

Predicting Appearance from Measured Microgeometry of Metal Surfaces

ZHAO DONG*, BRUCE WALTER*, STEVE MARSCHNER and DONALD P. GREENBERG

Cornell University

The visual appearance of many materials is created by micro-scale details of their surface geometry. In this paper, we investigate a new approach to capturing the appearance of metal surfaces without reflectance measurements, by deriving microfacet distributions directly from measured surface topography. Modern profilometers are capable of measuring surfaces with sub-wavelength resolution at increasingly rapid rates. We consider both wave- and geometric-optics methods for predicting BRDFs of measured surfaces and compare the results to optical measurements from a gonio-reflectometer for five rough metal samples. Surface measurements are also used to predict spatial variation, or texture, which is especially important for the appearance of our anisotropic brushed metal samples.

Profilometer-based BRDF acquisition offers many potential advantages over traditional techniques, including speed and easy handling of anisotropic, highly directional materials. We also introduce a new generalized normal distribution function, the ellipsoidal NDF, to compactly represent non-symmetric features in our measured data and texture synthesis.

Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—*Color, shading, shadowing, and texture*

General Terms: Appearance modeling and rendering, Profilometer micro-geometry measurement

Additional Key Words and Phrases: Surface reflectance, Microfacet models, Kirchhoff scattering, Spatially-variant anisotropic BRDF

ACM Reference Format:

canonical bib reference TBD

1. INTRODUCTION

Surface reflectance models based on the microfacet theory [Blinn 1977; Cook and Torrance 1982; Walter et al. 2007] have become predominant in computer graphics because they are simple, reason-

ably general, and they have consistently proven to be an excellent fit for real materials [Torrance and Sparrow 1967; Ngan et al. 2005a; Löw et al. 2012]. It has often been observed that the assumptions of microfacet theory—namely that the surface is smooth at the wavelength scale—are frequently violated, yet the models nevertheless fit measurements well.

The way a surface reflects light is represented by its bidirectional reflectance distribution function (BRDF). The shape of a microfacet BRDF is primarily determined by the distribution of micro-scale surface normals, which is represented as a normal distribution function (NDF). Traditionally, simple distributions based on statistical assumptions about the surface, such as the Beckmann distribution, are used so that the parameters can be obtained by fitting to a relatively sparse set of BRDF measurements.

The limited expressiveness of the traditional distributions has led more recently to methods that make denser BRDF measurements, and use more general models, such as tabulations and mixture models, to represent the NDF. This gives greater expressive power; but it requires dense angular measurement, leading to slow capture processes or assumptions of large-area surface homogeneity. For spatially varying surfaces, the problem is much more difficult, and it is generally infeasible to achieve high resolution NDF acquisition.

This paper explores a different avenue for capturing normal distributions: obtaining them by measuring the microscopic surface topography directly. Using a commercial profilometer, a standard laboratory instrument that uses white-light interferometry through a microscope to measure height fields, we measure surface heights at 110nm resolution in roughly 7 seconds per 70×55μm region (Figure 1 (b)). From these measurements we predict the BRDF of the surface, avoiding the need for scattering measurements.

We investigate two ways to predict BRDFs from surface height data. Microfacet theory is based on geometric optics and derives the BRDF from the geometric surface normals. However, it overestimates the effect of wavelength-scale roughness, and we find that it must be combined with spatial filtering for good results. A second approach, uses the wave optics based, scalar Kirchhoff diffraction theory, which makes weaker smoothness assumptions and treats small detail in a more rigorous way, without the need for filtering.

The resulting distributions, from both the microfacet and Kirchhoff approaches, are very detailed and we are easily able to calculate them for very small areas of surface—two things that are very difficult to do at once with conventional optical measurements. Despite their strong theoretical differences, we find both methods produce similar BRDF predictions for our test surfaces. We further present validation experiments that show close agreement between the BRDFs predicted from the surface microgeometry and measurements from a gonio-reflectometer.

Most of our samples are highly anisotropic, and the normal distributions for small regions often exhibit noticeable skew with the maximal direction shifted away from the average surface normal. Existing parametric NDF models assume mirror symmetry and cannot represent this effect. Thus we also propose the new ellipsoid NDF model, which supports anisotropy and skew as a gener-

*Joint first authors.

E-mail: {zd,bjw,dpg}@graphics.cornell.edu, srm@cs.cornell.edu

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 0730-0301/YYYY/15-ARTXXX \$10.00

DOI 10.1145/XXXXXXXX.YYYYYYY

<http://doi.acm.org/10.1145/XXXXXXXX.YYYYYYY>

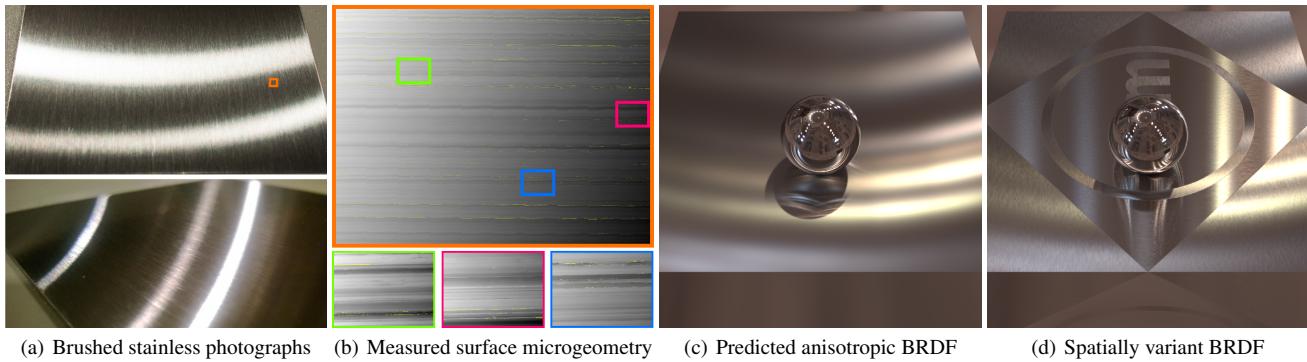


Fig. 1. We present a method to predict the appearance of complex surfaces, such as this stainless steel #4 sample, from its micro-scale geometry. Commercial profilometers can rapidly measure surfaces with sub-wavelength resolution, and we develop and validate methods to predict BRDFs from the surface data. The method can predict both average BRDFs and characteristic textures, or spatial variations, that are visually important for our example materials.

alization of the widely used GGX/Trowbridge-Reitz NDF [Walter et al. 2007; Trowbridge and Reitz 1975], and which can be efficiently evaluated and importance sampled. Using this new NDF, we explore an approach to statistically modeling spatial variation by computing ellipsoid NDF parameters for a small set of contiguous patches, and then synthesizing much larger textures with similar statistics.

This paper opens up a new avenue for capturing the appearance of surfaces with detailed normal distributions and spatial variation. It has significant advantages compared to the usual optical methods: it can achieve very high resolution in both space and angle, with relatively short measurement times and off-the-shelf instrumentation. Our current method is well suited to opaque single layer materials, such as rough metals, however due to the limitations of the profilometer and theories we utilize, it does not handle materials with multiple layers, or with strong multiple or sub-surface scattering. While surface geometry alone is insufficient to predict appearance those materials, such as paint or plastic, nevertheless surfaces form an important part of most materials. We hope our work also contributes to a better understanding of real world surfaces and how they affect light, to assist the future development of measurements strategies and models for a wider range of materials.

Some contributions of this paper include:

- (1) Propose a new approach to material appearance capture based on surface microgeometry measurement with off-the-shelf equipment, that has unique strengths compared with conventional optical BRDF measurement.
- (2) Demonstrate application of microfacet and Kirchhoff theory for measured rough surfaces including developing filtering modification for microfacet and formulation of Kirchhoff in terms of new NDF-analogous term D_K .
- (3) Validate predictions of both theories against gonireflectometer measurements for a set of challenging rough, anisotropic metal surfaces.
- (4) New parametric NDF model that supports both anisotropy and skew, called the ellipsoid NDF.
- (5) Simple texture synthesis method, based on separable assumption, that can extrapolate visually important NDF textures from a measurements of small regions.

2. RELATED WORK

Microfacet BRDF. Microfacet BRDFs [Torrance and Sparrow 1967; Blinn 1977; Cook and Torrance 1982; Walter et al. 2007]

are based on a geometric optics model of how light reflects from a rough surface, and have been shown to provide a good match to measured BRDF data for many real materials [Ngan et al. 2005b]. The surface is assumed to act as complicated curved mirror, or equivalently as a collection of many small flat reflective facets (microfacets). The resulting BRDF depends on three components: the microfacet Normal Distribution Function (NDF) term, a Fresnel term based on the material’s complex index of refraction, and a shadowing-masking term to ensure energy conservation. The NDF term is the most important in determining the pattern of reflected light, and many different parametric forms have been proposed. The Beckmann distribution, based on isotropic Gaussian random surfaces, is often used [Cook and Torrance 1982]. To model anisotropic materials, Ward [1992] used the anisotropic Beckmann distribution while omitting the Fresnel and shadowing/masking terms for simplicity. Ashikhmin and Shirley [2000] introduced an anisotropic reflection model using a Phong microfacet distribution. Walter et al. [2007] showed that the Beckmann and Phong distributions are very similar and introduced the GGX distribution to better fit a measured BRDF dataset. GGX is mathematically identical to the NDF proposed by Trowbridge and Reitz [1975] and an anisotropic extension was proposed by Burley [2012]. This paper introduces a new extension of the GGX NDF model that is better suited to modeling spatial variation in anisotropic materials. While convenient for rendering, parametric forms can be limiting and thus some methods have used more general representations such as tabulated NDFs [Ashikhmin et al. 2000; Wang et al. 2008].

