

Lightweight Estimation of Surface BRDFs

James A. Ferwerda

Chester F. Carlson Center for Imaging Science, Rochester Institute of Technology; Rochester, NY, USA
E-mail: jaf@cis.rit.edu

Abstract. *In this article, we introduce a method for using a smartphone to estimate the bi-directional reflectance distribution functions of real-world surfaces. Specifically, we use a combination of image-based and visual techniques to determine a surface's diffuse and specular reflectance parameters in the Ward light reflection model. We test the accuracy of the method by comparing our estimates to instrumental measurements of color and gloss standards. We demonstrate its utility by rendering synthetic images of method-estimated surfaces and comparing the renderings with photographs of the real surfaces. The method provides a lightweight approach to surface reflectance measurement. © 2018 Society for Imaging Science and Technology.*
[DOI: 10.2352/J.ImagingSci.Technol.2018.62.5.050407]

1. INTRODUCTION

Physically based computer graphics has enabled great advances in design, engineering, architecture, and entertainment. Physically based rendering algorithms can accurately simulate the appearances of real-world surfaces, but gathering data on surface reflectance properties is complicated, and typically involves specialized instruments such as goniospectrophotometers.

In this article, we introduce a method for using a smartphone to measure the reflectance properties of real-world surfaces so they can be incorporated into computer graphics simulations. Specifically we use a combination of image-based and visual measurement to estimate the parameters of the popular Ward bi-directional reflectance distribution function (BRDF) model [1]. We estimate a surface's diffuse reflectance ρ_d by linearizing and chromatically adapting a well-exposed and white-balanced image of the surface. We estimate a surface's specular reflectance ρ_s by applying the Fresnel equation for dielectric materials. Finally, we estimate a surface's micro-scale roughness α by displaying a high-contrast square-wave grating on the smartphone screen, reflecting that image in the surface and adjusting the spatial frequency of the grating until it is barely visible in the reflection. We then use this measurement to estimate α .

We test the method's accuracy by comparing physical reflectance measurements of the X-Rite ColorChecker chart [2] and NCS Gloss Standards [3] with estimates obtained by the method. Finally we demonstrate the utility and fidelity of the method by comparing rendered images of surfaces modeled using the method with real photographs of the same surfaces.

The impact of this work is that it provides a lightweight method for estimating the reflectance properties of real surfaces so they can be incorporated into computer graphics simulations in support of a wide range of design, engineering, and cultural applications.

2. BACKGROUND

Opaque surfaces absorb and reflect the light incident upon them. The spectral properties of absorption and reflection are associated with surface color. The directional properties of reflection are associated with surface gloss, and range from matte or diffuse, where light incident from a given direction is reflected equally in all directions, to glossy or specular, where to different concentrations, light is reflected in a direction equal and opposite to the incident direction.

Physical characterization of surface reflectance properties is embodied in the BRDF [4]. The BRDF is a four-dimensional function of wavelength that relates the light incident on a surface within a directional solid angle to the light reflected from the surface through another directional solid angle.


The goniospectrophotometer is an instrument designed to measure BRDFs. It consists of a light source that illuminates a surface from a given direction, a detector that measures the light reflected from the surface from another direction, and a stage that holds the surface being measured. Typically, these components are movable, allowing the orientations of the source, surface, and detector to be varied to sample the multi-dimensional BRDF.

One domain where BRDF data are used is in physically based rendering, a branch of computer graphics that seeks to generate synthetic images that are indistinguishable from images of the real world by simulating the physics of light transport in modeled scenes. Although it is possible to render images of surfaces directly from BRDF data, for reasons of efficiency and generality, BRDF data sets are typically approximated with BRDF models [5]. A wide variety of models have been developed, each with particular strengths and limitations, but most separate the BRDF into diffuse and specular components and model the magnitudes and distributions of these components using parameterized functions.

The Ward BRDF model is a highly regarded and popular light reflection model. Figure 1 shows the model, where the four-dimensional BRDF is approximated by a sum of functions scaled by three parameters, ρ_d the magnitude of the diffuse component, ρ_s the magnitude of the specular

Received Apr. 14, 2018; accepted for publication Aug. 29, 2018; published online Oct. 4, 2018. Associate Editor: Yaohua Xie.

