whether they target large-area average BRDFs or spatially-varying fine-scale details and "glints".

Large-area, geometric optics. Microfacet BRDFs have become a standard tool in rendering [Burley 2012; Butler et al. 2015; Cook and Torrance 1982; Walter et al. 2007; Westin et al. 2004]. The distribution of the normals of the microfacets is modeled using a smooth

normal distribution function (NDF), and the BRDF additionally contains Fresnel and shadowing-masking terms. This approach models only directional, not spatial variation; the latter is normally added by texturing and bump/normal mapping, which has limitations at high resolutions and under high-frequency lighting. Furthermore, geometric optics is theoretically accurate only if surface features are locally flat at the scale of microns; many real surfaces violate this assumption, but empirically the microfacet approach often still provides good results.

Large area, wave optics. Rough surface reflection models based on wave optics have been heavily studied in physics. Common approximations include Beckmann-Kirchhoff theory [Beckmann and Spizzichino 1968] and variations of Harvey-Shack theory[Harvey 1979]; a good overview is the thesis of Krywonos [2006]. In graphics, wave-based reflection models have been developed for surfaces with stationary statistics, either random [He et al. 1991] or periodic [Stam 1999], usually characterized by their power spectral density. A variety of methods have been proposed to measure such statistics for specific types of real surfaces, especially periodic ones [Dhillon et al. 2014; Lanari et al. 2017; Toisoul and Ghosh 2017]. Dong et al. [2015] acquired the surface microgeometry of real metallic surfaces using a profilometer, and applied Kirchhoff theory to successfully predict their large scale BRDFs. A combined microfacet-diffraction model was recently proposed by Holzschuch and Pacanowski [2017], demonstrating better fits to measured BRDF data for some materials than microfacet models alone. Levin et al. [2013] designed special multi-planar surfaces that can be lithographically fabricated to match a target BRDF using wave optics, essentially inverting the rendering process.

Wave optics has also been used to predict appearance from thin-film or layered materials (e.g., [Belcour and Barla 2017]), but we will only consider single-layer opaque surfaces here. Several methods support longer-range multi-surface interference effects (e.g., [Cuypers et al. 2012]), but that is beyond the scope of this paper.

Spatially-varying, geometric optics. Yan et al. [2014; 2016] presented algorithms for rendering glinty surfaces defined by explicit high-resolution heightfields (or normal maps), under geometric optics. These approaches are successful at simulating very high-resolution, spatially varying glinty behavior. The key idea is to extend the NDF from microfacet theory to a patch-based $\mathcal{P}$-NDF, essentially replacing the large-area average for the whole surface by a unique solution per given small patch of the surface. In the first paper, they introduce an algorithm that evaluates $\mathcal{P}$-NDFs by turning the problem into integration over the patch $\mathcal{P}$, finely discretized into triangles. The second paper considers the same problem, but proposes a higher-performance algorithm. Instead of discretizing the normal map into triangles, they fit small Gaussian elements to texels of the normal map. Our approach is related, but instead of real Gaussians, we use complex Gabor kernels, which are better matched to approximating the underlying complex integrals. Jakob et al. [2014] also simulated glinty surfaces but used a statistical distribution of tiny mirror-like flakes rather than an explicit surface. In addition to these general approaches, other methods [Bosch et al. 2004; Mérillou et al. 2001; Raymond et al. 2016] focus specifically on scratched surfaces, also under geometric optics.

Spatially-varying, wave optics. The only previous work we are aware of in this area is the recent paper by Werner et al. [2017] (with a real-time extension by Velinov et al. [2018]), rendering surfaces with collections of randomly oriented scratches using a Harvey-Shack-based wave optics model. This work represents the surface as a collection of one-dimensional scratches over a smooth BRDF. Under this assumption, they are able to compute the reflection efficiently and analytically. In contrast, our method can render arbitrary heightfields (e.g. Figure 2 and 13), including but not limited to ones containing scratches. Additionally, our scratched heightfields can contain more variety and imperfections, resulting in glinty highlights that only roughly align in lines, compared to the smooth line highlights of Werner et al (see Figure 1, esp. insets).

3 WAVE BRDF THEORY

In wave optics, light is described by fields that satisfy appropriate boundary conditions and governing differential equations (e.g., wave or Helmholtz equations). We will consider each wavelength (denoted λ) separately and use complex-valued fields to encode both magnitude and phase. The local light energy is related to the squared magnitude of the field at that point. Scalar diffraction models, such as Harvey-Shack [Krywonos 2006] or Kirchhoff [Ogilvy 1991], can be used to estimate the reflected field from a rough surface. Unlike in geometric optics, the contributions from different parts of the surface can sum non-linearly due to interference effects, to create the characteristic diffraction effects of wave optics.

Let us assume we have a surface heightfield $H(s)$ (as in Figure 3) such that for a given 2D point $s = [s_x, s_y]$, the corresponding 3D point on the rough surface is $[s_x, s_y, H(s)]$. In our approach, the
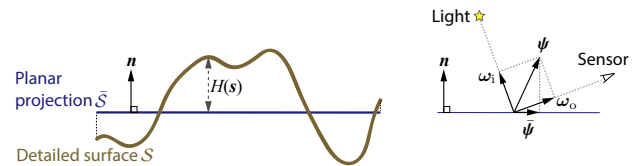

Fig. 3. Heightfield surface and BRDF directions example.

|	
i	Imaginary unit for complex numbers, $i^2 = -1$
λ	Wavelength of light
$\mathbf{n}$	Average surface normal (equal to z-axis)
$\mathbf{s}$	2D point (on the XY plane)
$H(\mathbf{s})$	Height of surface above $\mathbf{s}$
$H'(\mathbf{s})$	Gradient of height function
$\bar{S}$	Domain of height function (region on XY plane)
$A_{\bar{S}}$	Area of $\bar{S}$
ω_i	Direction from which light arrives (3D unit vector)
ω_o	Direction of reflected light (3D unit vector)
ψ	$\psi = \omega_i + \omega_o$
$\bar{\psi}$	2D projection ψ (removing its z-component)
f_r	Bidirectional reflectance distrib. function (BRDF)
F	Surface reflectance (e.g., from Fresnel equations)
ξ_1, ξ_2, ξ_3	See Figure 6

Fig. 4. List of symbols.
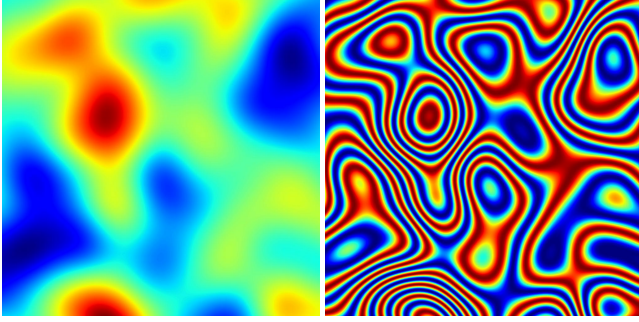
Fig. 5. Left: A discretized surface heightfield at 1 micron resolution, showing an area of about 64×64 microns. For visualization purposes, we complete the heightfield into a continuous function $H(s)$ by bicubic interpolation. Right: The real component of the reflection function $R(s)$ of this surface patch, specifying the spatially varying phase shift. The imaginary component looks similar.