By using geometric optics, microfacet theory assumes that diffraction effects can be safely neglected. This is only a safe assumption if the micro-surfaces are locally flat compared to the wavelength. In other words, the surface should not contain any roughness at scales near the wavelengths of visible light. Real surfaces often violate this assumption. Lacking access to actual surface microgeometry, prior microfacet work has not addressed this issue.

Kirchhoff Scattering. Another approach for modeling BRDFs is to use Kirchhoff theory [Beckmann and Spizzichino 1968], which is based on wave optics and can correctly predict many diffraction effects. Kajiya [1985] suggested using Kirchhoff theory to model anisotropic surface appearance but did not test against any measured data. He et al. [1991] derived a BRDF model for a broad class of Gaussian random surfaces using vector Kirchhoff theory. Stam [1999] used scalar Kirchhoff theory to derive BRDF models for periodic and Gaussian random surfaces based on the Fourier transform of their height correlation functions. His solu-

tion for periodic structures was improved in [Dhillon et al. 2014] to better reproduce diffraction-grating-like effects seen in measured reptile skins. Levin et al. [2013] used Kirchhoff theory to predict the BRDFs of a specific class of surfaces that can be manufactured using photolithography. Our measured microgeometries are irregular and not likely to fit into any of these prior categories, so we have chosen to numerically estimate the scalar Kirchhoff integrals directly rather than using one of these prior approximate solutions. Cuypers et al. [2012] extend the definition of a BRDF using Wigner distributions which can predict macroscopic near-field diffraction effects, but that lies beyond the scope of this paper and we will use the standard BRDF definition.

Some prior work has compared measured BRDFs and microgeometry statistics for particular kinds of surfaces to test Kirchhoff-based predictions. For example, Marx and Vorburger [1990] and Li and Torrance [2005] tested surfaces specially prepared to be approximately 1D or isotropic gaussian respectively, and found good agreement with Kirchhoff theory. Another study [Schröder et al. 2011] tested isotropic surfaces with varying roughness (though with lower roughness than our examples) against a more advanced Kirchhoff-based model, called Generalized-Harvey-Shack theory, and found good agreement. They extracted a statistical model of their surfaces using a combination of optical profilometer and atomic force microscopy measurements. In contrast to this prior work, we use the measured microgeometry directly rather than assuming any particular statistical model for the surfaces, and demonstrate that our method works even for highly anisotropic surfaces.

Spatially Varying BRDF. To capture important spatial variation, many methods have been proposed for capturing densely sampled BRDF measurements across surfaces. Dana et al. [1999] proposed to use a spatial gonioreflectometer to directly measure spatially varying BRDFs and bidirectional texture functions (BTFs) of real surface materials. By measuring many pixels at once, camera-based acquisition systems can be used to measure SVBRDFs [Gu et al. 2006], BTFs [Dana 2001; Müller et al. 2005; Han and Perlin 2003], and reflectance fields [Garg et al. 2006] with dense sampling of view and light directions over the hemisphere. Because of the need to sample a 6D space, all these methods require complicated and specialized setups, and they either require time-consuming acquisitions or achieve low angular resolution. In comparison, we measure NDFs directly, avoiding many of the difficulties and limitations of radiometric measurements and achieving very high resolution in both space and angle, though not over large surface areas.

Data-driven BRDF Fitting. To limit the need for high angular resolution, many methods fit parametric models to their spatially varying BRDF measurements. Gardner et al. [2003] used a linear light source to scan a surface and fit an isotropic Ward model at each pixel in a fixed view. Lensch et al. [2003] separate surfaces into different materials and use Lafortune BRDF [Lafortune et al. 1997] fit to each material cluster as a basis in which to represent spatial variation. Similarly, Goldman et al. [2005] use a basis of isotropic Ward BDRFs to represent SVBRDFs. These methods capture detailed spatial variation, but their low angular resolution requires assuming that particular low-parameter BRDF models are sufficient. Using device setup similar to Gardner et al., Wang et al. [2008] capture a denser 2D slice of the 4D BRDF at each pixel in a fixed view, then combine sparse data from many pixels to estimate tabulated NDFs. Their NDFs are much lower resolution than with our method, and their merging of NDFs relies on rotational symmetries that our data shows do not always hold.

Dong et al. [2010] takes another approach to leveraging the limited variation of BRDFs on a sample. They assume that reflectance over a given material sample forms a low-dimensional manifold,

allowing a two-phase process in which isolated measurements with relatively high angular resolution, performed using a hand-held BRDF capture device, are followed by a standard fixed-view image capture under a varying area source. Based on microfacet BRDF theory, a small number of measured NDFs then serve to infer a spatially varying BRDF over the whole surface.

Modeling Microgeometry. Zhao et al. [2011] used X-ray computed tomography (CT) to measure the volumetric micro-structure of cloth, which was combined with estimated appearance parameters from photographs to generate highly realistic images. In contrast, our focus is on surfaces and surface microgeometry, which can be captured much more efficiently, and at higher resolution, with profilometry than by CT. Yan et al. [2014] demonstrated a method to compute the scattering directly from a detailed microgeometry using geometric optics, however they did not use measured microgeometry or validate against BRDF measurements.

McKnight et al. [2001] used optical profilometers to measure isotropic dielectric surfaces and compared the results to geometric optics and Kirchhoff-based predictions for in-plane BRDF measurements. They found both methods matched to the data well, but unlike our results, they found no need to filter the geometric results, likely due to the lower resolution of their geometric data. Sung et al. [2002] used confocal microscopy to measure the microgeometry of embedded flakes in an isotropic metallic paint and compare it to in-plane measured BRDF values under a geometric optics assumption. They use a local least squares plane fitting to compute surface normals that is similar to our filtering, however the paper states this is being used for data interpolation purposes while we explicitly use it to account for wavelength scale effects missing from geometric optics. Unlike these prior methods, we consider more general anisotropic surfaces, including ones with higher roughness, validate against both in and out of plane measurements, and provide essential tools for graphics applications such as practical BRDF models and consideration of spatial variation.

3. METHOD OVERVIEW

Our goal in this work is to reproduce the complex appearance of metal surfaces by measuring their surface microscale geometry (with sub-wavelength resolution). An example is shown in Figure 1, where we scan a brushed stainless steel sample to obtain its surface microgeometry from which we are able to predict its highly anisotropic average BRDF as well as visually important aspects of its characteristic spatial texture.

Compared to more conventional direct BRDF measurement, our indirect microgeometry approach has several advantages, especially for anisotropic surfaces. Our approach provides extremely high angular and spatial resolution in the BRDF, much higher than is typically feasible with direct approaches. It requires only off-the-shelf equipment (optical profilometers are available from multiple manufacturers and used in photolithography and nanofabrication industries). And it provides insight into the underlying physical causes of a material's appearance.

There are also some potential disadvantages to our microgeometry approach. Scanning time is proportional to spatial extent, so is not well suited to capturing low-frequency spatial features. Our test surfaces contain considerable variation at small scales but are homogeneous at large scales. Our method currently only attempts to model first-surface reflections, so it only applies to materials, such as metals, where this is the dominant effect. Also as an indirect method, it requires validation especially since real surfaces generally do not satisfy all the assumptions of the surface scattering theories used. A major contribution of this paper is to explore




Fig. 2. A photograph of the optical profilometer used for this paper (image courtesy of Zygo®) and an illustration of how it measures surface heights.

the feasibility of this approach and validate the BRDF results for a small set of challenging, extremely anisotropic metal samples.

The rest of the paper is organized as follows. We begin by discussing optical profilometers and how we use them to scan surface microgeometry in Section 4. Then we discuss two different theories for predicting BRDFs from surface geometry. Section 5 describes microfacet theory, which is based on geometric optics and is widely used in graphics, as well as the spatial filtering modification we found necessary when applying it to our data. Section 6 reviews scalar Kirchhoff theory, a popular wave optics approximation for predicting scattering from rough surfaces. Both theories include approximations and local smoothness assumptions that our surfaces do not fully obey, so it is important to test their accuracy. In Section 7, we measure five metal surfaces, one isotropic and four brushed, and compare profilometer-based predictions against data from our gonireflectometer, a specialized device that directly measures BRDF values. Despite their considerable differences, we find that both theories make similar predictions and show generally good agreement with the gonireflectometer data for our test surfaces.

Both microfacet and Kirchhoff theory predict that the primary determinant of a BRDF's shape is the surface's normal distribution function (NDF) or a function of a similar form that we can call its effective NDF. Since in general BRDFs are 4D functions while NDFs are only 2D, this is a significant simplification and much easier to measure, store, and visualize. For generality, we represent our NDFs as 1025×1025 tabulated data, which can represent even extremely anisotropic and narrow NDFs such as found in some of our samples. However the tabulated NDF form can be data-intensive and inconvenient for rendering, especially if we want to represent spatially varying BRDFs. In Section 8 we introduce a new 5 parameter extension of the popular GGX distribution that captures the most significant NDF features in our data. This new distribution is called the ellipsoid NDF, and supports both anisotropy and skew (or shifting of the maximum of the distribution away from the center). While not present in the large-scale NDFs, skew is a feature that we see in our fine-scale NDF data and believe is important in modeling their spatial variation.