1062-3701/2018/62(5)/050407/8/\$25.00



$$\rho(\theta_i, \phi_i, \theta_o, \phi_o) = \frac{\rho_d}{\pi} + \rho_s \frac{\exp[-\tan^2 \delta / \alpha^2]}{4\pi\alpha^2 \sqrt{\cos\theta_i \cos\theta_o}}$$

Figure 1. The Ward BRDF model.

component, and α , a term that describes the width of the specular lobe. The model parameters can be varied on a wavelength basis, but in practice dielectric surfaces are described using a ρ_d triple in RGB to represent diffuse reflectance, and single ρ_s and α values to describe the specular lobe. For metallic surfaces ρ_s is also an RGB triple to reflect the fact that the specular component of metals is spectrally selective.

A perennial problem in physically based rendering is how to set the parameters of a BRDF model to accurately simulate the appearance of a desired real-world surface. If BRDF data is available, the parameters can be directly calculated by sampling and integration, and while BRDF databases are becoming more widely available (e.g. [6]), the number of materials included is still small and the computations are complex. Recognizing this limitation, a number of graphics researchers have developed alternative tools and methods for surface measurement and modeling [7], and while these innovations are promising, they still typically require significant expertise in device fabrication and data analysis.


A lightweight method for surface reflectance measurement was described by Wang et al. [8]. They used a digital camera to capture the reflected image of a high-contrast edge, and then analyzed that image to estimate both micro-scale specular reflectance properties as well as meso-scale surface texture. While this work provided an inspiration for the current project, in the sense of using a smartphone as a measurement device, this approach still requires significant device calibration and data analysis, and only estimates the specular parameters of the light reflection model.

3. DEVELOPING THE METHOD


The goal of our method is to characterize the BRDFs of real-world surfaces by estimating the Ward model parameters of the surfaces using a smartphone. To achieve this we use a combination of image-based and visual measurement techniques. The following sections describe how we estimate each of the parameters.



(a)



(b)



(c)

Figure 2. Estimating diffuse reflectance ρ_d : (a) high and low gloss paint samples used to test the method, (b) original smartphone image, (c) linearized and chromatically adapted image with estimated ρ_d 's.

3.1 Estimating Diffuse Reflectance (ρ_d)

The diffuse reflectance of a surface describes the fraction of incident light that is reflected uniformly over the hemisphere of outgoing directions. In the Ward model, diffuse reflectance is represented by ρ_d , a parameter that varies between 0 and 1. In the model, ρ_d can vary on a wavelength basis, to represent spectral variation in diffuse reflectance. In practice, a surface's diffuse reflectance is usually represented by a triple of ρ_d values relative to a set of RGB primaries.

To estimate a surface's RGB ρ_d values we take advantage of the fact that most smartphone cameras encode their sensor responses with respect to the sRGB standard [9]. In sRGB, image values represent sensor responses coded with respect to a defined set of primaries, a D65 whitepoint, and system gamma of 2.2. We invert this encoding to estimate a surface's ρ_d values relative to the sRGB primaries.

Figure 2 shows the method applied to a set of high and low gloss paint samples. Starting from a well-exposed, white-balanced, off-specular, sRGB-encoded image of a surface, we first linearize the image by inverting the 2.2 gamma correction factor. This yields RGB values that are approximately linear with respect to surface luminance. We then apply the Bradford chromatic adaptation transform [10] to change the values' whitepoint from D65 to Commission Internationale de l'Eclairage (CIE) illuminant E. This equalizes the channels and yields pixel values that approximate the surface's reflectance factors in RGB. Finally we average the pixel values within an image region to produce the final estimates of the RGB ρ_d values.




Figure 3. Testing the accuracy of the diffuse reflectance estimation method. Insets show the ΔE_{ab} color differences with respect to ground truth.

3.2 Testing the ρ_d Method

To test the accuracy of the method we used an iPhone5s to image the X-Rite ColorChecker chart, and attempted to estimate the diffuse reflectance properties of the patches. The chart was placed on a white piece of paper on the floor of an atrium and photographed from above. The illumination was indirect northern daylight. Imaging parameters were determined by the camera’s auto-focus, auto-exposure, and auto-whitebalance algorithms. The captured image was downloaded as a high-quality sRGB JPEG.

We used the method described above to estimate the patch reflectances. To evaluate the accuracy of these estimates we then converted these values to D50 referenced XYZs and CIELabs and then calculated ΔE_{ab} color differences with respect to CIELabs calculated from reflectance spectra provided with the X-Rite ColorChecker chart.

The results are shown in Figure 3. The min, max, and mean ΔE_{ab} values are 1.6, 18.8, and 7.1, respectively. While most of these color differences are above threshold, with the exception of the outlier values for the dark blue patches, most differences are below 10.0, which is remarkable given that no attempt was made to characterize or calibrate the camera. Thus while the method produces estimates of ρ_d that are certainly acceptable, performance could easily be improved with some simple camera modeling.