Diffraction	Equation Components		
BRDF Model	ξ_1	ξ_2	ξ_3
1. OHS	$\dfrac{\|\omega_o \cdot n\| F}{\lambda^2 \|\omega_i \cdot n\|}$	1	2
2. GHS	$\dfrac{\|\omega_o \cdot n\| F}{\lambda^2 \|\omega_i \cdot n\|}$	1	$\psi \cdot n$
3. R-OHS	$\dfrac{\|\psi \cdot n\|^2 F}{4\lambda^2 \|\omega_i \cdot n\|\|\omega_o \cdot n\|}$	1	2
4. R-GHS	$\dfrac{\|\psi \cdot n\|^2 F}{4\lambda^2 \|\omega_i \cdot n\|\|\omega_o \cdot n\|}$	1	$\psi \cdot n$
5. Kirchhoff	$\dfrac{\|\psi \cdot n\|^2 F}{4\lambda^2 \|\omega_i \cdot n\|\|\omega_o \cdot n\|}$	$1 - \dfrac{\bar{\psi} \cdot H'(s)}{\psi \cdot n}$	$\psi \cdot n$

Fig. 6. BRDF integrals for five scalar diffraction models (see equations (1) and (2)). The first two are based on the Original-Harvey-Shack (OHS) and the Generalized-Harvey-Shack (GHS) models. The next two are reciprocal versions of these models we created by substituting Kirchhoff propagation instead of Fourier: Reciprocal OHS (R-OHS) and Reciprocal GHS (R-GHS). The fifth is a fully Kirchhoff-based BRDF model. Detailed derivations of these models can be found in the supplemental material.

heightfield is typically discretized at the resolution of 1 μm per texel. Figure 5 (left) illustrates a small example heightfield. Our goal is to estimate the surface's Bidirectional Reflectance Distribution Function (BRDF) $f_r(\omega_i, \omega_o)$, which is defined as the ratio between the reflected radiance in direction ω_o and the incident irradiance from direction ω_i. Light reflecting from different parts of the surface will travel different distances depending on the local surface height. This causes phase shifts in reflected waves which then interfere with each other to determine the BRDF.

These phase shifts can be approximated using a planar surface that reflects light with a spatially-varying phase shift, specified by its reflection function:

$$R(s) = \xi_2 \, e^{-i\frac{2\pi}{\lambda} \xi_3 H(s)}. \qquad (1)$$

Figure 5 (right) shows a visualization of the real component of this function. The values of ξ_2 and ξ_3 depend on which diffraction model is chosen (see Figure 6 for examples). We represent the directions ω_i and ω_o as 3D unit vectors. Let $\psi = \omega_i + \omega_o$ and $\bar{\psi}$ be its 2D projection (by discarding its z-component). The BRDF of this planar

proxy can be computed using a surface integral of the form:

$$f_r(\omega_i, \omega_o) = \frac{\xi_1}{A_{\bar{S}}} \left| \int_{\bar{S}} R(s) \, e^{-i\frac{2\pi}{\lambda}(\bar{\psi} \cdot s)} \, ds \right|^2 \qquad (2)$$

where $\bar{S}$ is the domain of the heightfield (i.e. the projection of the rough surface onto the XY plane), $A_{\bar{S}}$ is its area, and ξ_1 depends on the chosen diffraction model (see Figure 6) .

The parameters for five different diffraction models are listed in Figure 6 and detailed derivations of these models are provided in the supplemental material. These models are closely related and often produce similar results, especially for low-slope surfaces and paraxial directions. One advantage of our approach is that it can be used to compute any of these models. The first four are derived from the Harvey-Shack family of diffraction models [Harvey and Pfisterer 2016]. The first uses the phase shift approximation from Original-Harvey-Shack (OHS) and the second uses the more accurate phase shift from Generalized-Harvey-Shack (GHS). These produce non-reciprocal BRDFs (i.e. $f_r(\omega_i, \omega_o) \neq f_r(\omega_o, \omega_i)$). Reciprocal BRDF estimates are often preferred in rendering, since real world BRDFs are reciprocal, and reciprocity also simplifies some light transport algorithms. Therefore we created reciprocal versions (R-OHS and R-GHS) by keeping the same planar proxy and phase shift approximations, but using the Kirchhoff propagation integral instead of the usual Fourier-based propagation. The fifth model is equivalent to the Kirchhoff-based BRDF from Dong et al. [2015] and is also reciprocal. In our results we use the third method (R-OHS), with its convenient simplicity and reciprocity, except where otherwise noted. In the supplementary material, we include some comparisons using the other models, showing that (for computer graphics purposes) the results are often quite similar.

3.1 Coherence area

The spatial size over which the incident light's phase remains correlated (i.e. coherent) is known as its coherence area. Equation 2 was derived using incident light with an infinite coherence area, but realistic sources have finite ones (typically inversely related to their solid angle [Mandel and Wolf 1995]). For example, sunlight [Mashaal et al. 2012] has a measured coherence area diameter of roughly one hundred wavelengths, or ~ 50 microns.

Coherent contributions must be summed using their complex field values, while incoherent ones are accumulated by summing their energy (or equivalently, averaging their BRDF values). This is commonly simulated (e.g., [Dong et al. 2015; Levin et al. 2013; Werner et al. 2017]) by spatially limiting the surface integrals using a coherence kernel $w(s)$, and then averaging multiple such BRDF evaluations over the region of interest (e.g., the pixel footprint). The principal effect of limiting the coherence area is a small angular blurring of the BRDF. The BRDF estimate for one coherence area becomes:

$$f_r(\omega_i, \omega_o) = \frac{\xi_1}{A_c} \left| \int_{\bar{S}_c} R^\star(s) \, e^{-i\frac{2\pi}{\lambda}(\bar{\psi} \cdot s)} \, ds \right|^2 \qquad (3)$$

$$R^\star(s) = w(s - x_c) R(s) \qquad (4)$$

where $\bar{S}_c$ is the portion of $\bar{S}$ within the support of the coherence kernel centered at x_c, the corresponding normalization factor is

$A_c = \int |w(s)|^2\, ds$, and $R^\star$ is the product of $R(s)$ and the coherence kernel. This has the advantages of limiting our integrals to small surface regions and effectively prefiltering the BRDF to remove high frequency angular features that we expect are too small to be resolved. Generally we do not need to exactly match the real coherence area. Overestimating it leads to high angular frequency aliasing that can be resolved by using more light samples, while underestimating it causes some angular over-blurring of the BRDF.