Finally we explore a simple proof-of-concept method for modeling the characteristic textures of brushed metals in Section 9 by extrapolating from a relatively small set of measurements. Computing NDFs for many smaller regions provides information about how the BRDF varies spatially. We capture this data for two short, narrow stripes aligned with the surface's principal directions, fit the data to the ellipsoid NDF, and then use a simple Fourier technique to synthesize much larger textures with similar statistics. Our results show that even this simple approach produces significantly more realistic images than just using the average or large-scale NDF.




Fig. 3. Profilometer measured surface microgeometry for aluminum #4 and stainless steel #4 samples. The blue regions correspond to a single profilometer scan and contain measured heights for 640×480 surface points at 110nm horizontal resolution. Larger areas (orange) are measured by combining multiple scans. Grayscale intensity corresponds to surface height, and yellow indicates points where the profilometer could not determine a height.

4. OPTICAL PROFILOMETER

Optical profilometers are commonly used in nanofabrication and related areas to obtain accurate surface measurements. The model we use is a Zygo® NewView™ 7300. This device, illustrated in Figure 2, combines microscope optics with a white light interferometer. Light is split into two beams, with one reflecting off the surface and the other from a reference mirror, and then recombined causing interference effects dependent on the relative lengths of the two paths. A camera observes the surface while the interferometer is scanned vertically allowing it determine the surface height at each pixel with sub-nanometer vertical resolution. Horizontal resolution depends on the camera and microscope optics used. In our case, we use a 640×480 camera set for 110nm horizontal resolution to scan a $70 \times 53 \mu\text{m}$ region of the surface, and each scan took approximately 7 seconds. A horizontal translation stage moves the sample between scans for rapidly capturing larger regions. Detailed info about this device can be found online [Zygo® 2011].

4.1 Profilometer Surface Measurement

To test our approach, we selected a set of five metal samples: QPanel, stainless steel #4, stainless steel #3, aluminum #4, and copper #4. The QPanel is a steel plate with an isotropic rough finish, while the other four are commercially available brushed metal surfaces. The number refers to the brushing type, with lower numbers corresponding to coarser brushing. Surface roughness also depends strongly on material properties; aluminum #4 is much rougher than stainless steel #4.

For each sample, we measure its surface microgeometry using the profilometer. Each scan generates a 640×480 grid of height measurements covering a $70 \times 53 \mu\text{m}$ region of the surface, and we perform multiple contiguous scans. For example, a set of 10×10 scans generates roughly 30 million height values covering an area of 0.37mm^2 in about 12 minutes. The 110nm horizontal resolution allows us to measure surface detail down to and slightly below the wavelength range of visible light. There is a roughly $2 \mu\text{m}$ misalignment between the boundaries of neighboring scans due to

ψ	Direction from which light is incident
ω	Direction in which light is scattered
\mathbf{h}	Half-direction for reflection, $\mathbf{h} = (\psi + \omega) / \ \psi + \omega\ $
\mathbf{n}	Large-scale, or average, surface normal
\mathbf{m}	Microsurface, or local, surface normal
D_M	Microfacet Normal Distribution Function (NDF)
D_K	Kirchhoff distribution (analogous to NDF, eq. 6)
G	Shadowing-masking function
F	Fresnel term
f_r	Bidirectional reflectance distribution function (BRDF)
λ	Wavelength of light
λ_q	Effective Kirchhoff wavelength (eq. 5)
\mathcal{X}_+	$\mathcal{X}_+(x) = 1$ if $x \geq 0$ and is zero otherwise

Fig. 4. List of important symbols.

limitations of the mechanical translation table, but we process each scan individually and are not sensitive to this error.

Sample profilometer measurements for two of our samples are shown in Figure 3. The data reveals the surfaces in great detail, but also contains some artifacts. In some cases, the profilometer does not return an height for a given point (shown as yellow points). This can happen in locations with very steep slopes or too much local surface complexity, creating ambiguous fringing patterns in the interferometer. Such points are treated as missing data in subsequent processing and constitute less than 5% of the measured area in each of our datasets. More rarely, the profilometer sometimes returns height values that are inconsistent with their neighbors, which we refer to as outliers. These are likely points that should have been marked as missing data but did not quite meet the interferometer’s ambiguity threshold. We experimented with a variety of filters to remove these artifacts, such as hole-filling and outlier-rejection. However we noted these filters had very little effect on either our microfacet or Kirchhoff-based BRDF predictions, and we ultimately discarded the filters as unnecessary for our purposes.

Next we discuss two theoretical approaches for predicting BRDFs from our measured surface data: microfacet and Kirchhoff. These theories only consider first-surface reflection, which we believe is the dominant mode in our test samples. They do not attempt to model multiple or sub-surface scattering, which can be very important in other types of materials.

5. MICROFACET THEORY

Microfacet theory treats a surface as a curved mirror, or equivalently as a set of tiny flat mirror facets, that obey geometric optics. At each point, the surface reflects incident light into the corresponding mirror reflection direction according to the local surface normal. Therefore light coming from direction ψ will be reflected into a given direction ω only by those parts of the surface where local surface normal is equal to the half direction \mathbf{h} , defined as: $\mathbf{h} = (\psi + \omega) / \|\psi + \omega\|$. The area density with a given local normal is described by the surface’s normal distribution function (NDF).

The BRDF corresponding to a microfacet model is given by:

$$f_r(\psi, \omega) = \frac{D_M(\mathbf{h}) F(\psi \cdot \mathbf{h}) G(\psi, \omega)}{4 |\psi \cdot \mathbf{n}| |\omega \cdot \mathbf{n}|} \quad (1)$$

where D_M is the surface’s normal distribution function, F is a Fresnel term, G is a shadowing-masking term, and \mathbf{n} is the large-scale or average normal of the surface. See Walter et al. [2007] for a detailed derivation of this equation. The Fresnel term depends on the material type and can be computed based on its refractive index and




Fig. 5. Comparison of naive and filtered geometric NDF estimation for our Al#4 sample. The NDF on the left is computed using the simple geometric NDF definition, while the one in the middle is computed using our gaussian weighting filter. Using the filtered NDFs results in a much better match to NDF data inferred from our gonioreflectometer BRDF measurements on the right. Color intensity corresponds the NDF value for the corresponding normal direction, after projection from the hemisphere to the unit disk.

extinction coefficients. The shadowing-masking term is needed for energy conservation but its value is typically close to one except for near-grazing angles. The most interesting and important term in determining the BRDF shape is the normal distribution function.

Microfacet theory neglects wave effects, such as diffraction, and may not be accurate for surfaces with roughness at scales near the wavelength of light. All our test surfaces have roughness at all the scales we could measure, including down to and below the wavelengths of visible light. Most prior microfacet work has estimated effective NDFs by fitting to BRDF data, without actually measuring the real surface geometry or examining its relationship to the fitted NDF. However our goal is to predict the BRDF from the measured surface geometry, and we find that naively computing the NDF according to its geometric definition does not work well in our case. From wave optics it is well known that surface detail below wavelength scale has much less effect on the reflection pattern, but we did not find any consensus on how best to account for this effect in a geometric optics context. Thus with some experimentation, we developed a filtered geometric NDF estimation method, described below, that uses a smoothing kernel to more accurately predict the BRDFs for our test samples.

5.1 Filtered Geometric NDF Estimation

A surface’s normal distribution function (NDF) is a density function over the sphere of directions that is proportional to the surface area with a given surface normal \mathbf{m} . The NDF can be defined geometrically as:

$$D_M(\mathbf{m}) = \lim_{|\Omega_m| \rightarrow 0} \frac{A(\Omega_m)}{|\Omega_m| A_S^\perp} \quad (2)$$

where Ω_m is a small solid angle containing the direction \mathbf{m} , $A(\Omega_m)$ is the area of the subset of the surface with normals inside Ω_m , and A_S^\perp is the total projected surface area in the direction of the average surface normal \mathbf{n} . Histograms provide a convenient way to numerically estimate NDFs. Instead of taking the limit, we discretize the sphere of directions into bins and evaluate Equation 2 for the finite Ω_m corresponding to each bin. For heightfield surfaces such as ours, the NDF is restricted to the hemisphere around the average surface normal. To represent our tabulated NDFs, we project this hemisphere onto the unit disk, and then embed it into a square which is discretized into a 1025×1025 regular grid.

The straightforward way to compute the NDF would be to triangulate the height data and evaluate the histogram using each triangle’s area and normal. However, this naive approach estimates broad NDFs, such as the one shown on the left side of Figure 5, which predict BRDFs that are much wider than the actual BRDFs

as measured by our gonireflectometer. The naive approach ignores wavelength whereas we know from wave optics that surface detail below wavelength-scale has much reduced influence on the BRDF.

Instead we developed a modified NDF estimation method that uses a Gaussian filter to reduce the influence of small scale features. At each surface data point, we fit a plane to the local height data that minimizes the squared vertical distances to the plane, weighted by a 2D Gaussian kernel based on horizontal distance from the point. This process is also known as locally-weighted linear least squares fitting. The normal to the fitted plane is taken as the effective local surface normal at that point, and by accumulating a histogram over many surface points, we estimate the effective NDF for the surface. The Gaussian kernel tends to filter out small scale roughness, resulting in narrower NDFs that much more closely match our measured BRDF data as shown in Figure 5.

