Accuracy Validation of Specular Normals and Roughness from Spherical Statistics

Ryan Webb and Pieter Peers

Department of Computer Science

The staggering advances over the past few decades in computational resources and in digital imaging have spurred rapid progress in Computer Graphics. Thanks to the performance and affordability of commodity hardware, computer generated imagery can now be synthesized at a level of photorealism unimaginable just a decade ago. Generating such photorealistic imagery requires one to compute the full light transport through a digital 3D scene. This not only requires powerful simulation algorithms, but also an accurate and detailed description of the scene geometry, the light sources, and a description of the material properties. Modeling all of these components is a difficult and arduous process. The most successful techniques to date rely on measurements of the physical world. In this work, we focus on digitizing the appearance of real-world materials.

Digital cameras are a convenient tool for obtaining information from physical scenes and materials, and nowadays have become the cornerstone of measurement-based modeling techniques. However, a digital photograph only provides a 2-dimensional "slice" of the appearance of a material—it only captures an instance of the appearance of the material under a fixed lighting condition and from a fixed viewpoint. One approach for obtaining a full description (i.e., with variable lighting and viewpoint) of the reflectance properties of a physical material is to capture many such 2-dimensional slices, each under different lighting conditions and from different viewpoints. While such a data-driven digital representation enables photoreal image synthesis, it also has the disadvantage that it: (1) requires significant storage, (2) can suffer from discretization artifacts, and (3) cannot be easily modified by an artist (e.g., to change color, make it a little bit shinier, etc.). To avoid these issues, a data fitting step is often applied after data capture. In this step, a low-parameter physically-based appearance model is fitted to the measured data. While such fitted models do not exactly replicate the measured data, they often provide a good (visual) match, require very little storage, and are much easier to modify by hand (by virtue of only depending on just a few parameters). Satisfactory models can often be constructed with less than 10 parameters (color, shininess or roughness, etc.) compared to millions of data points for the raw measurements. However, the relation between these few model parameters and the raw measured data points is often non-linear. In order to find the best matching parameters, fragile and computationally expensive algorithms are required.

Recently we developed a novel theory, based on spherical statistics, that enables us to separate the non-linear fitting process of these low-parameter appearance models into a robust linear matching phase and a non-linear mapping phase. The advantage of this strategy is that the linear fitting step can be rapidly and robustly computed on any measured dataset, while the non-linear mapping only needs to be precomputed once for each low-parameter appearance model using a small range of parameters which are independent of the nature of the measured material. This two-step process makes the fitting process more robust and significantly reduces the computational cost.

In this project we aim to validate this theory by precomputing the non-linear mappings for four different low-parameter models, and for a large set of appearance measurements of approximately 100 different materials. This requires us to (1) precompute the four different non-linear mappings, (2) compute the linear fits for the 100 different materials, and (3) compute a non-linear fit using prior methods (for validation) on each of the four models for each of the 100 material measurements. This significant computational effort would not have been practical without the SciClone cluster.

The results of our analysis show that for material models that fit the physical material well, we can robustly obtain equally accurate fits in a fraction of the time required by prior non-linear fitting methods. However, our analysis also shows that in the case that the physical model does not fit the data well, our linear fitting technique underestimates the specularity of the measured material. While this is an unexpected result, we have now identified the source of this bias, and are working towards a refined theory.