During rendering, rather than trying to estimate each source's coherence area, we use a fixed size, which should be at least as large as for any expected light source. For w we use a Gaussian with standard deviation of 10 microns (similar to [Werner et al. 2017]).

3.2 Fourier Interpretation

Let us denote the Fourier transform of a 2D function $f(s)$ as:

$$\mathcal{F}[f](\boldsymbol{v}) \equiv \widetilde{f}(\boldsymbol{v}) \equiv \int_{\mathbb{R}^2} f(s)\, e^{-i2\pi(s \cdot \boldsymbol{v})}\, ds \qquad (5)$$

where $\boldsymbol{v}$ is a 2D frequency vector. Equation (3) can be rewritten as:

$$f_r(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o) = \frac{\xi_1}{A_c} \left| \widetilde{R^\star}\left(\frac{\bar{\boldsymbol{\psi}}}{\lambda}\right) \right|^2 \qquad (6)$$

Thus the BRDF can be computed using the Fourier transform of $R^\star(s)$ evaluated at $\bar{\boldsymbol{\psi}}/\lambda$. One approach could be to compute and store the full Fourier transform, either analytically or numerically via the Fast Fourier Transform (FFT) algorithm. However we use tabulated heightfields which have no simple analytic Fourier transform, and precomputing FFTs for each surface position would require far too much storage. Computing full FFTs at render time would also be very inefficient as we typically only need one, or at most a few, values for each BRDF evaluation. Also $R^\star(s)$ typically contains very high frequencies, much higher than those in the original heightfield, so using an FFT would require an extremely fine discretization step of 0.1 microns or less. We could evaluate just $\widetilde{R^\star}(\bar{\boldsymbol{\psi}}/\lambda)$ as needed using numerical quadrature, but this would similarly be expensive and require high sampling rates. However, we do use the FFT approach as a ground truth for checking the correctness of our approach in Section 6.

4 EFFICIENT BRDF EVALUATION

In this section, we discuss how to evaluate the BRDF integrals for our wave optics diffraction models. Our high-level idea for efficiently approximating the integral in equation (3) is to approximate the phase-delay reflection function $R^\star(s)$ by a weighted combination of *Gabor kernels*, which are products of a 2D Gaussian with a complex exponential (plane wave). These kernels are well suited to representing the high-frequency features found in typical $R^\star(s)$, while also having other desirable properties.

Notably, Gabor kernels have an analytical Fourier transform that is itself a Gabor kernel. This means that the kernels and their transforms both have spatially localized support (ignoring negligibly small values of the Gaussian component), which is a key property for designing an efficient pruning algorithm.

4.1 Gabor kernels

Let us define a Gabor kernel as the product of a 2D Gaussian and a complex exponential:

$$g(s; \boldsymbol{\mu}, \sigma, \boldsymbol{a}) = G_{2D}(s; \boldsymbol{\mu}, \sigma)\, e^{-i2\pi(\boldsymbol{a} \cdot s)} \qquad (7)$$

where $G_{2D}(s; \boldsymbol{\mu}, \sigma) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{\|s-\boldsymbol{\mu}\|^2}{2\sigma^2}\right)$ is a normalized 2D isotropic Gaussian. Here $\boldsymbol{\mu}$ is the center, σ the width and $\boldsymbol{a}$ the plane wave parameter. This definition is similar to others used in the literature; the normalization constant of the Gaussian and the additional 2π factor in the complex exponential are chosen to simplify the following derivations.

The Fourier transform of a Gabor kernel can be written as another Gabor kernel:

$$\mathcal{F}[g(s; \boldsymbol{\mu}, \sigma, \boldsymbol{a})](\boldsymbol{v}) = e^{-i2\pi(\boldsymbol{\mu} \cdot (\boldsymbol{v}+\boldsymbol{a}))}\, e^{-2\pi^2\sigma^2\|\boldsymbol{v}+\boldsymbol{a}\|^2}$$

$$= \frac{1}{2\pi\sigma^2} e^{-i2\pi(\boldsymbol{\mu} \cdot \boldsymbol{a})}\, g\left(\boldsymbol{v}; -\boldsymbol{a}, \frac{1}{2\pi\sigma}, \boldsymbol{\mu}\right) \qquad (8)$$

4.2 Approximating R with Gabor kernels

We first subdivide the heightfield domain $\bar{S}$ into a grid of cells. We use a uniform grid so all the cells are identically-sized squares, matching the original heightfield texels, but an adaptive subdivision could also be used. Then we select a set of cells, with centers $\boldsymbol{m}_k$, that covers the support of the current coherence kernel. Since the cells are much smaller than the coherence area, we approximate the coherence kernel as being constant over a cell with value $w_k = w(\boldsymbol{m}_k - \boldsymbol{x}_c)$. Then we place a Gabor kernel centered on each grid cell designed to approximate $R(s)$ in its neighborhood. Together this gives us an approximation for $R^\star(s)$ of the form:

$$R^\star(s) \approx \sum_k w_k R_k(s) = \sum_k w_k C_k\, g(s; \boldsymbol{m}_k, \sigma_k, \boldsymbol{a}_k) \qquad (9)$$

where C_k is a complex constant, incorporating an appropriate scaling coefficient and phase shift.

We choose $\sigma_k = l_k/2$, where l_k is the side length of the cell. This choice was found to give good results experimentally. A sum of Gaussians is not an exact partition of unity; this leads to slight approximation error that manifests itself as spurious periodic copies of the main transform image in the Fourier domain. However, for the choice $\sigma_k = l_k/2$, the copies are weak enough that we do not observe them in practice.

Next, we approximate the heightfield $H(s)$ in each cell by its first order expansion around $\boldsymbol{m}_k$:

$$H(s) \approx H(\boldsymbol{m}_k) + H'(\boldsymbol{m}_k) \cdot (s - \boldsymbol{m}_k) \qquad (10)$$

$$= H'(\boldsymbol{m}_k) \cdot s + (H(\boldsymbol{m}_k) - H'(\boldsymbol{m}_k) \cdot \boldsymbol{m}_k) \qquad (11)$$

where $H'(\boldsymbol{m}_k)$ is the gradient of the heightfield at $\boldsymbol{m}_k$. Substituting this approximation into the definition of $R(s)$, we can approximate a single grid cell's contribution as:

$$R_k(s) = B_{2D}(s; \boldsymbol{m}_k, l_k)\, \xi_2\, e^{-\frac{i2\pi\xi_3}{\lambda} H(s)} \qquad (12)$$