3.3 Estimating Specular Reflectance (ρ_s)

The specular reflectance of a surface describes the fraction of light incident on the surface from a given direction that is reflected in the specular direction. In the Ward model specular reflectance is represented by ρ_s a parameter that varies between 0 and 1 and indicates the magnitude of the specular lobe. In the model, ρ_s can vary on a wavelength basis to represent spectral variation in the specular reflectance. However most dielectric surfaces are not selective, so a single ρ_s value suffices to describe the surface’s specular reflectance.

To estimate ρ_s we take advantage of the fact that a surface’s specular reflectance can be calculated directly from the Fresnel equation shown in Figure 4.




Figure 4. The Fresnel equation for specular reflectance.



Figure 5. Effects of variations in micro-scale surface roughness on a reflected image. Note the changes in the contrast and sharpness of the reflection as the roughness increases (and gloss declines) from left to right.

Fig. 4 also plots values of the equation as a function of incident angle. The two curves are for metals and dielectrics. The curves show that for incident angles below 60° the specular reflectance of typical dielectric surfaces is roughly constant at 5% but increases dramatically at higher (grazing) angles. The functions that comprise the Ward model approximate this increase in specular reflectance at grazing so using the decimal value of 0.05 for ρ_s should provide a good approximation of the magnitude of specular reflection for most dielectric surfaces.

3.4 Estimating Surface Roughness (α)

In addition to its magnitude, the specular component of surface reflectance has a distribution that is affected by the micro-scale roughness of a surface. It is this roughness and its effects on the scattering of specularly reflected light that account for the range of glossy surface appearances, from near-matte to mirror-like. In the Ward model surface roughness is described by the parameter α , which represents the standard deviation (SD) of the Gaussian distribution used to model the specular lobe.

To estimate α we leverage the insight that glossy surfaces behave like imperfect mirrors, reflecting the patterns of light incident upon them, but scrambling those patterns to different degrees. This effect can be seen in Figure 5, where a high-contrast square-wave grating image is reflected in a paint sample card showing gloss levels from “high gloss” through “eggshell.” Note how the contrast and sharpness of the reflected image drops as the finish becomes less glossy.

We take advantage of this property to estimate α by generating grating patterns on the smartphone screen, reflecting those patterns in the surface under consideration, and having the observer adjust the spatial frequency of the grating until its visibility is at threshold. We then use the grating frequency and knowledge of the human contrast sensitivity function [11] in a model that determines the SD of

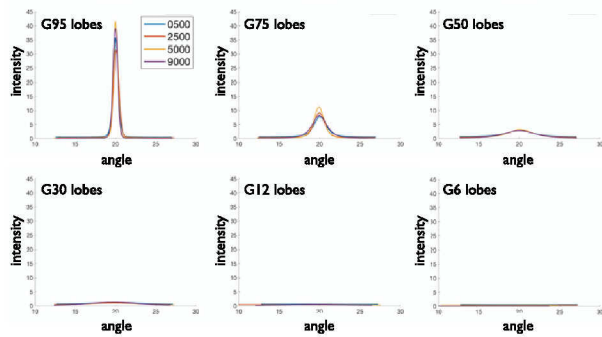


Figure 6. Specular lobes of the NCS gloss standards.




Figure 7. (a) Sample grating image and (b) physical setup used for the psychophysical measurements.

a Gaussian low-pass filter that would produce the same effect on the input pattern and then rescales this value to estimate α . The following sections describe how we developed the model.

4. PHYSICAL MEASUREMENTS

We first needed to relate a physical measure of surface roughness to the surface’s visual appearance. We started by measuring the reflectance properties of the Natural Color System (NCS) gloss standards [3]. The NCS gloss standards are a set of 28 paint samples that have known reflectance properties. The set is divided into four groups that vary in lightness (0500,2500,5000,9000) and gloss (2,6,12,30,50,75,95). The intervals between samples were selected to provide perceptually uniform appearance scales in both dimensions.

We measured these samples using an Elcometer 408 multi-angle glossmeter [12], which has a linear sensor that provides goniometric measurements of a surface’s specular lobe at 20° relative to the surface normal. The graphs in Figure 6 show the lobe distributions at gloss levels ranging from 95 (high gloss) to 6 (low gloss). Each curve in a given graph represents a different surface lightness. We omitted surfaces with gloss level 2 because the measurements showed no variation with angle.