We experimented with different Gaussian kernel sizes and found that a Gaussian with a standard deviation of one micron ($\sigma = 1\mu\text{m}$) works well, so we have used this fixed kernel size for all our results. However, we observe that the results are not very sensitive to small changes in this parameter. One micron corresponds to roughly two wavelengths for green light, which fits well with the intuition of filtering out surface details that are wavelength scale and smaller.

6. KIRCHHOFF THEORY

Scalar Kirchhoff theory is a scattering approximation based on wave optics which is often used to predict diffraction effects. Conceptually, the incident light induces secondary sources on the surface, based on the Fresnel equations for a locally flat surface. These generate secondary waves that radiate outward and combine to form the scattered distribution. The contributions of these secondary sources are summed using complex numbers to account for phase and interference.

Using scalar Kirchhoff theory, one can derive the following equation for the BRDF:

$$f_r(\psi, \omega) = \frac{\left| \int_{\mathcal{S}} R(\vec{s}, \psi, \omega) (\vec{q} \cdot \mathbf{m}(\vec{s})) e^{-i(\vec{q} \cdot \vec{s})} d\vec{s} \right|^2}{16\pi^2 A_S^\perp |\psi \cdot \mathbf{n}| |\omega \cdot \mathbf{n}|} \quad (3)$$

where the integral is over the surface \mathcal{S} , \vec{s} are points on the surface, $\mathbf{m}(\vec{s})$ are the local surface normals, R is the local Fresnel reflectance coefficient, and A_S^\perp is the total projected area of the surface in the direction of the average surface normal \mathbf{n} . A derivation of this equation is included in the supplemental material and, other than notation, is similar to derivations in standard texts [Beckmann and Spizzichino 1968; Ogilvy 1991; Stam 1999]. We define the vector \vec{q} as:

$$\vec{q} = \frac{2\pi(\psi + \omega)}{\lambda} = \frac{4\pi}{\lambda_q} \mathbf{h} \quad (4)$$

where λ is the wavelength of the light. The length of \vec{q} depends on both the wavelength λ and the angle θ_d between the incident direction ψ and exitant direction ω . The vector \vec{q} is closely related to, and collinear with, the half direction \mathbf{h} from microfacet theory. To make this clearer, we define an effective wavelength λ_q as:

$$\lambda_q = \frac{\lambda}{\frac{1}{2} \|\psi + \omega\|} = \frac{\lambda}{\cos\left(\frac{\theta_d}{2}\right)} \quad (5)$$

Let us assume that the Fresnel coefficient R is independent of surface position, which is a common simplifying assumption used

in Kirchhoff scattering. Then we define the following function:

$$D_K(\vec{q}) = \frac{1}{4\pi^2 A_S^\perp} \left| \int_{\mathcal{S}} (\vec{q} \cdot \mathbf{m}(\vec{s})) e^{-i(\vec{q} \cdot \vec{s})} d\vec{s} \right|^2 \quad (6)$$

which we call the Kirchhoff distribution, and rewrite equation 3 as:

$$f_r(\psi, \omega) = \frac{D_K(\vec{q}) |R(\psi, \omega)|^2}{4 |\psi \cdot \mathbf{n}| |\omega \cdot \mathbf{n}|} \quad (7)$$

Now the microfacet and Kirchhoff BRDFs (equations 1 and 7) have a remarkably similar form. The Kirchhoff distribution D_K is analogous to the microfacet NDF D_M . The same Fresnel equations are typically used for both R^2 and F , making them interchangeable. Kirchhoff has no term matching G because in order to make the problem tractable, Kirchhoff derivations usually neglect non-local effects including shadowing-masking. However this causes Kirchhoff models to violate energy conservation for near-grazing angles, and the same shadowing-masking approximation terms used in microfacet theory are sometimes added as a practical enhancement. After adding such a term and some rewriting, we finally get the following Kirchhoff-based BRDF model which we will use in our results.

$$f_r(\psi, \omega) = \frac{D_K(4\pi\mathbf{h}/\lambda_q) F(\psi \cdot \mathbf{h}) G(\psi, \omega)}{4 |\psi \cdot \mathbf{n}| |\omega \cdot \mathbf{n}|} \quad (8)$$

One significant difference is that $D_K(\vec{q})$ is a 3D function while $D_M(\mathbf{h})$ is a 2D function because \mathbf{h} is required to be a unit vector. However, for sufficiently rough surfaces, D_K is often only very weakly dependent on the length of \vec{q} and our tests indicate this is true for our surfaces. In such cases, we can approximate \vec{q} as a fixed length vector (e.g., by treating λ_q as a constant), thus effectively reducing D_K to a 2D function. Then we can view the Kirchhoff distribution D_K as being an alternate way to compute effective NDFs for a microfacet-like BRDF.

In our implementation, we evaluate Equation 6 using a Monte Carlo solver for a dense set of values and then store the results using the same discrete 1025×1025 format as for the geometric NDFs. This is significantly more expensive than the geometric NDF estimation, but only needs to be done once per measured surface patch. More details about this process are given in Appendix A

Despite their very different definitions, we find that D_K and our filtered D_M estimation generate fairly similar predictions for our test surfaces. The strong similarity between the Kirchhoff and microfacet models may help explain why microfacet models have been so successful despite their assumptions being frequently violated. Prior works have derived microfacet BRDFs as approximate solutions to the Kirchhoff integrals for particular statistical roughness models [Beckmann and Spizzichino 1968; Stam 1999], but we have not seen the relationship expressed in the more general form of equation 6 before.

7. GONIOREFLECTOMETER AND NDF VALIDATION

Thus far we have discussed how we can measure the microgeometry of a surface using an optical profilometer and how we can apply either microfacet or Kirchhoff theory to predict its BRDF from this measured data, at least for surfaces where first surface reflection is the dominant scattering effect. Next we need to test the accuracy of these predictions, especially since our test surfaces do not fully satisfy the assumptions of either theory.

One important theoretical assumption is that the surfaces are sufficiently smooth. The literature gives various “rules of thumb” for




Fig. 6. **Left:** Average fraction of the surface visible as a function of polar viewing angle for our smoothest (stainless #4) and roughest (aluminum #4) surfaces. Visibility is estimated by tracing rays against our measured surface data. Our gonioreflectometer measurements are concentrated at small angles (e.g., $< 40^\circ$) where the visibility is close to one and G can be neglected without introducing much error. **Right:** Example spectral reflectance data from gonioreflectometer measurements of our stainless steel #4 surface for a variety of angles. The spectra are all fairly simple and similar which supports our finding that the NDFs for our surfaces do not show significant variation with wavelength.




Fig. 7. Illustration of gonioreflectometer measurement setup and NDF projection onto unit disk for visualization. The red rectangle represents the measured angle coverage over the projected hemisphere (green circle).

when these theories can be safely applied, which can be roughly summarized as: at wavelength scales the surface should either have low curvature for Kirchhoff, or be essentially planar for microfacet. Kirchhoff-based predictions are often assumed to be more accurate because of its weaker assumptions and ability to predict a wider range of phenomena (e.g., diffraction gratings). However since our test surfaces violate the smoothness assumptions of both theories, it is not immediately obvious if either theory will be accurate, much less which will perform better.

To test the theoretical BRDF predictions, we also measure our test surfaces using a gonioreflectometer, which is a specialized instrument for directly measuring BRDF values. Our gonioreflectometer [Li et al. 2005] consists of a collimated light source, a spectrometer detector, and a sample holder. The motor-controlled mechanical design moves the components with high angular resolution to measure the scattered light over a wide range of incident and reflected directions. One limitation of our device is that it only has 3 motors so it cannot reach every possible configuration the 4D BRDF space. It also does not measure spatial variation but instead the average reflection from a roughly 4mm diameter region on the sample, where the actual area varies with viewing angle. Its high angular resolution and radiometric accuracy make it an excellent tool for validating our predicted BRDFs. For these measurements, the light source and detector are kept fixed while rotating the sample, which allows us measure NDF details down to a resolution of

roughly 1.5 degrees, limited primarily by the aperture sizes of the source and detector.

Based on the theoretical importance of NDFs, we select a set of measurements that span a wide range of half directions, \mathbf{h} . Since 4D BRDF space is hard to visualize, we also project these measurements into the 2D NDF space by inverting the Equation 1 to get measurements of the effective NDF:

$$D(\mathbf{h}) \approx \frac{4 |\psi \cdot \mathbf{n}| |\omega \cdot \mathbf{n}| f_r(\psi, \omega)}{F(\psi \cdot \mathbf{h})} \quad (9)$$

where $f_r(\psi, \omega)$ is a gonioreflectometer measurement and we assume the shadowing-masking term can be neglected (i.e. $G \approx 1$), which is generally true for non-grazing angles (see Figure 6 left). We estimate the Fresnel term F using published refractive index and extinction coefficients for the metal type (i.e. steel, aluminum, or copper). For these tests, we fix the angle between the incident, ψ , and reflected, ω , directions ($\theta_d = 14.14^\circ$) while varying the half direction, \mathbf{h} , as illustrated in Figure 7. For each sample, we measure around 4000 half directions arranged in a customized pattern designed to cover the most interesting parts of the NDFs, which takes roughly 3 hours on our gonioreflectometer.