$$\approx l_k^2 G_{2D}(s; \boldsymbol{\mu}_k, \sigma_k)\, \xi_2\, e^{-\frac{i2\pi\xi_3}{\lambda}\left(\alpha_k + H'(\boldsymbol{m}_k) \cdot s\right)} \qquad (13)$$

where $\alpha_k = H(\boldsymbol{m}_k) - H'(\boldsymbol{m}_k) \cdot \boldsymbol{m}_k$. B_{2D} is a binary box function indicating the domain of the grid cell, which integrates to the cell's

ground truth   2x2 kernels / texel   1 kernel / texel

1 kernel / 2x2 texels   1 kernel / 4x4 texels   1 kernel / 8x8 texels
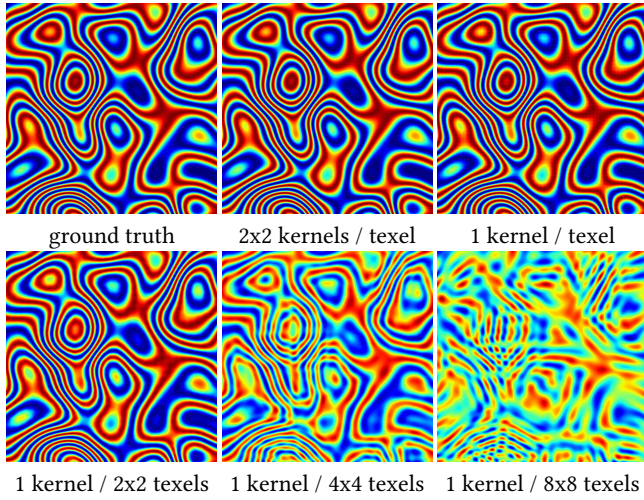
Fig. 7. A color-mapped (range [-1,1]) visualization of the real component of $R(s)$ for the isotropic noise heightfield (the imaginary component looks similar). The area depicted is about 64×64 texels, using the resolution of 1 micron / texel. Note the common structure seen in these functions: high-frequency ripples aligned with slopes of the original heightfield, with frequency increasing proportional to slope. This structure is ideal for approximation by Gabor kernels. These images show the approximation quality for various kernel sampling densities. All our results use 1 kernel per texel (i.e. per micron). Note that as the number of kernels decreases, the approximation degrades, as expected.

area l_k^2. Then we replace the box function with a 2D Gaussian of the equal area.

Comparing Eqn. 13 with the definition in Eqn. 9, we have

$$C_k = l_k^2 \xi_2 e^{-\frac{i2\pi\xi_3}{\lambda}\left(H(m_k) - H'(m_k) \cdot m_k\right)} \tag{14}$$

$$a_k = \frac{\xi_3 H'(m_k)}{\lambda} \tag{15}$$

which completes our Gabor approximation for $R^\star(s)$.

Figure 7 shows $R(s)$ for an example heightfield compared to its approximation as a sum of Gabor kernels. At the density of 1 kernel per texel, though the sampling pattern is just visible, we obtain a sufficiently good approximation that reproduces the relevant details.

4.3 BRDF approximation

Finally we use our Gabor kernel approximation to evaluate the BRDF. Starting from Equation 6 we have:

$$f_r(\omega_i, \omega_o) = \frac{\xi_1}{A_c} \left| \widetilde{R^\star}\left(\frac{\bar{\psi}}{\lambda}\right) \right|^2 \tag{16}$$

$$\approx \frac{\xi_1}{A_c} \left| \sum_k w_k C_k \mathcal{F}\left[g(s; m_k, \sigma_k, a_k)\right]\left(\frac{\bar{\psi}}{\lambda}\right) \right|^2 \tag{17}$$

where we can use the above definitions of the quantities C_k and a_k, and equation (8) to evaluate the Fourier transform of the Gabor kernel. Thus we can evaluate the sum in a straightforward manner by iterating over all the cells within the coherence area. We also apply pruning to non-contributing cells, as detailed in Section 5.

Appendix A briefly sketches an alternate derivation of our method where the surface is approximated by overlapping planar elements.

5 IMPLEMENTATION

In this section, we provide key implementation details of our Gabor kernel solution.

Heightfields and Gabor kernels. We use pre-defined high resolution ($8K \times 8K$) heightfields as texture maps to specify the microgeometry, where each texel represent a fixed size of 1 square micron in the real world. The heightfields are tiled repeatedly to achieve a high resolution over a surface. The texels in a heightfield form a uniform grid naturally, so, we convert each texel into a Gabor Kernel, as specified in Sec. 4.

For simplicity, we assume no distortion from the texture map, i.e. the texture coordinates are defined to be area preserving and orthogonal in terms of u and v directions in the world coordinate system.

Acceleration by pruning. To accelerate computation, the key is to quickly decide whether a Gabor kernel contributes to the desired outgoing direction ω_o. Regardless of the cancellation from the complex numbers, each Gabor kernel is bounded by a Gaussian $G_{2D}(s; m_k, \sigma_k)$ positionally and by $G_{2D}(v; -a_k, \frac{1}{2\pi\sigma_k})$ directionally. Although in theory Gaussians have infinite support, in practice we limit them to within ±3 standard deviations, and clamp them to zero outside this region so they have only localized support.

We pre-generate a mipmap-style hierarchy for each heightfield, where each node contains both positional and directional bounding boxes of its 4 child nodes. For each BRDF query, we perform a top-down traversal of this hierarchy, discarding nodes that are not within the coherence region $\bar{S}$ using their positional bounding boxes. At the same time, we use the directional bounding boxes to prune the nodes that will not contribute to the query direction $\bar{\psi}/\lambda$.

Figure 9 shows the number of evaluations towards different directions to generate an example BRDF image. In general, each Gabor kernel contributes to a much larger range directionally in wave optics than the elements in geometric optics [Yan et al. 2016]. This explains the soft appearance in these images, as well as slower performance of wave optics. However, our hierarchical pruning is still efficient. In practice, we have a more than 50× speedup as compared to the un-accelerated implementation.

Importance sampling. With the Gabor kernels defined to represent a heightfield, it is straightforward to perform BRDF importance sampling to get the outgoing ray for global illumination. First, we randomly pick a Gabor kernel within the coherence region according to its weighting function $w(s)$. Then, we immediately know that the chosen Gabor kernel contributes to a Gaussian directionally, as analyzed in the acceleration part. By sampling this Gaussian, we have the sampled query direction $\bar{\psi}/\lambda$ and thus the corresponding outgoing direction ω_o.