We then calculated the SDs of the lobes as a correlate of surface roughness. Prior to this calculation we subtracted the minimum value in each distribution from every other value in the distribution to minimize the effects of a surface’s diffuse reflectance. Table 1a shows the SD values for each of the measured surfaces. Note first, that the values are

Table 1. Roughness model data, NCS gloss standards.

(a) Measured lobe SDs		NCS gloss number					
		95	75	50	30	12	6
NCS diffuse	White (500)	0.321	0.839	1.915	2.751	7.523	221.721
	lt gray (2500)	0.335	0.698	1.881	2.759	5.640	155.755
	dk gray (5000)	0.278	0.686	1.761	2.761	3.458	5.010
	Black (9000)	0.296	1.024	2.055	2.870	3.437	6.613
(b) Threshold PPBs		NCS gloss number					
		95	75	50	30	12	6
NCS diffuse	White (500)	2.5	6	11	18	50	76
	lt gray (2500)	2	5.5	10	17	45	65
	dk gray (5000)	1.5	5	9	16	25	45
	Black (9000)	1	5	9	15	22	40
(c) Estimated alphas		NCS gloss number					
		95	75	50	30	12	6
NCS diffuse	White (500)	0.003	0.010	0.020	0.033	0.096	0.146
	lt gray (2500)	0.002	0.009	0.018	0.033	0.090	0.131
	dk gray (5000)	0.002	0.009	0.017	0.031	0.049	0.088
	Black (9000)	0.001	0.009	0.017	0.029	0.043	0.080

more or less the same at each gloss level, suggesting relative independence between the diffuse and specular components of reflectance. Also note that the SDs increase systematically with decreasing surface gloss. While these patterns are good news, the shaded portions of the table indicate where limitations in our measurements made us question the accuracy of the statistics. For the dark gray (5000) and black (9000) surfaces, the lobe statistics at the lowest gloss levels (12 and 6) are in line with the patterns and trends seen at higher gloss levels, but the statistics for the light gray and white (2500 and 0500) surfaces are inconsistent with the other estimates. We believe that the SD estimates for these surfaces are in error due to the facts that the lobes are very broad and extend beyond the field of view of the glossmeter’s sensor, and because the variations in lobe intensity from peak to tails are very small and the measurements are probably corrupted by sensor noise and the relatively high levels of diffuse reflection.

5. PSYCHOPHYSICAL MEASUREMENTS

We then conducted a psychophysical study to investigate the effects of surface roughness (and by correlation specular lobe width) on the visibility of reflected gratings. A sample grating and the experimental setup are shown in Figure 7. Images of black/white square-wave gratings were shown on the smartphone screen and reflected in the surface under consideration. The phone was placed against the surface and held at a 45° angle to the surface normal. The surface was also viewed from 45° at a distance of 18 inches, yielding an overall eye-to-screen distance of 19.1 inches. To avoid resampling issues on display, the gratings were scaled in terms of integer pixels per bar (PPBs) and ranged from 1 to 120 PPB, which

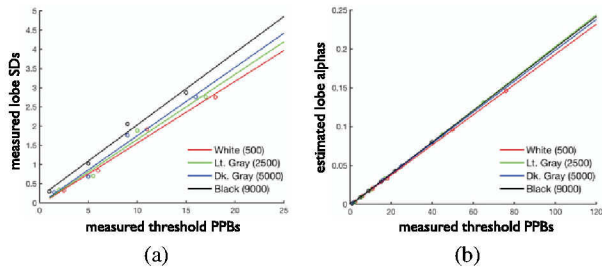


Figure 8. (a) PPB-to-SD and (b) PPB-to- α models.

corresponds to 54.2–0.45 cycles per degree (CPD) at the eye. The maximum and minimum luminances of the display were 505.5 and 0.5 cd/m^2 , respectively, and the effects of light scattering prior to reflection were negligible with measured grating contrasts nominally 1.0 except for the 1 PPB/54.2 CPD grating.

For each of the surfaces in the NCS set, the observers varied the spatial frequency of the grating until its reflected image was at the limit of visibility. They achieved this by holding the phone in both hands and using their thumbs to swipe through a set of grating images that had been loaded into the phone’s iPhoto application. While more sophisticated interfaces can be imagined, this approach was simple to implement and easy to use.

The grating PPBs at threshold are shown in Table 1b. The data represents the averaged responses of three observers ranging from 16 to 58 years in age, all with normal vision.