Besides the potential inaccuracies of microfacet and Kirchhoff theory (e.g., both assume smoothness and neglect multiple scattering), there are several other issues to keep in mind when comparing data from these two instruments. We currently cannot match the exact measurement region, either in location or size, between the profilometer and gonioreflectometer. Instead we try to measure large enough regions to provide a reasonable approximation to the average BRDF for each test sample. The profilometer measurements cover smaller areas and thus exhibit greater variability. There are also small angular misalignments between the two instrument's coordinate systems. The Fresnel term neglects surface imperfections (e.g., oxidation) and relies on generic metallurgical data, which causes some uncertainty in the absolute magnitudes of the results. The gonioreflectometer's source and detector apertures limit its ability to resolve fine NDF features below about 1.5 degrees. This limit (sometimes called the instrument signature) was determined empirically by measuring the NDF of a first-surface mirror with the gonioreflectometer. In these comparisons, our profilometer predictions have already been convolved by the measured gonioreflectometer instrument signature to account for its resolution limits.

Our gonioreflectometer measures spectral data at 10nm resolution over the visible range; however, in our measurement data the NDFs for our samples are mostly independent of the wavelength over this range (Figure 6 right). Thus to simplify the presentation, we only show data for a wavelength λ of 550nm in these comparisons. For the Kirchhoff results we have set $\lambda_q = 555\text{nm}$ to match the θ_d configuration used in the gonioreflectometer measurements.

7.1 Large-scale NDF Results

Next we test the predictions for our five test surfaces by comparing the effective NDF inferred from our gonioreflectometer measurements to the profilometer-derived NDFs predicted by the (microfacet-based) filtered geometric and Kirchhoff approaches.

Our QPanel surface has an isotropic appearance. We measured a set of 20×20 scans on the profilometer, and Figure 8 shows the measured effective NDF from our gonioreflectometer data and the predicted NDFs from our profilometer data using our geometric and Kirchhoff methods. As the plots show, both methods predict a NDF that is nearly isotropic and in close agreement with the go-

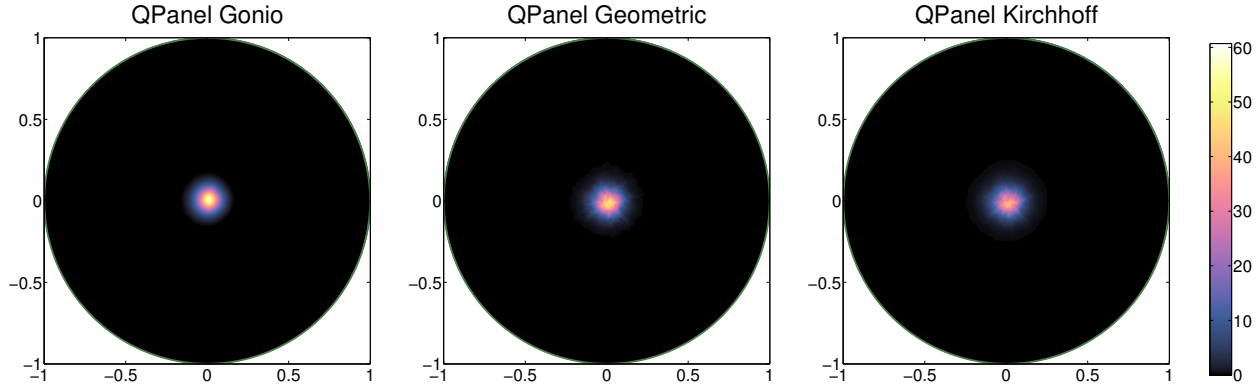


Fig. 8. Comparison of measured and predicted NDF estimates for isotropic QPanel sample. (Left) NDF inferred from gonioreflectometer measurements, (middle) NDF predicted by our geometric filtered normal method, and (right) predicted by Kirchhoff integral. For this sample the predicted NDFs are quite close to each other and to the gonioreflectometer data.

nioreflectometer data. In this case, the geometric NDF is the better match while the Kirchhoff NDF is slightly wider.

The four brushed metal samples are highly anisotropic, and for these we used sets of 4×100 profilometer scans, with the long axis aligned perpendicular to the brushing direction to better estimate the average surface NDF. Figure 9 compares results from our geometric filtered normal estimation method and Kirchhoff-based predictions to the gonioreflectometer-derived NDFs. Overall the geometric and Kirchhoff approaches produce very similar predictions and both generally match the highly anisotropic shapes of the gonioreflectometer well. In general the geometric method does a slightly better at matching the NDF lobes in the wide (or cross-scratch) direction, particularly for aluminum#4, while in the Kirchhoff is slightly better at matching the narrow (or along-scratch) widths. However this data does not present a clear case for preferring the geometric or Kirchhoff methods and both seem sufficiently accurate for computer rendering applications. For the remainder of this paper we will use the geometric method, mainly because it is easier to compute.

Rendered images generated from our profilometer-based tabulated NDFs for each of our samples are shown in Figure 10. For each material, we render sections of two cylinders with the BRDF, or brushing direction, rotated 90 degrees on the left cylinder, and lit by the St. Peters HDR environment map. For the isotropic QPanel, both cylinders look similar, but for the brushed metals, the brushing direction has an enormous effect on the appearance. Such narrow and highly anisotropic NDFs are quite difficult to measure with conventional techniques. The BRDFs in these images are spatially homogenous and the texture-like patterns come from high frequency variations in the estimated NDFs, which would likely be reduced by measuring larger areas with the profilometer.

8. ELLIPSOID NORMAL DISTRIBUTION FUNCTION

In this section, we introduce a new parametric NDF model with the goal of compactly representing the anisotropy of our NDF data as well as the asymmetric skew which occurs in tabulated NDFs for smaller surface regions such as in the examples in Figure 11. While not generally present in the average NDFs, we believe skew is important for modeling spatial variation and will be essential in the next section. None of the prior parametric models allow all these kinds of asymmetric features. The ellipsoid NDF model is con-

trolled by five parameters with intuitive meanings and is efficient for rendering while still providing a good fit to our measured data.

The GGX NDF is a single parameter NDF model proposed [Walter et al. 2007] to fit some measured data, from ground glass, better than the then standard Beckmann NDF [Cook and Torrance 1982]. It has since become widely used along with a two-parameter, anisotropic extension called GTR2aniso [Burley 2012]. GGX is mathematically identical to an earlier NDF model, Trowbridge-Reitz [1975]. The Trowbridge-Reitz NDF was derived from the normal distribution of a spheroid, which is a particular type of ellipsoid with two axes of equal length.

The ellipsoid NDF corresponds to the normal distribution of an arbitrary ellipsoid. It generalizes the GGX/Trowbridge-Reitz NDF to more degrees of freedom while also providing a geometric interpretation that is useful for computing related quantities. The ellipsoid NDF is defined as:

$$D(\mathbf{m}) = \frac{\mathcal{X}_+(\mathbf{m} \cdot \mathbf{n})}{\pi |\mathbf{A}| \|\mathbf{A} \mathbf{n}\| \|\mathbf{A}^{-T} \mathbf{m}\|^4} \quad (10)$$

where \mathbf{m} is a microsurface normal, \mathbf{n} is the large-scale, or average, normal of the surface, and \mathbf{A} is a 3×3 matrix with determinant $|\mathbf{A}|$ and inverse transpose denoted as \mathbf{A}^{-T} . The normals are represented as unit length column vectors (i.e. $\|\mathbf{m}\| = \|\mathbf{n}\| = 1$). Surface NDF models are typically restricted to be zero outside of a hemisphere centered on the average surface normal \mathbf{n} , which is equivalent to using only half of the ellipsoid. The numerator formalizes this restriction using the indicator function for positive numbers (i.e. $\mathcal{X}_+(x) = 1$ if $x \geq 0$ and is zero otherwise). Equation 10 corresponds to the normal distribution of an ellipsoid defined as the points $\bar{\mathbf{p}}$ that satisfy $\bar{\mathbf{p}}^T \mathbf{A}^T \mathbf{A} \bar{\mathbf{p}} = C_e^2$ for some constant C_e .