The sampling weight can be calculated as the BRDF evaluation with sampled ω_o, divided by the sampling pdf. However, in practice, we found that wave optics effects are essentially not observable in indirect lighting. So, we assume that our sampling weight is always 1, i.e. discarding the complex cancellations and assuming
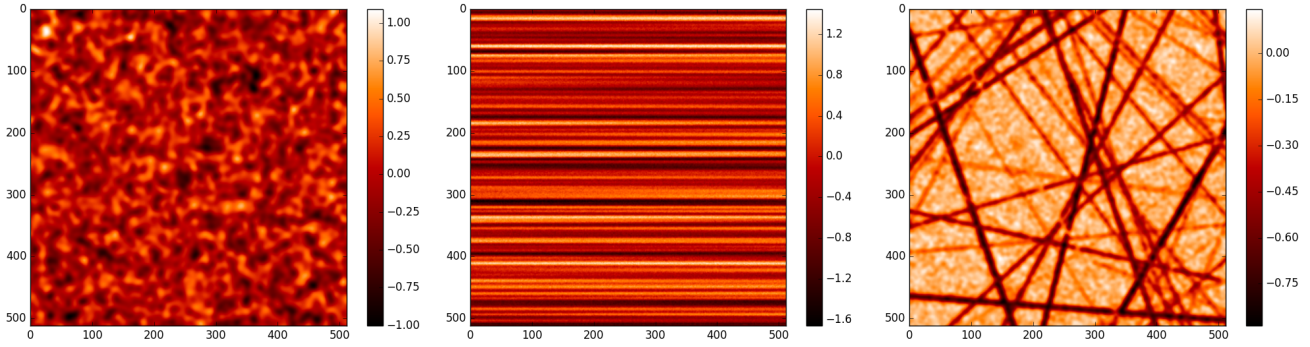
Fig. 8. The heightfields used in this paper. Left to right: isotropic bumps, brushed metal, scratched metal. These are 512^2 crops of the full 8192^2 maps. The units (horizontal and vertical) are microns (μm), so the full maps cover a square area about 8.2 mm × 8.2 mm large.

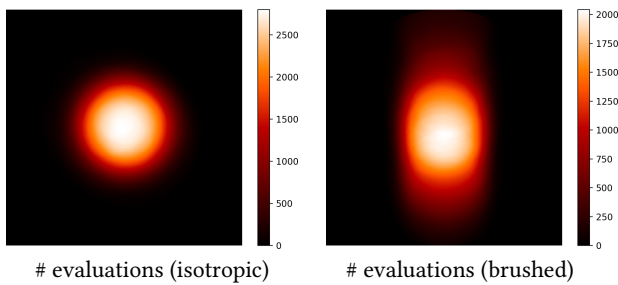# evaluations (isotropic)   # evaluations (brushed)

Fig. 9. Visualization of the numbers of Gabor kernels that are evaluated to calculate the BRDF values toward different directions. Note that the shapes of the corresponding BRDFs are captured well, and that a large number of evaluations are successfully pruned.

contribution only from the Gaussian part of each Gabor kernel. This gives us significant speed-up, allowing the indirect illumination to use more samples to converge. This is roughly equivalent to reducing the coherence area used for indirect illumination to the size of our Gabor kernels.

Practical rendering pipeline. For convenience, we separate the final rendered image into three components. First, direct illumination from point lights. Second, indirect illumination from point lights. Third, illumination from other lights, including the environment lighting, both direct and indirect. The separation allows us to use very few samples per pixel to render the most time-consuming first part, usually only 4 to 25 samples on a regular sub-pixel grid.

Spectral rendering. For each BRDF evaluation, we compute for different wavelengths ranging from 0.36 microns to 0.83 microns, i.e. the visible spectrum. We find that using 8 spectral samples is generally good enough to produce identical results to those generated using more samples. We split the wavelength range into bins and use the midpoints (not endpoints) of the bins as the samples. We follow the standard spectrum samples→XYZ→RGB method to eventually convert the spectral values to the sRGB color space.

We also find it useful to perform the top-down pruning only once using the largest wavelength. In this way, we record all contributing Gabor kernels first, then evaluate them for all spectrum samples at once. As a result, our computation time scales sub-linearly with the

number of spectrum samples, which gives us another 3× speedup, compared with brute force spectral rendering.

6 RESULTS

6.1 Heightfield and BRDF visualizations

Figure 8 shows a color-mapped visualization of the heightfields used in our results. For all heightfields, we use a discretization step of 1 micron. The heightfields were generated procedurally by inverse FFT noise generation, and (in the case of scratches) by drawing lines with randomized positions, depths, and widths.

In Figure 10, we show visualizations of the outgoing BRDF lobes of our model and geometric optics, for a fixed incoming direction and footprint (coherence area). This illustrates the differences between geometric and wave optics, and also the differences between a single-wavelength and spectral simulation. Note that the appearance of high-frequency features is clearly different in the geometric and wave solutions: the geometric solutions contain sharp folds in areas where the normal map Jacobian becomes singular [Yan et al. 2014]. The wave optics solutions have no such features, and the high frequencies in them are more reminiscent of laser speckle. Also note the significant color effects in the full spectral wave optics version.

In Figure 11, we show BRDF lobes computed with our approach (for a single wavelength) side-by-side with lobes computed using the FFT algorithm applied to equation 6. Note the close match, despite our method taking a completely different approach of Gabor kernel approximation.

6.2 Rendered results

In this section, we illustrate our method's capability to render actual scenes using wave optics, as shown in Figures 1, 2, 12, and 13. To show that the results computed separately at each frame are temporally coherent, please see the accompanying video. We implement our method using Mitsuba [Jakob 2010], and run all renderings in 720p (1280 × 720) on a 6-core Intel i7-4770K desktop at 3.5 GHz, hyperthreaded to 12 threads.

Scene configurations and performance comparisons are listed in Table 1. In general, for direct illumination from point lights, our method with a single wavelength is about 5 − 20× slower than geometric optics, and about another 3.5× slower with 8 spectral samples.

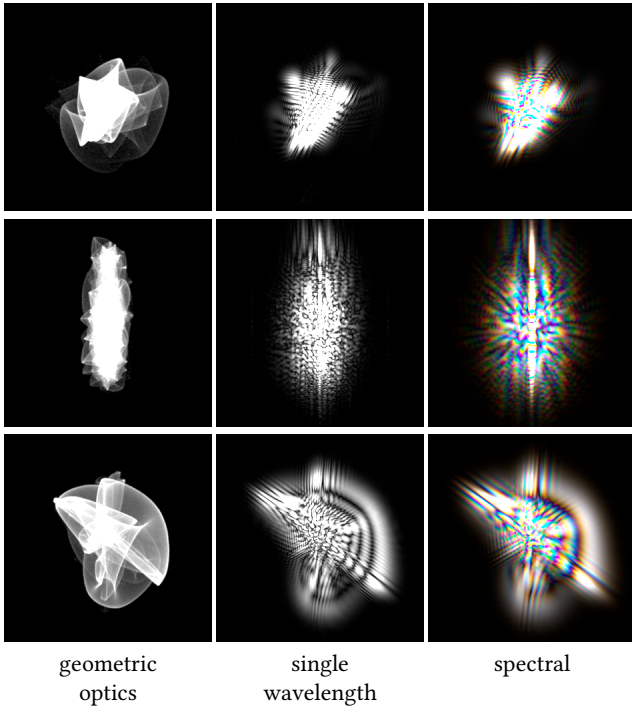| geometric | single | spectral |
| optics | wavelength | |

