Experimental Validation of Analytical BRDF Models

Addy Ngan Frédo Durand Wojciech Matusik*

Massachusetts Institute of Technology

1 Introduction

The Bidirectional Reflectance Distribution Function (BRDF) describes the interaction of light with matter and is an important property for shading. A variety of analytical models have been proposed to represent BRDFs. However, validation of these models have been scarce due to the lack of high resolution measured data. Building upon a recent dataset of over a hundred BRDFs acquired at high resolution [Matusik et al. 2003], we evaluate the performance of several popular BRDF models in terms of their ability to fit the measured data. While previous work in BRDF modeling have validated their models with some measurements, our work is the first to quantitatively compare different models based on one sizable dataset covering a wide class of materials. We believe this work can serve as a guide for practitioners in the field of appearance modeling to choose the right model that suits their needs.

2 Analysis

Methodology Our dataset consists of 108 isotropic BRDFs including metals, plastics, painted surfaces, cloth, etc. The angular resolution for the measurement is higher than 1°, and the dynamic range is 1 in 10^8. The BRDFs are acquired by capturing images of a sphere sample lit by a point source from a dense set of directions. We choose five analytical models to compare: Ward, Blinn-Phong, Cook-Torrance, Lafortune and Ashikhmin-Shirley. The models are fitted to the measurement using a nonlinear optimization routine, minimizing the squared difference between the BRDFs multiplied by the cosine of the incident angle, integrated over the hemispher es.

Single Lobe First we fit the dataset using a single specular lobe. The Ashikhmin and Cook-Torrance models perform the best in most cases, while Blinn-Phong and Ward yield the highest errors. The errors for the Lafortune model is systematically higher than Cook-Torrance and Ashikhmin, despite having the same number of free parameters. The differences between the BRDF fits are also perceptually significant in rendered images.

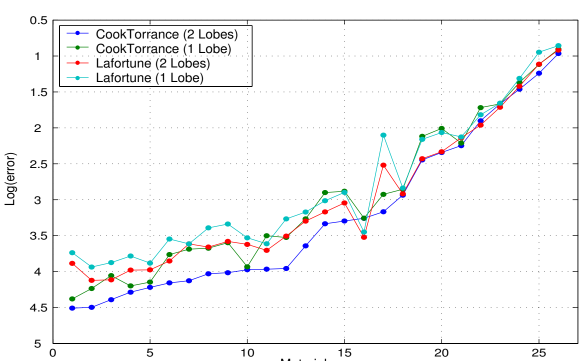

Figure 1: The fitting errors (logarithmic scale) with one/two lobes of the Cook-Torrance and the Lafortune models.

Multiple Lobes Cook and Torrance suggest that two or more specular lobes are necessary to model surfaces where there are multiple scales of roughness. Lafortune's model serve as basis functions for representing BRDFs and thus multiple lobes are also expected in general. To study the effect of an extra lobe, we choose 26 materials from the dataset, typically materials that have multiple layers of finishing, e.g. metallic paints. The fitting errors with the Cook-Torrance and the Lafortune models with one and two lobes are shown in Figure 1. In 14 cases for Cook-Torrance and 11 cases for Lafortune, an extra lobe reduces the error by more than 20%. Again, we note that the Lafortune model perform worse than the Cook-Torrance model.

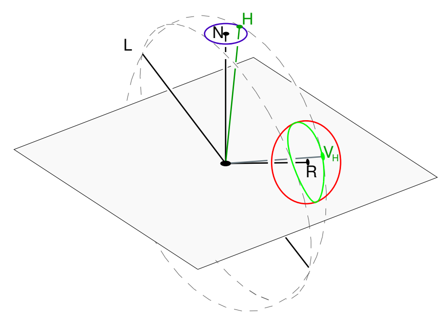

Figure 2: The $V \cdot R$ lobe compared with the $H \cdot N$ lobe remapped to the outgoing directions.

Mirror direction considered harmful Our experimental validation highlights the profound difference between two popular formulations of the specular lobe: 1. defined around the mirror direction symmetrically ($V \cdot R$), and 2. defined by the normal and the half-vector ($H \cdot N$). Most previous discussions of the behavior of reflectance have focused on the increase in intensity of the reflection when the incident angle moves toward grazing. We found that the shape of the lobe also has important implications. We illustrate the difference between the two representations in Figure 2. Observe that the lobe corresponding to $H \cdot N$ (light green) is more circular when L is near N, and becomes narrower and asymmetric as L moves towards grazing angle. On the contrary, the $V \cdot R$ (red) lobe is always circular around the mirror direction R.

Plots of the measured BRDFs verify that the shapes of the specular lobes are more consistent with the $H \cdot N$ lobe in almost all cases. To quantify the differences, we re-express the Blinn-Phong model using $V \cdot R$, and we find that the fitting errors for the $V \cdot R$ version is higher than the original Blinn-Phong model in almost all cases, and 40% higher on average. We believe this is one reason for the inferior performance of the Lafortune model, as it uses a formulation similar to $V \cdot R$, where R is replaced by a generalized mirror direction. The superiority of the $H \cdot N$ lobe is in strong opposition to some beliefs that the half-vector representation is just another way to express the specular lobe.

References

MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics 22*, 3 (July), 759–769.

*e-mail:addy,fredo,wojciech@graphics.csail.mit.edu