The matrix \mathbf{A} controls the distribution shape. An ellipsoid has six degrees of freedom (three axis lengths and three for orientation), but the NDF normalization constraint (i.e. $\int D(\mathbf{m}) |\mathbf{m} \cdot \mathbf{n}| d\mathbf{m} = 1$) reduces this to five. It is convenient to specify \mathbf{A} as the product of a rotation matrix \mathbf{R} and a scaling matrix \mathbf{S} as:

$$\mathbf{A} = \mathbf{S} \mathbf{R} \quad \text{where} \quad \mathbf{S} = \begin{bmatrix} \alpha_x & 0 & 0 \\ 0 & \alpha_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad \mathbf{R}^T \mathbf{R} = \mathbf{I} \quad (11)$$

The scaling matrix provides two degrees of freedom in addition to three for the rotation, to span the space of all possible ellipsoid distributions. Let us assume we are working in a coordinate system where the surface normal \mathbf{n} is aligned with the z-axis. If the rotation

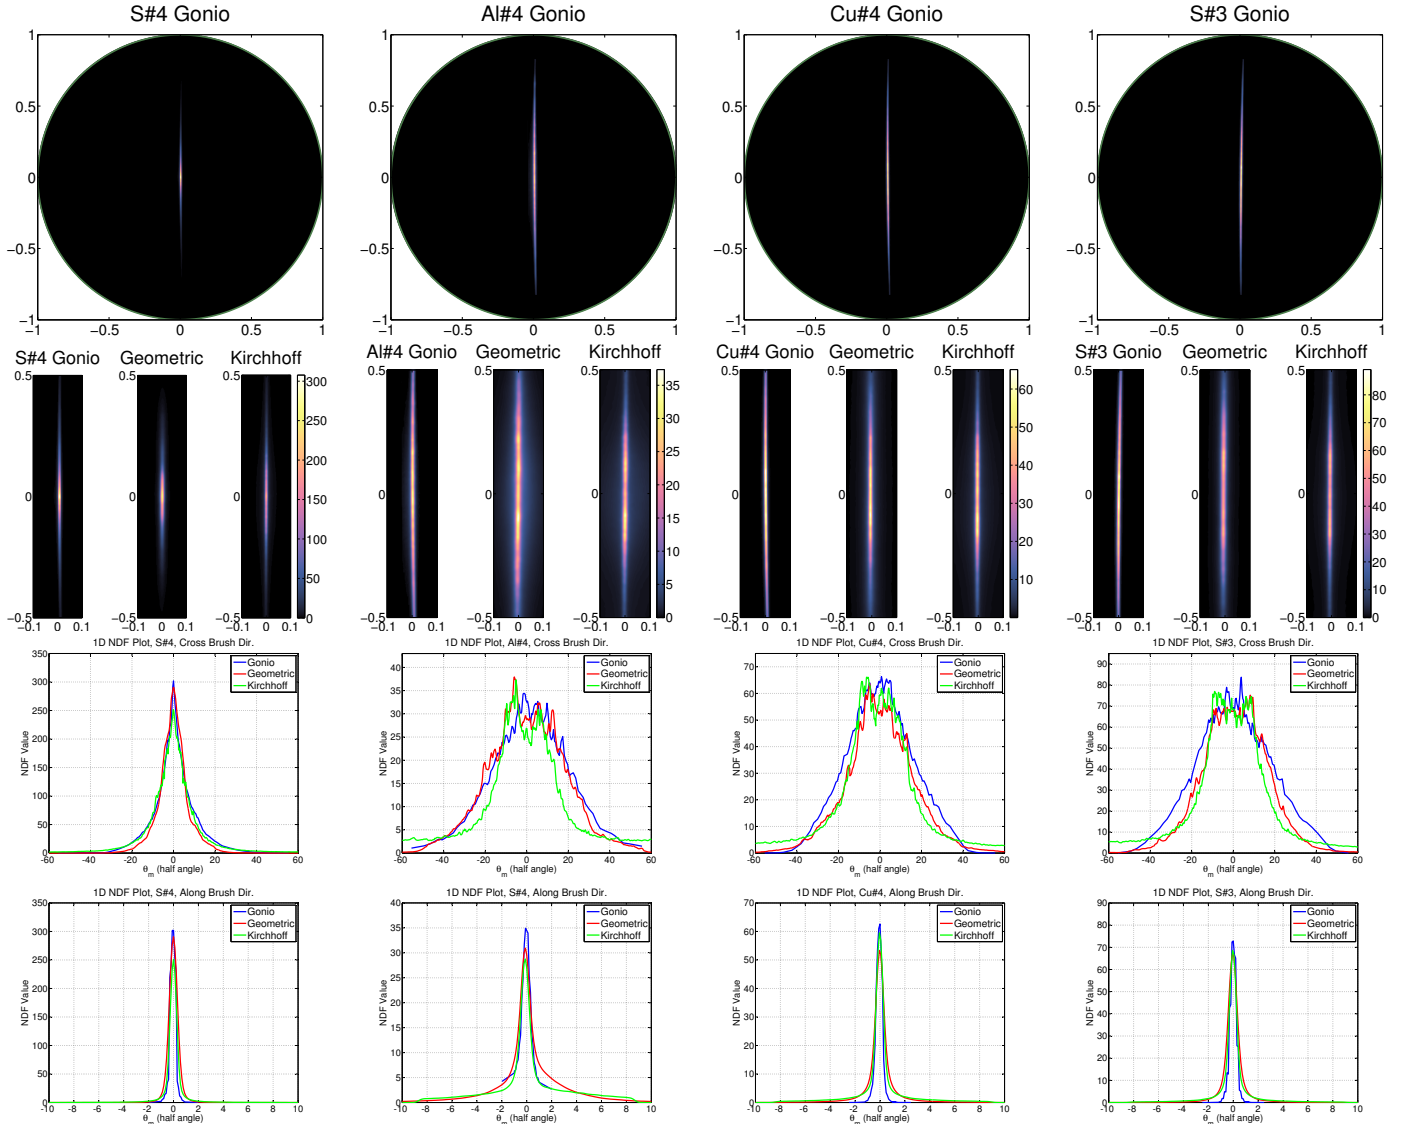


Fig. 9. Comparison of measured and predicted NDF estimates for brushed metal samples. For each sample, the 1st row shows the NDF inferred from gonioreflectometer measurements. The 2nd row includes the 2D zoom-in plots of the NDF from the gonioreflectometer (left) along with those predicted by our filtered normal method (middle), and by Kirchhoff integral (right). We further show the 1D NDF plots of two 1D perpendicular slices through the 2D NDFs: cross brush direction (3rd row) and along brush direction (4th row).

matrix is the identity (i.e. $\mathbf{R} = \mathbf{I}$ or no rotation), then the ellipsoid distribution reduces to be exactly the same as the anisotropic distribution GTR2aniso in [Burley 2012][Eq. 13]. And if we further set $\alpha_x = \alpha_y = \alpha$, then the distribution becomes identical to the GGX/Trowbridge-Reitz distribution. As in those previous models, the α parameters control the width of the distribution in two orthogonal directions and correspond to notions of surface roughness.

One convenient way to specify the rotation matrix is as the product of three axis-aligned rotations:

$$\mathbf{R} = \mathbf{R}_x(\theta_x) \mathbf{R}_y(\theta_y) \mathbf{R}_z(\theta_z) \quad (12)$$

This parameterization provides nicely intuitive controls when θ_x and θ_y are small, which is true for our data. The θ parameters in this space each have a simple meaning. θ_z rotates the major axis

of the BRDF (e.g., rotates the brushing direction for brushed metals), while θ_x and θ_y rotate the peak of the distribution away the direction of the average surface normal. This shifting of the NDF peak away from \mathbf{n} is something that we observe when we estimate NDFs for small areas and is important to be able to represent to capture the spatial variation over our samples. This ability to skew the distribution is not supported by any of the prior parametric NDF models, such as GGX and Beckmann, and is why we developed the ellipsoid NDF.

We recommend the following energy-conserving shadowing masking term for use with the ellipsoid NDF:

$$G(\psi, \omega) = G_1(\psi) G_1(\omega) \quad (13)$$




Fig. 10. Images rendered using our profilometer-derived geometric NDFs. Each image shows sections of two cylinders, where the BRDF has been rotated 90 degrees on the second cylinder. For the anisotropic samples, the brushing direction is horizontal on the left cylinder and vertical on the right. The cylinders use spatially uniform BRDFs with any apparent texturing caused by angular variations within the BRDFs. Copper’s color is due to its Fresnel term.

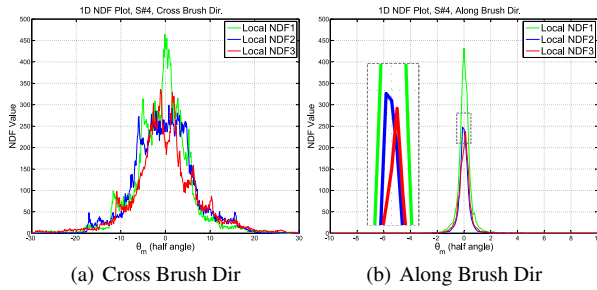


Fig. 11. Example of the NDF variations for different sub-regions in stainless steel #4 sample.

$$G_1(\mathbf{u}) = \min\left(1, \frac{2 \|\mathbf{An}\|^2 |\mathbf{u} \cdot \mathbf{n}|}{\|\mathbf{Au}\| \|\mathbf{An}\| + (\mathbf{Au}) \cdot (\mathbf{An})}\right) \quad (14)$$


A supplemental document provides more details including derivations of the NDF, this shadowing-masking term, and efficient sampling strategies for the ellipsoid NDF.

Figure 12 shows data from fitting the Ellipsoid and Beckmann distributions to the large area tabulated NDFs for our metal samples. For the large area NDFs, the ellipsoid θ parameters are all approximately zero, making it equivalent to the GTR2aniso distribution. The ellipsoid NDF generally provides a better fit to our measured NDFs than the Beckmann distribution, such as in the example shown in the figure.

9. SPATIAL VARIATION

So far we have worked primarily with large-area average NDFs, in order to have a meaningful comparison between the gonioflectometer and profilometer results. However, for rendering, the spatial variation of the BRDF over the metallic surface is important for the overall appearance of all our samples, especially the brushed finishes. The profilometer data contains extraordinarily detailed information about spatial variation over the small areas we measured, and since the spatial structure is essentially random and stationary we take the approach of using a Gaussian random field model to generate spatially varying NDFs that resemble the spatial variation observed in the measured areas.