Fig. 10. Visualizations of the outgoing BRDF lobes on the projected hemisphere. Top: isotropic bumps, middle: brushed, bottom: scratched. Note the clearly different high-frequency features predicted by geometric and wave optics. Also note the significant color effects predicted by wave optics (here we are using 8 spectral samples).
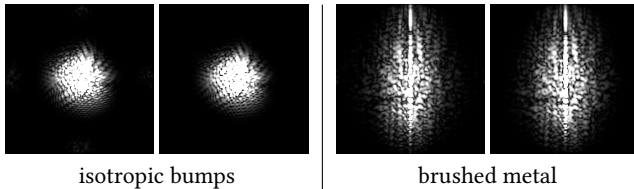

| isotropic bumps | brushed metal |

Fig. 11. Comparison of BRDF lobes computed using our Gabor kernel approach (left image in each pair) to ground truth computed by evaluating equation 6 using the FFT algorithm (right image).

This is because of the wide directional spread predicted by wave optics, as analyzed in Section 5. However, direct illumination only takes up about half of the overall computing time. Considering indirect lighting and environment lighting together, the performance of our wave optics method is within 1.5× of geometric optics, and is thus a practical solution. In the rest of this section, we will discuss individual scenes.

Patch. This is a simple scene showing a 5 *cm* × 5 *cm* patch. The camera is looking towards the center of the patch from an elevation angle of 45°. The point light is on the opposite side, and moves left and right in the video. Figure 12 shows renderings of three different heightfields (isotropic noise, brushed and scratched), each rendered using multiple wavelengths, single wavelength (0.4 microns) and geometric optics for comparison. We added isotropic noise on top

Table 1. Scene configurations including materials of the main objects and number of samples per pixel, and performance comparisons between geometric optics and wave optics with 1 and 8 spectral samples. For the performance of the *Patch* scene, we use the isotropic noise heightfield as representative.

Scene	Patch	Cutlery	Laptop	Tumbler
# Point light(s)	1	1	1	2
# Env. light	0	1	1	1
Material	Al	Ag	Al	Fe
# Samples (direct)	4	9	25	25
# Samples (ind.+env.)	N/A	256	1024	1024
Direct (geom.)	9.6s	3.1s	37.4s	19.8s
Direct (single)	3.7m	0.8m	6.4m	1.9m
Direct (spectral)	13.1m	2.4m	21.1m	6.4m
Indirect + env.	N/A	4.0m	25.1m	9.7m
All (geom.)	N/A	4.2m	25.7m	10.0m
All (single)	N/A	4.8m	31.5m	11.6m
All (spectral)	N/A	6.4m	46.2m	16.1m

of the brushes and scratches to make them more visible under the point light.

From these images, we can clearly see that our method is able to produce characteristic structures from the underlying heightfields: intuitively, round highlight for isotropic, vertical anisotropic highlight for brushed and spiderweb-like highlight for scratched. These shapes indicate the correctness of our method. Also, since different wavelengths behave differently in wave optics, colors are expected from spectral rendering.

Cutlery. This scene shows silver cutlery with strong scratches, rendered using a point light with static grayscale environment lighting, in order to make sure that the colors are from diffraction. In Figure 1, we can clearly see the colored scratches rendered using multiple wavelengths. Also, even with a single wavelength, our method is able to generate a more convincing result, as we compare with the geometric method by Yan et al. [2016]. The geometric method arguably produces harsher glints, due to the sharper folds in the BRDF lobes predicted by the $\mathcal{P}$-NDF theory.

In the video, we move the point light back and forth, so that we can see the scratches and highlights changing consistently. We also compare with geometric optics. We can clearly see strong cross-shaped highlights from geometric optics, but they're more subtle in wave optics.

Laptop. This scene shows a laptop with a roughened aluminum matte finish (modeled as a Gaussian random heightfield). It is rendered using a point light and environment lighting. We can observe colored glints in Figure 2. Albeit subtle, these colored glints are pervasively observed in the real world. To further verify this effect, we illuminated a MacBook using an LED light from a cell phone in a dark room; this leads to obvious colored highlights. Our method is able to produce perceptually similar appearance.

In the video, we move the light to show how the colored glints change. We also show comparisons with geometric optics, which produces much more "noisy" glints. This can also be observed from the BRDF images in Figure 10, where geometric optics preserves every detail from the heightfield, including sharp edges and corners,
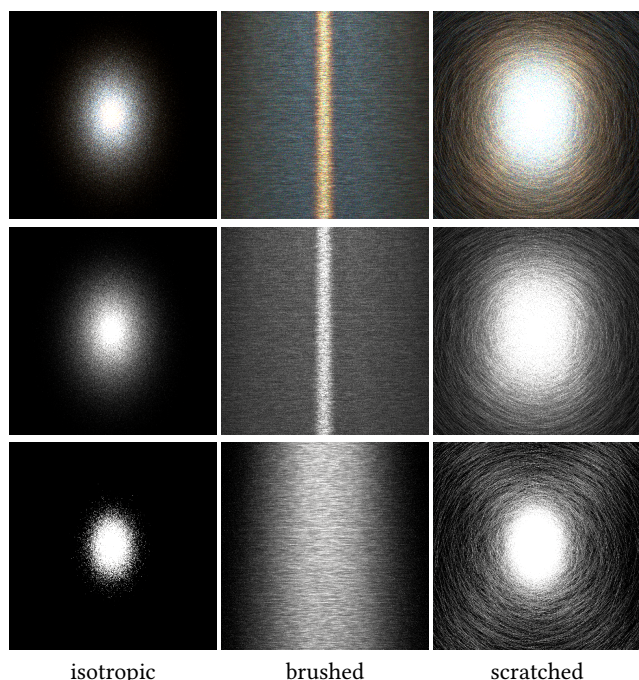
Fig. 12. The *Patch* scene showing renderings of different heightfields with a point light. (Top row) Spectral. (Middle row) Single wavelength. (Bottom row) Geometric.

isotropic   brushed   scratched

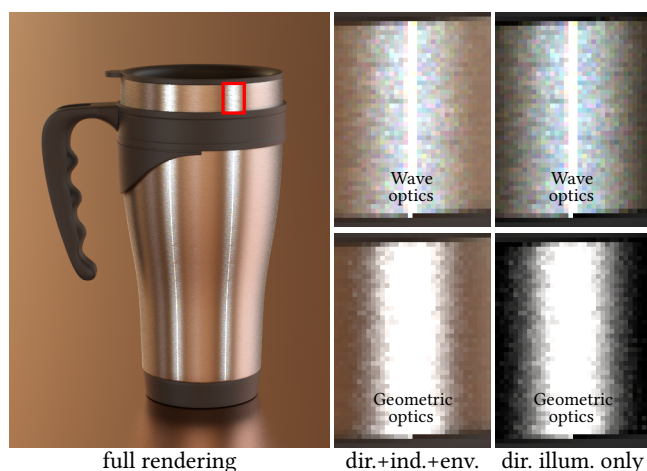full rendering   dir.+ind.+env.   dir. illum. only

Fig. 13. The *Tumbler* scene rendered with two point lights and environment lighting, showing brushed aluminum with strong anisotropy. Insets compare our spectral method (top row) with the geometric method (bottom row). We can see that the geometric method produces wider highlight peaks but narrower highlights overall, and misses the colored glints.

while wave optics smooths the details out. This is also expected in theory, because wave optics tells us that the light will "ignore" structures whose sizes are comparable or smaller than the wavelength.