6. DEVELOPING THE α MODEL

The relationships between grating PPBs at threshold and measured specular lobe SDs are plotted in Figure 8(a). The data for different surface lightnesses are plotted separately. Linear fits to each data set are also shown. Overall the fits are excellent with all R^2 values exceeding 0.96. We interpret the differing slopes of the curves as a small effect of “contrast gloss,” where dark and light surfaces with the same specular reflectance produce reflected images with differing contrast due to the effects of a diffuse “floor” on image contrasts. Although we tried to compensate for this in our calculations of lobe statistics, it persists as a small effect, so we incorporate it in our model.

To estimate Ward model α , we note that the parameter describes the SD of the Gaussian used to represent the specular lobe. As such, estimating α from the PPB-to-SD model simply involves rescaling units. The resulting PPB-to- α model is shown in Fig. 8(b). Note that the x-axis scale (PPBs) has been extended to show the full range of threshold PPBs and their predicted α s. The data points in the 0–20 PPB range are the same as those shown in Fig. 8(a). Above 20 PPB, the data points correspond to the threshold PPBs measured for the low gloss NCS surfaces (12 and 6). It is interesting and reassuring to note that these data points are consistent with the linear PPB-to- α model even though they were not used to develop the model due to the aforementioned concerns about issues with measuring the specular lobes of these surfaces.

7. TESTING THE α MODEL

To test the accuracy of the model we compared real and rendered images of the NCS materials reflecting a high-contrast, low-frequency grating (contrast ≈ 1.0 , frequency = 109 PPB, 0.5 CPD). The parameters of the NCS materials were estimated using the methods/models described above (estimated ρ_d values were 0.853, 0.485, 0.245, and 0.010 from white to black, ρ_s was set to 0.05 and the model-estimated α values are shown in Table 1c. Figure 9(a) shows the real images, which were taken with a Canon Xsi DSLR and converted from RAW format to uncompressed TIFF images. Camera parameters were selected so they did not limit the contrast and resolution of the reflected gratings. Fig. 9(b) shows the rendered images, which were generated using the radiance rendering system [13], with surface parameters specified by the methods described above and scene properties selected to match the physical environment. The images were rendered as floating point High Dynamic Range (HDR) images and converted to uncompressed TIFFs. As with the photographs, the rendering parameters were set so they did not limit the contrast and resolution of the reflected gratings. Fig. 9(c) shows the results of an analysis that compared the sharpness of the reflected gratings in the real and rendered images using the slanted-edge modulation transfer function (MTF) algorithm [14]. Each graph in the panel shows the spatial frequency spectra of corresponding edges in the real and rendered images. The similarity of the corresponding curves indicates that the α model shows good accuracy with respect to photographic ground truth.

8. USING THE METHOD

Our “method” for estimating surface reflectance properties is to combine the three techniques described above. To demonstrate the utility of this method we used it to estimate the Ward model parameters of the set of high and low gloss red, green, and blue paint samples shown in Fig. 2(a) and then rendered images of those materials using the parameters. To estimate the ρ_d values, we took a well-exposed, white-balanced, off-specular image of the samples with an iPhone5s camera and then linearized and chromatically adapted the image as described above. Figure 10(b) shows the ρ_d ’s estimated by this method for the high and low gloss red samples (first three values in each number set). To estimate ρ_s ’s we used the 0.05 value indicated by the Fresnel equation (fourth value in each set). Finally to estimate α ’s we used the reflected grating method and the grating threshold-to- α model described above. Fig. 10(b) shows the α values for the high and low gloss red samples estimated by the model (last value in each set).

Using the estimated parameters we can now incorporate these materials into computer graphics simulations. To test the fidelity of the method we rendered images of the paint samples from their parameters and compared them with photographs of the real samples.

Fig. 10 compares the real and rendered images of the red sample reflecting the homescreen of the iPhone 5s. Overall the match is quite good in terms of color, contrast and