Our basic strategy is to fit the ellipsoid NDF model to small-area NDFs derived from subsets of the profilometer data corresponding to texels at the desired texture resolution. This produces a 5-vector of ellipsoid NDF parameters at each texel over two very small texture images. We assume that the variation in these parameters



Sample	Ellipsoid (α_x, α_y)	Beckmann (α_x, α_y)
QPanel	(0.0823, 0.0967)	(0.0893, 0.1028)
Stainless #4	(0.0074, 0.142)	(0.008, 0.16)
Stainless #3	(0.0173, 0.39)	(0.0181, 0.418)
Aluminum #4	(0.0208, 0.469)	(0.0265, 0.497)
Copper #4	(0.0183, 0.414)	(0.0209, 0.459)

Fig. 12. Comparison of stainless steel #4 large area tabulated NDF with fits to the Ellipsoid and Beckmann parametric NDFs, and fitted parameters for all our samples. The Ellipsoid NDF usually provides a better fit to our measured data than Beckmann as in the Stainless #4 example (top).

comes from a Gaussian random field with a separable Fourier amplitude spectrum, calculate a spectrum that fits the data, and then generate a large area of 5-channel texture from our random field using the inverse Fourier transform method. The resulting texture is then used to define a spatially varying BRDF via the microfacet model and the ellipsoid NDF.

We compute our textures from two sets of profilometer images: one is 1×100 (a $53\mu\text{m} \times 7\text{mm}$ area aligned across the scratches) and the second is 100×1 (a $5.3\text{mm} \times 70\mu\text{m}$ area aligned with the scratches). We predict NDFs from the surface data in each block, producing two samples of the texture we want: 100×1 and 1×100 5-channel parameter textures sampled at ~ 400 texels per inch.

Synthesizing the texture from these exemplars is simple. We calculate the 1D Fourier spectrum (retaining only the amplitude) of each 1D texture and compute their outer product; under the separability assumption this gives the 2D Fourier spectrum of a $5.3 \times 7 \text{ mm}^2$ area. Assuming this is a large enough area that the same spectrum would be observed anywhere on the surface, we simply up-sample to the size of the desired texture, randomize the phase, and use the inverse FFT to obtain the parameter texture. The whole process is carried out independently per parameter, corresponding to an assumption that the 5 parameters vary independently.

To test the feasibility and visual quality of this simple texture synthesis procedure, in Figure 13 we compare a photograph of our stainless steel #4 sample to the images rendered with matched light-

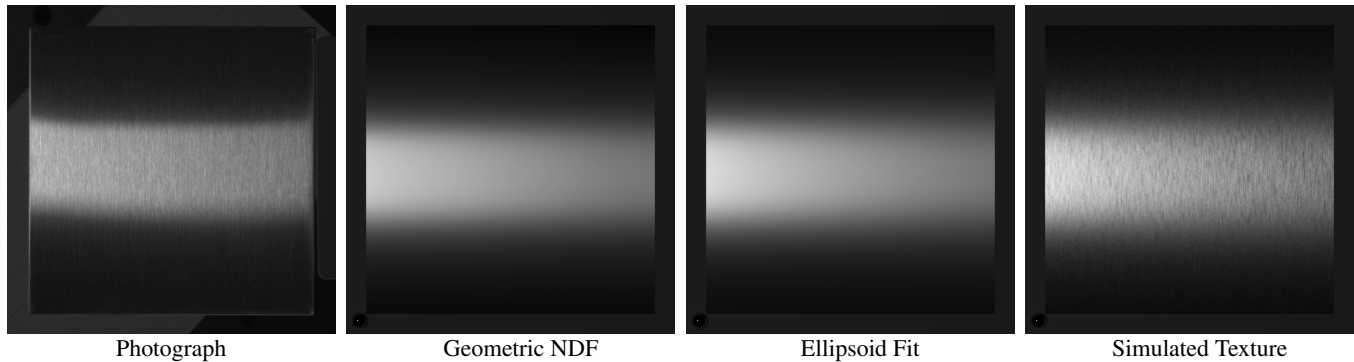


Fig. 13. Comparison of a photograph of our stainless steel #4 samples and three images rendered under matched lighting using profilometer-based predictions of its BRDF. The geometric NDF image uses the average BRDF estimated in Section 7.1 without any spatial variation. The next image is rendered using the ellipsoid fit to the average NDF. The last image is rendered using our spatial variation, or texture synthesis method. While not a perfect match to the real texture, using the synthetic texture significantly improves the visual realism of the renderings.

ing. The sample is roughly 10cm square, the camera is directly above it, and the light source spans roughly 9 degrees at a distance of 64cm and is rotated 23 degrees from the camera (toward the left). To avoid problems with camera color demosaicing, we use grayscale images. The three rendered images shown in the figure use the average NDF (geometric method) from Section 7.1, the ellipsoid fit to the average NDF, and the synthesized texture, respectively. The highly anisotropic nature of the BRDF, creates a bright stripe highlight on the surface. The geometric average NDF does a good job of reproducing the general location and shape of this highlight but obviously lacks the finer scale details visible within the lit stripe. The rendered stripe is also slightly too wide, likely because we have somewhat overestimated the narrow width of the NDF lobe. The ellipsoid fit to the average NDF is nearly identical to it. Using the synthesized texture adds detail within the highlight and significantly improves its visual realism. While generally similar, the texture exhibits a less structured appearance than the photograph and contains some other kinds of artifacts. We suspect this is because the texture is based on measurements span only ~6mm and a tiny fraction of the sample's area which limits its ability to reproduce larger structures. This could likely be improved by measuring larger areas with the profilometer.

10. CONCLUSION

In this paper we have presented an alternative approach for reproducing real world material appearance based on measuring its surface microgeometry and using that data to predict its BRDF. Our approach leverages the increasing speed and resolution of surface profilometer tools from the photolithography and nanofabrication industries, and has many advantages over traditional BRDF capture strategies. It handles highly anisotropic surfaces and is capable of rapid BRDF acquisition with both high spatial and angular resolution, albeit with limited spatial extent. We explored both microfacet and Kirchhoff-based theoretical approaches to predicting BRDFs from surface detail, developed a new geometric NDF estimation scheme, and validated the results against gonireflectometer measurements. Our results demonstrate that both the geometric and Kirchhoff-based approaches can successfully predict the appearance of highly anisotropic surfaces such as our brushed metal examples. We also presented the new ellipsoid NDF model, which supports both anisotropic and asymmetric features, and have shown how we can use it to model spatial variation for our sample materials in a compact format that is convenient for rendering. We be-

lieve our method has demonstrated state of the art results for our test metal samples and opens a new avenue for future BRDF acquisition methods based on microgeometry measurement.

Limitations and Future Work. Our methods only predict the first surface reflection for the surfaces and thus is currently only suitable for materials where this is the dominant effect, such as metal surfaces. Materials with important subsurface scattering, such as paint, or more complex geometry, such as hair, are beyond the scope of this paper. Although only knowing the surface is not sufficient to predict the appearance of such materials, nearly all materials contain important surfaces. In future we would like to extend our method to handle surface transmission and multiple scattering so that it can be used as a component in appearance modeling for a broader class of materials. We are also limited in the total spatial extent of the regions we can measure at such high resolution, and our approach is not well suited for capturing large scale spatial patterns. Commercial profilometers are continuing to improve in speed, so in future may be feasible to capture much larger regions to get better spatial statistics. Combining our technique with camera-based acquisition for larger scales, is also an interesting avenue for future work.

Acknowledgements

We would like to thank the Cornell NanoScale Science & Technology Facility (CNF) for use of their profilometer. Thanks to Hurf Sheldon for the technical support during experiments. This work was generously supported by funding from Autodesk.

REFERENCES

- ASHIKHMIN, M., PREMOŽE, S., AND SHIRLEY, P. 2000. A microfacet-based brdf generator. In *Proc. of SIGGRAPH '00*. 65–74.
- ASHIKHMIN, M. AND SHIRLEY, P. 2000. An anisotropic phong brdf model. *J. Graph. Tools* 5, 2 (Feb.), 25–32.
- BECKMANN, P. AND SPIZZICHINO, A. 1968. *The Scattering of Electromagnetic Waves from Rough Surfaces*. Artech House Radar Library. Books on Demand.
- BLINN, J. F. 1977. Models of light reflection for computer synthesized pictures. In *Computer Graphics (Proceedings of SIGGRAPH 77)*. 192–198.
- BURLEY, B. 2012. Physically-based shading at disney. In *ACM SIGGRAPH 2012 Course: Practical Physically-based Shading in Film and Game Production*. SIGGRAPH '12.