Tumbler. This scene illustrates a tumbler with brushed metal on the body under two point lights and environment lighting. As shown

in Figure 13, our method is able to handle the anisotropic appearance, resulting in two vertical lines of highlights. From the insets, we can also see that the geometric method generates wider highlight peaks but narrower highlights overall, while our method is able to produce thin peaks but with much wider spread. This observation corresponds to the brushed BRDF images in Figure 10, where most energy concentrates in the central vertical line for wave optics. A similar observation was noted by Dong et al. [2015]. Moreover, the observation is also in accordance with the geometric GGX BRDF [Walter et al. 2007], well known for its "long tail" and its ability to better represent slow falloffs of highlights than other geometric BRDFs such as Beckmann. The question of why wave optics often leads to longer tails is complex; a simple though incomplete explanation is that wave optics is less influenced by finer scale roughness, leading to sharper peaks, while off-peak dropoff depends on destructive interference which tends to be somewhat random and incomplete, leading to stronger tails.

Limitations and future work. Currently our approach only looks at single specular reflection, ignoring inter-reflection, layered materials, refraction, or complex 3D structures (e.g., in biological iridescence). Our computational expense currently requires separation of direct illumination due to small lightsources from other components, and using a different (cheaper) BRDF for these components; further accelerations should be explored. Improved importance sampling techniques would also help. Physical measurement to acquire the heightfields would be another interesting addition.

7 CONCLUSION

As computer graphics has pursued ever increasing realism in material appearance, there has been a trend from strictly using geometric optics to introducing wave optics in the derivation of BRDF models and in the simulation of specific iridescence effects. In this paper we have taken this process to a new level, introducing the first practical technique for simulating full diffraction effects in completely arbitrary micron-scale height-field geometry. The result is a dramatic change in the predicted BRDF due to reflection from small areas of surface: the unrealistically sharply defined structures of geometric optics give way to softer results that depend on wavelength, introducing color into the BRDF. In practice, the new model produces softer, more natural looking reflections from microgeometry, with subtle color effects visible under sharp lighting.

BRDF models that are derived from wave optics predict different appearance than microfacet models, but the results can often be nearly matched by microfacet models with somewhat different parameters. In the case of surface detail, the geometric model will always predict reflection patterns with hard edges induced by folds in the set of reflected rays. Here moving to wave optics—the model appropriate to the scale in question—produces fundamentally different results, with smooth highlight edges and color effects.

An exciting implication of our work is that it provides the ability to represent surface features at all scales with the appropriate type of model: large features can be handled with geometry; smaller ones down to a fraction of a millimeter can be represented using geometric normal maps; and features down to wavelength scale are represented as diffracting height fields. Smaller features than

that are not optically relevant and are not needed in any visual simulation. Though the speed and memory footprint of our initial implementation can be further reduced in future work, our method is already efficient enough to use routinely in offline rendering. We can represent the appearance of rough surfaces with no excuses about wave effects being assumed irrelevant or surface details being left out of the model.

ACKNOWLEDGEMENTS

REFERENCES

P. Beckmann and A. Spizzichino. 1968. *The Scattering of Electromagnetic Waves from Rough Surfaces*. Books on Demand. http://books.google.com/books?id=nn92AAAACAAJ

Laurent Belcour and Pascal Barla. 2017. A Practical Extension to Microfacet Theory for the Modeling of Varying Iridescence. *ACM Trans. Graph.* 36, 4, Article 65 (July 2017), 14 pages. https://doi.org/10.1145/3072959.3073620

Carles Bosch, Xavier Pueyo, Stéphane Mérillou, and Djamchid Ghazanfarpour. 2004. A Physically-Based Model for Rendering Realistic Scratches. In *Computer Graphics Forum*, Vol. 23. Wiley Online Library, 361–370.

Brent Burley. 2012. Physically-Based Shading at Disney. *Technical Report* (2012).

Samuel D Butler, Stephen E Nauyoks, and Michael A Marciniak. 2015. Comparison of microfacet BRDF model to modified Beckmann-Kirchhoff BRDF model for rough and smooth surfaces. *Optics Express* 23, 22 (2015), 29100–29112.

R. L. Cook and K. E. Torrance. 1982. A Reflectance Model for Computer Graphics. *ACM Trans. Graph.* 1, 1 (1982), 7–24.

Tom Cuypers, Tom Haber, Philippe Bekaert, Se Baek Oh, and Ramesh Raskar. 2012. Reflectance Model for Diffraction. *ACM Trans. Graph.* 31, 5 (2012), 122:1–122:11.

Daljit Singh Dhillon, Jeremie Teyssier, Michael Single, Michel Milinkovitch Iarsolav Gaponenko, and Matthias Zwicker. 2014. Interactive Diffraction from Biological Nanostructures. *Computer Graphics Forum* 33, 8 (2014), 177–188. https://doi.org/10.1111/cgf.12425

Zhao Dong, Bruce Walter, Steve Marschner, and Donald P. Greenberg. 2015. Predicting Appearance from Measured Microgeometry of Metal Surfaces. *ACM Trans. Graph.* 35, 1, Article 9 (2015), 13 pages. https://doi.org/10.1145/2815618

James Harvey. 1979. Fourier treatment of nearâĂŘfield scalar diffraction theory. *American Journal of Physics* 47, 11 (1979), 974–980. https://doi.org/10.1119/1.11600

James E. Harvey and Richard N. Pfisterer. 2016. Evolution of the transfer function characterization of surface scatter phenomena. *Proc.SPIE* 9961 (2016), 9961 – 9961 – 17. https://doi.org/10.1117/12.2237083

Xiao D. He, Kenneth E. Torrance, François X. Sillion, and Donald P. Greenberg. 1991. A Comprehensive Physical Model for Light Reflection. *SIGGRAPH Comput. Graph.* 25, 4 (1991), 175–186. https://doi.org/10.1145/127719.122738

Nicolas Holzschuch and Romain Pacanowski. 2017. A Two-scale Microfacet Reflectance Model Combining Reflection and Diffraction. *ACM Trans. Graph.* 36, 4, Article 66 (2017), 12 pages. https://doi.org/10.1145/3072959.3073621

Wenzel Jakob. 2010. Mitsuba renderer. http://www.mitsuba-renderer.org.