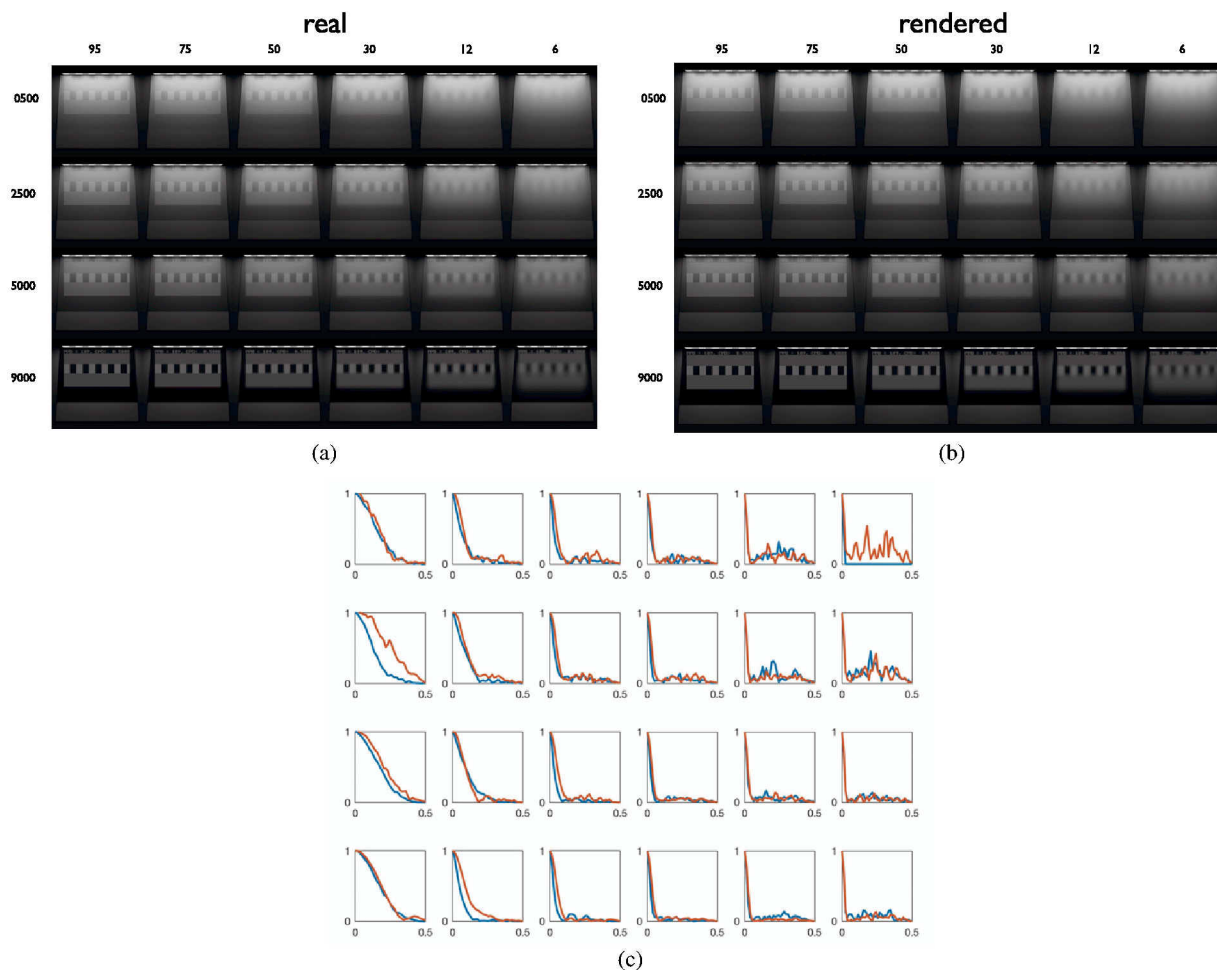


Figure 9. Testing the accuracy of the threshold-to- α model: (a) real images of the NCS standards, (b) renderings of the NCS standards based on the model, (c) spatial frequency spectra of the real and rendered images.

sharpness for both the high and low gloss samples. One significant difference is the sharper drop-off in reflected contrast from top to bottom seen in the real image of the low gloss surface relative to the rendering. While conducting our tests we noticed that light emission from the iPhone screen was quite directional, and we believe this accounts for the difference, because the illumination map used in the rendering emitted in a purely diffuse manner. More accurate modeling of the screen as a light source should reduce these differences.

Figure 11 shows another set of renderings generated using the material parameters estimated by the method. In this case we have applied the materials to spheres and placed them in an environment different than the one they were measured in. While correspondence in color and gloss between the real samples and the rendered objects is good, the experience of doing the visual comparison between the real materials and the renderings is not quite as definite as it was in Fig. 10. This shows the significant impacts that object shape and scene geometry and illumination have on surface appearance. These are important topics that are


outside of the scope of this project, but investigating them could be facilitated by using the method to estimate material properties and then using renderings with varying shape and illumination properties as stimuli in psychophysical experiments.

9. VALIDATION

To evaluate the accuracy of the method, we conducted a series of tests that compared measurements of surface properties made using traditional instrumentation with estimates of surface properties made using the method describe above. The surfaces tested were the low and high gloss red, green, and blue paint samples shown in Fig. 2 and rendered in Figs. 10 and 11.