- COOK, R. L. AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Trans. Graph.* 1, 1 (Jan.), 7–24.
- CUYPERS, T., HABER, T., BEKAERT, P., OH, S. B., AND RASKAR, R. 2012. Reflectance model for diffraction. *ACM Trans. Graph.* 31, 5 (Sept.), 122:1–122:11.
- DANA, K. 2001. Brdf/btf measurement device. In *Proc. of ICCV*. Vol. 2. 460–466.
- DANA, K. J., VAN GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and texture of real-world surfaces. *ACM Trans. Graph.* 18, 1–34.
- DHILLON, D., TEYSSIER, J., SINGLE, M., GAPONENKO, I., MILINKOVITCH, M., AND ZWICKER, M. 2014. Interactive diffraction from biological nanostructures. *Computer Graphics Forum* 33, 8, 177–188.
- DONG, Y., WANG, J., TONG, X., SNYDER, J., LAN, Y., BEN-EZRA, M., AND GUO, B. 2010. Manifold bootstrapping for svbrdf capture. *ACM Trans. Graph.* 29, 98:1–98:10.
- GARDNER, A., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2003. Linear light source reflectometry. *ACM Trans. Graph.* 22, 3 (July), 749–758.
- GARG, G., TALVALA, E.-V., LEVOY, M., AND LENSCH, H. P. A. 2006. Symmetric photography: Exploiting data-sparseness in reflectance fields. In *Rendering Techniques*. 251–262.
- GOLDMAN, D., CURLESS, B., HERTZMANN, A., AND SEITZ, S. 2005. Shape and spatially-varying brdfs from photometric stereo. In *Proc. of ICCV*. Vol. 1. 341–348.
- GU, J., TU, C.-I., RAMAMOORTHI, R., BELHUMEUR, P., MATUSIK, W., AND NAYAR, S. 2006. Time-varying surface appearance: acquisition, modeling and rendering. *ACM Trans. Graph.* 25, 3 (July), 762–771.
- HAN, J. Y. AND PERLIN, K. 2003. Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Trans. Graph.* 22, 3 (July), 741–748.
- HE, X. D., TORRANCE, K. E., SILLION, F. X., AND GREENBERG, D. P. 1991. A comprehensive physical model for light reflection. *SIGGRAPH Comput. Graph.* 25, 4 (July), 175–186.
- KAJIYA, J. T. 1985. Anisotropic reflection models. *SIGGRAPH Comput. Graph.* 19, 3 (July), 15–21.
- LAFORTUNE, E. P. F., FOO, S.-C., TORRANCE, K. E., AND GREENBERG, D. P. 1997. Non-linear approximation of reflectance functions. In *Proc. of SIGGRAPH*. 117–126.
- LENSCH, H. P. A., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2003. Image-based reconstruction of spatial appearance and geometric detail. *ACM Trans. Graph.* 22, 2 (Apr.), 234–257.
- LEVIN, A., GLASNER, D., XIONG, Y., DURAND, F., FREEMAN, W., MATUSIK, W., AND ZICKLER, T. 2013. Fabricating brdfs at high spatial resolution using wave optics. *ACM Trans. Graph.* 32, 4 (July), 144:1–144:14.
- LI, H., FOO, S. C., TORRANCE, K. E., AND WESTIN, S. H. 2005. Automated three-axis gonioreflectometer for computer graphics applications.
- LI, H. AND TORRANCE, K. E. 2005. An experimental study of the correlation between surface roughness and light scattering for rough metallic surfaces.
- LÖW, J., KRONANDER, J., YNNERMAN, A., AND UNGER, J. 2012. Brdf models for accurate and efficient rendering of glossy surfaces. *ACM Trans. Graph.* 31, 1 (Feb.), 9:1–9:14.
- MANDEL, L. AND WOLF, E. 1995. *Optical Coherence and Quantum Optics*. Cambridge University Press.
- MARX, E. AND VORBURGER, T. V. 1990. Direct and inverse problems for light scattered by rough surfaces. *Appl. Opt.* 29, 25 (Sep), 3613–3626.
- MASHAAL, H., GOLDSTEIN, A., FEUERMANN, D., AND GORDON, J. M. 2012. First direct measurement of the spatial coherence of sunlight. *Opt. Lett.* 37, 17 (Sep), 3516–3518.
- MCKNIGHT, M. E., VORBURGER, T. V., MARX, E., NADAL, M. E., BARNES, P. Y., AND GALLER, M. A. 2001. Measurements and predictions of light scattering by clear coatings. *Appl. Opt.* 40, 13 (May), 2159–2168.
- MÜLLER, G., MESETH, J., SATTLER, M., SARLETTE, R., AND KLEIN, R. 2005. Acquisition, synthesis, and rendering of bidirectional texture functions. *Comput. Graph. Forum* 24, 1, 83–109.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2005a. Experimental analysis of brdf models. In *Rendering Techniques 2005: 16th Eurographics Workshop on Rendering*. 117–126.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2005b. Experimental analysis of brdf models. In *Proc. of EGSR*. 117–126.
- OGILVY, J. 1991. *Theory of wave scattering from random rough surfaces*. A. Hilger.
- SCHRÖDER, S., DUPARRÉ, A., CORIAND, L., TÜNNERMANN, A., PENALVER, D. H., AND HARVEY, J. E. 2011. Modeling of light scattering in different regimes of surface roughness. *Optics Express* 19, 9820.
- STAM, J. 1999. Diffraction shaders. In *Proc. of SIGGRAPH*. 101–110.
- SUNG, L.-P., NADAL, M., MCKNIGHT, M., MARX, E., AND LAURENTI, B. 2002. Optical reflectance of metallic coatings: Effect of aluminum flake orientation. *Journal of Coatings Technology* 74, 932, 55–63.
- TORRANCE, K. E. AND SPARROW, E. M. 1967. Theory for off-specular reflection from roughened surfaces. *J. Opt. Soc. Am. A* 9, 1105–1114.
- TROWBRIDGE, T. S. AND REITZ, K. P. 1975. Average irregularity representation of a rough surface for ray reflection. *J. Opt. Soc. Am.* 65, 5 (May), 531–536.
- WALTER, B., MARSCHNER, S. R., LI, H., AND TORRANCE, K. E. 2007. Microfacet models for refraction through rough surfaces. In *Proc. of EGSR*. 195–206.
- WANG, J., ZHAO, S., TONG, X., SNYDER, J., AND GUO, B. 2008. Modeling anisotropic surface reflectance with example-based microfacet synthesis. *ACM Trans. Graph.* 27, 3 (Aug.), 41:1–41:9.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection. *SIGGRAPH Comput. Graph.* 26, 2 (July), 265–272.
- YAN, L.-Q., HAŠAN, M., JAKOB, W., LAWRENCE, J., MARSCHNER, S., AND RAMAMOORTHI, R. 2014. Rendering glints on high-resolution normal-mapped specular surfaces. *ACM Trans. Graph.* 33, 4 (July), 116:1–116:9.
- ZHAO, S., JAKOB, W., MARSCHNER, S., AND BALA, K. 2011. Building volumetric appearance models of fabric using micro ct imaging. *ACM Trans. Graph.* 30, 4 (July), 44:1–44:10.
- ZYGO®. 2011. Newview™ 7300 3d optical surface profiler. <http://www.zygo.com/?/met/profilers/newview7000/>.

APPENDIX

A. COMPUTING KIRCHHOFF DISTRIBUTION, D_K

In section 6 we described how Kirchhoff theory can be used to derive a BRDF model (equation 8) that is very similar in form to microfacet BRDF models (equation 1). In this context, the Kirchhoff distribution D_K (equation 6) can be viewed as an alternate way to compute an effective NDF. The definition of D_K already includes a notion of wavelength scale, so there is no need to add a filtering scale parameter as we did for the geometric-based NDF estimation. However there are some practical issues related to its computation.

Kirchhoff scattering models the incident light as a coherent plane wave, but the light from real sources is more complicated. With a plane wave, the phase of the incident light remains fully correlated

regardless of distance between receiving points. In general, however, phase correlation decreases with separation distance, becoming incoherent for sufficiently distant receiver points. The region size over which the phase remains correlated is called the coherence area. For a simple uniform source the coherence area is approximately given by: $A_c \approx \lambda^2/\Omega_\ell$, where Ω_ℓ is the solid angle subtended by light source in steradians [Mandel and Wolf 1995]. Scattered light under coherent illumination exhibits interference effects (both constructive and destructive), which are modeled by the Kirchhoff integral, that are not seen with incoherent light. Limiting the surface integral to regions matching the size of the coherence area is one way to account for this difference (e.g., [Levin et al. 2013]). Unfortunately the actual coherence area depends on details of the lighting configuration and usually is not known in advance.

Experimentally we observe that the principal effect of limiting the coherence area is a blurring of the effective NDF. This makes sense, as the coherence area is closely related to the light source solid angle which similarly limits our ability to observe fine details in the BRDF and NDF. In practice, it suffices to choose a coherence area large enough to preserve the major features in D_K or to encompass likely lighting configurations. Our profilometer patch size ($70 \times 53 \mu\text{m}$) is just large enough to avoid broadening our narrowest NDF (stainless steel #4) and also matches the measured coherence area for sunlight ($\sim 10^4 \lambda^2$ [Mashaal et al. 2012]). Thus in this paper, we compute D_K for each profilometer patch individually and then average its value over all the patches in a dataset.

Equation 6 is a complex, highly oscillatory integral that must be recomputed for each value of $\vec{\mathbf{q}}$. We triangulate the profilometer height data to create a surface, and estimate the integral for each triangle, and sum over the patch. We tried different triangulations as well as both analytic and Monte Carlo solvers, and found they all produced essentially identical results. In our implementation, a Monte Carlo solver with a simple triangulation is the fastest solver and is used here. To compute each NDF, we fix λ_q and then evaluate the surface integral for roughly 10^5 values of \mathbf{h} to sufficiently resolve the NDF in the region where it is significantly non-zero. The result is then resampled and stored in the same discrete 1025×1025 format as for the geometric NDFs.