Wenzel Jakob, Miloš Hašan, Ling-Qi Yan, Jason Lawrence, Ravi Ramamoorthi, and Steve Marschner. 2014. Discrete Stochastic Microfacet Models. *ACM Trans. Graph.* 33, 4 (2014).

Andrey Krywonos. 2006. *Predicting Surface Scatter Using A Linear Systems Formulation Of Non-paraxial Scalar Diffraction*. Ph.D. Dissertation. University of Central Florida.

Ann M Lanari, Samuel D Butler, Michael Marciniak, and Mark F Spencer. 2017. Wave optics simulation of statistically rough surface scatter. In *Earth Observing Systems XXII*, Vol. 10402. International Society for Optics and Photonics, 1040215.

Anat Levin, Daniel Glasner, Ying Xiong, Frédo Durand, William Freeman, Wojciech Matusik, and Todd Zickler. 2013. Fabricating BRDFs at High Spatial Resolution Using Wave Optics. *ACM Trans. Graph.* 32, 4 (2013), 144:1–144:14.

L. Mandel and E. Wolf. 1995. *Optical Coherence and Quantum Optics*. Cambridge University Press. http://books.google.com/books?id=FeBix14iM70C

Heylal Mashaal, Alex Goldstein, Daniel Feuermann, and Jeffrey M. Gordon. 2012. First direct measurement of the spatial coherence of sunlight. *Opt. Lett.* 37, 17 (2012), 3516–3518. https://doi.org/10.1364/OL.37.003516

Stéphane Mérillou, Jean-Michel Dischler, and Djamchid Ghazanfarpour. 2001. Surface scratches: measuring, modeling and rendering. *The Visual Computer* 17, 1 (2001), 30–45.

J.A. Ogilvy. 1991. *Theory of wave scattering from random rough surfaces*. A. Hilger.

Boris Raymond, Gael Guennebaud, and Pascal Barla. 2016. Multi-Scale Rendering of Scratched Materials using a Structured SV-BRDF Model. *ACM Transactions on Graphics* (July 2016). https://doi.org/10.1145/2897824.2925945

Jos Stam. 1999. Diffraction Shaders. In *SIGGRAPH 99*. New York, NY, USA, 101–110. https://doi.org/10.1145/311535.311546

Antoine Toisoul and Abhijeet Ghosh. 2017. Practical Acquisition and Rendering of Diffraction Effects in Surface Reflectance. *ACM Trans. Graph.* 36, 5, Article 166 (July 2017), 16 pages. https://doi.org/10.1145/3012001

Zdravko Velinov, Sebastian Werner, and Matthias B. Hullin. 2018. Real-Time Rendering of Wave-Optical Effects on Scratched Surfaces. *Computer Graphics Forum 37 (2) (Proc. EUROGRAPHICS)* 37, 2 (2018).

Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet Models for Refraction Through Rough Surfaces *(EGSR 07)*. 195–206.

Sebastian Werner, Zdravko Velinov, Wenzel Jakob, and Matthias B. Hullin. 2017. Scratch Iridescence: Wave-optical Rendering of Diffractive Surface Structure. *ACM Trans. Graph.* 36, 6, Article 207 (2017), 14 pages. https://doi.org/10.1145/3130800.3130840

Stephen H Westin, Hongsong Li, and Kenneth E Torrance. 2004. A comparison of four brdf models. In *Eurographics Symposium on Rendering. pags.* 1–10.

Ling-Qi Yan, Miloš Hašan, Wenzel Jakob, Jason Lawrence, Steve Marschner, and Ravi Ramamoorthi. 2014. Rendering Glints on High-resolution Normal-mapped Specular Surfaces. *ACM Trans. Graph.* 33, 4, Article 116 (2014), 9 pages. https://doi.org/10.1145/2601097.2601155

Ling-Qi Yan, Miloš Hašan, Steve Marschner, and Ravi Ramamoorthi. 2016. Position-normal Distributions for Efficient Rendering of Specular Microstructure. *ACM Trans. Graph.* 35, 4, Article 56 (2016), 9 pages. https://doi.org/10.1145/2897824.2925915

A ALTERNATE METHOD DERIVATION

Here we briefly sketch an alternative, but mathematically equivalent, derivation of our method. One can also view our method as approximating the rough surface by a set of small planar elements, or flakes, with one per grid cell. These flakes have soft overlapping boundaries in $\tilde{S}$, defined by a Gaussian flake shape function $K(s) = l_k^2 G_{2D}(s; 0, \sigma_k)$, which assuming a uniform grid, is the same for all flakes. K is normalized so that its integral is equal to the grid cell area. The flakes are centered at the grid cell centers m_k.

Approximating the coherence kernel as constant over a flake with $w_k = w(m_k - x_c)$, and evaluating Equation (3) as a sum over the flake approximation of the surface gives:

$$f_r(\omega_i, \omega_o) \approx \frac{\xi_1}{A_c} \left| \sum_k w_k \int_{\tilde{S}_c} K(s - m_k) R(s)\, e^{-i \frac{2\pi}{\lambda} (\bar{\psi} \cdot s)}\, ds \right|^2 \quad (18)$$

Then we can expand $R(s)$ using (1), substitute the per-flake planar height approximation, $H(s) = H(m_k) + H'(m_k) \cdot (s - m_k)$, and after some rearranging of terms, recognize that the integral has the form of a Fourier transform (5) of K. Combining these steps we get:

$$f_r(\omega_i, \omega_o) = \frac{\xi_1}{A_c} \left| \sum_k w_k \xi_2\, e^{-i2\pi \phi_k}\, \widetilde{K}\left(\frac{\xi_3 H'(m_k) + \bar{\psi}}{\lambda} \right) \right|^2 \quad (19)$$

$$\phi_k = \frac{\xi_3 H(m_k) + (\bar{\psi} \cdot m_k)}{\lambda} \quad (20)$$

$$\widetilde{K}(v) = l_k^2 e^{-2\pi^2 \sigma_k^2 \|v\|^2} \quad (21)$$

where we have assumed ξ_2 and ξ_3 are constant per flake.

Since $\widetilde{K}$ is a Gaussian, the expression in the sum can be interpreted as evaluations of Gabor kernels. If we expand out the expression inside the sum from our prior derivation (17), we find that they match exactly. This alternate derivation may provide some additional intuition about our approximation and can also be applied to other types of flake shape functions.