9.1 Estimating Diffuse Reflectance

To evaluate the accuracy of estimates of diffuse reflectance properties we first measured the XYZ and Lab values of the samples using an X-Rite Colormunki colorimeter. We then used our method to estimate the XYZ and Lab values from images of the samples using the technique illustrated in



(a)



(b)

Figure 10. Using the method: (a) real and (b) rendered images of the high/low gloss red paint sample reflecting the iPhone homescreen. Number sets indicate the estimated RGB ρ_d , and ρ_s and α 's of the high/low gloss surfaces.



(a)

(c)

(e)



(b)

(d)

(f)

Figure 11. High/low gloss red, green, and blue spheres rendered using parameters estimated with the method for the paint samples shown in Fig. 2(a). Number sets indicate the estimated RGB ρ_d , and ρ_s and α 's of the surfaces.

Fig. 2. Since the two data sets were referenced to different illuminants (D50 and D65, respectively), we applied the Bradford chromatic adaptation transform to convert the reference illuminant for both data sets to CIE illuminant E before calculating CIE ΔE_{ab} color differences.

The results are summarized in Table II. Each column represents a given sample. The first and second rows show the measured and estimated Lab values for that sample, and the third row shows the ΔE_{ab} color differences between the two values. To aid comprehension of the magnitudes of these




Figure 12. Measurements and estimates of specular lobe properties, (a) low gloss samples, (b) high gloss samples.

Table II. Measured and estimated diffuse reflectance properties of the test samples.

Lab values	Low gloss samples			High gloss samples		
	Red	Green	Blue	Red	Green	Blue
Measured	51.277	59.723	58.252	49.809	58.911	58.716
	47.724	-32.376	-10.469	51.158	-34.346	-11.419
	15.447	28.201	-44.210	17.814	31.656	-44.530
Estimated	50.450	57.418	55.041	51.687	58.164	56.581
	43.358	-33.870	3.871	44.154	-34.182	1.755
	4.177	14.620	-60.479	7.881	17.905	-59.122
ΔE_{ab}	12.11	13.86	21.92	12.30	13.77	19.77

differences, each of the panels in the table has been colored using sRGB values calculated from the corresponding Labs. Thus the colored table shows that while the numerical color differences are all above threshold, the visual differences are relatively small, and that the estimates are usefully close to the measurements, especially considering the simplicity of the method.

9.2 Estimating Specular Reflectance

To evaluate the accuracy of estimates of specular reflectance properties we measured the specular lobes of each sample using the Elcometer glossmeter described earlier and also estimated the lobe properties using our method. The results are summarized in Figure 12 and Table III.

Fig. 12(a) plots the measured and estimated lobes for the low gloss samples. The first thing to notice is that the measured lobes are all very broad, with only modest peaks. This is reflected in the relatively large lobe SDs for these samples summarized in Table III. The second thing to notice is that the measured lobes peak at approximately 17–18°, showing the off-specular characteristic often found in low gloss surfaces [15]. While this property can be significant physically, it has little impact on appearance and is typically not accounted for by conventional light reflection models such as the Ward model, whose lobes will peak at the specular angle.

Fig. 12(a) and Table III also summarize the method-estimated lobe properties of the low gloss samples. The plots of the estimated lobes appear similar to the measured lobes, with broad spreads and modest peaks; however

Table III. Measured and estimated specular lobe statistics of the test samples.

Lobe SDs	Low gloss samples			High gloss samples		
	Red	Green	Blue	Red	Green	Blue
meas.	13.69	4.23	4.06	2.73	2.49	2.83
est.	10.46	7.95	6.75	2.51	2.09	2.46

the statistics in Table III show that the estimated SD values are substantially different than the measured values, with the estimates for the red, green, and blue samples being off by 31%, 87%, and 66%, respectively, from the measured values. This reflects the general difficulties with accurately measuring/estimating the lobe properties of low gloss samples due to the low signal levels produced by the broad lobes. However the good news is that due to the low signal levels, even large numerical errors in lobe estimates produce only small changes in appearance, which can be seen by comparing the visual similarities of the low gloss samples shown in Figs. 10 and 11 with the numerical differences in their lobe SDs summarized in Table III.

The picture is clearer for the high gloss samples whose lobe properties are summarized in Fig. 12(b) and Table III. Here the shapes of the measured and estimated lobes are quite similar and the lobe statistics are much closer as well, with estimated SDs for the red, green, and blue samples off by 9%, 19% and 15%, respectively, from the measured values. These errors reflect a limitation of the current implementation of the method, where to avoid display aliasing issues, the bars in the test gratings vary in integer pixel increments, but which means that at high grating frequencies, a one pixel change in PPB is equivalent to 8% change in estimated lobe SD. Thus the limiting granularity of the method is relatively coarse for estimating the physical properties of high gloss surfaces. That said, as Figs. 10 and 11 show, the accuracy appears to be sufficient for representing the appearance properties of these surfaces.

In sum, the testing described in this section suggests that while the proposed method is not a replacement for traditional instrumentation when highly accurate measurement of the physical reflectance properties of surfaces is important, it does appear to be sufficiently accurate to represent the appearance properties of surfaces for applications such as computer graphics visualization.

10. CONCLUSION

In this article, we have introduced a lightweight method for estimating surface reflectance properties using a smartphone. The method leverages the smartphone camera and display and uses both image-based and visual measurement techniques to estimate the parameters of the Ward BRDF model for real surfaces. We have described our approaches to

estimating each of the parameters, evaluated the accuracy of the techniques and demonstrated the utility of the method by rendering images of surfaces from the estimated parameters and comparing them with photographs of the real surfaces.

Overall the method appears to be sufficiently accurate to produce renderings that capture the appearances of real surfaces, but there is clearly much room for improvement. First, we have only tested the method on dielectric surfaces. Estimating the properties of metals is an obvious and seemingly achievable next step. Second, many surfaces are textured and/or patterned, and developing methods for measuring spatially varying materials would be a valuable advance. Finally, further testing and validation both in terms of physical accuracy and visual fidelity across a range of materials and scenes would provide more information on the general utility of the method as well as greater knowledge of the effects of object and scene geometry and illumination on surface appearance.

REFERENCES

- G. J. Ward, "Measuring and modeling anisotropic reflection," *Proc. ACM SIGGRAPH '92* (Association for Computing Machinery, Chicago, IL, 1992), Vol. 26, pp. 265–272.
- XRite, <http://xritephoto.com/colorchecker-classic>.
- NCS Colour AB, <http://ncscolour.com/product/ncs-gloss-scale/>.
- United States. National Bureau of Standards, and F. E. Nicodemus. *Geometrical Considerations and Nomenclature for Reflectance*, Vol. 160. (US Dept. of Commerce, National Bureau of Standards, 1977).
- S. H. Westin, H. Li, and K. E. Torrance, "A field guide to BRDF models," *Cornell Program of Computer Graphics*, PCG-04-01, (Cornell Program of Computer Graphics, Ithaca, NY, 2004) <https://www.graphics.cornell.edu/pubs/2004/WLT04b.pdf>.
- W. Matusik, H. Pfister, M. Brand, and L. McMillan, "A data-driven reflectance model," *ACM Trans. Graph.* **22**, 759–769 (2003).
- D. Guarnera, G. C. Guarnera, A. Ghosh, C. Denk, and M. Glencross, "BRDF Representation and Acquisition," *Comput. Graph. Forum* **35**, 625–650 (2016).
- C. Wang, N. Snavely, and S. Marschner, "Estimating dual-scale properties of glossy surfaces from step-edge lighting," *ACM Trans. Graph.* **30** 172, 12 (2011).
- M. Anderson, R. Motta, S. Chandrasekar, and M. Stokes, "Proposal for a standard default color space for the internet—RGB," *Proc. IS&T Color and Imaging Conf. '96* (IS&T, Springfield, VA, 1996), pp. 238–245.
- K. M. Lam, "Metamerism and Colour Constancy," Ph.D. thesis, (University of Bradford, 1985).
- P. G. J. Barten, "Physical model for the contrast sensitivity of the human eye," *Proc. SPIE, Human Vision, Visual Processing, and Digital Display III* (International Society for Optics and Photonics, San Jose, CA, 1992), Vol. 1666.
- Elcometer, <http://www.elcometer.com/en/appearance/distinctiveness-of-image/elcometer-408-triple-angle-gloss-doi-meter.html>.
- G. J. Ward, "The RADIANCE lighting simulation and rendering system," *Proc. 21st Annual Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH '94)* (ACM, New York, NY, USA, 1994), pp. 459–472.
- P. D. Burns, "Slanted-edge MTF for digital camera and scanner analysis," *Proc. IS&T PICS 2000: Image Processing, Image Quality, Image Capture Systems Conf.* (IS&T, Springfield, VA, 2000), pp. 135–138.
- K. E. Torrance and E. M. Sparrow, "Theory for off-specular reflection from roughened surfaces," *J. Opt. Soc. Am.* **57**, 1105–1114 (1967).