

Estimating Surface Reflectance from Images

Ron Dror & Ted Adelson

Artificial Intelligence Laboratory
Massachusetts Institute Of Technology
Cambridge, Massachusetts 02139

<http://www.ai.mit.edu>



The Problem: To determine the reflective properties of an object's surface, such as glossiness, mattiness, roughness, etc, given a single image taken with unknown lighting.

Motivation: Physical surfaces such as metal, plastic, or paper have different optical qualities that lead to different characteristics in images. Information on the reflectance of a visual surface can be a valuable cue for inferring the properties or identity of a material. Such information is necessary to render a surface realistically in computer graphics and could also allow a vision algorithm to compensate for the effects of non-Lambertian reflectance in visual tasks such as motion estimation, stereo reconstruction, and object recognition. A complete specification of a surface involves the bidirectional reflectance distribution function (BRDF), which describes the proportion of light reflected from any given incident direction to any given view direction. It is a massive four-dimensional function, which can be carefully measured in the laboratory. Humans are sensitive to a small subset of surface properties, such as the color, gloss, or roughness, and they are able to infer these properties from single images taken in unknown, uncontrolled lighting.

Previous Work: Although the computer graphics community has shown significant recent interest in reflectance measurement, most previous work has used carefully controlled laboratory lighting [1, 2], or has relied on color histogram tricks that do not generalize to monochrome images [2]. Yu et al. [4] determined reflectance and illumination for all surfaces in a room given multiple photographic images of all the surfaces and light sources. We wish to determine reflectance without complete knowledge of illumination, because in most vision applications such illumination information is not readily available.

Approach: We assume initially that we observe a surface of known geometry, such as a sphere. The image of a sphere is determined by the sphere's surface properties combined with the distribution of surrounding light. A chrome-plated sphere simply presents a distorted picture of the environment. Every image of a chrome sphere looks different at a pixel level, but it still looks like chrome (Figure 1). This invariance of appearance must depend on the statistics of the visual scene surrounding the sphere. We use machine learning techniques to find relationships between features of observed images and surface reflectance. After rendering large sets of images of various materials under illumination maps acquired photographically in common visual settings, we compute image features in the image and wavelet domains. We use these features to classify surfaces or to estimate the parameters of finite-parameter BRDF models from computer graphics [3]. We must extract a small set of informative features from the image to achieve acceptable performance.

Difficulty: This problem presents a major challenge for two primary reasons. First, it is highly ill-posed. Different combinations of illumination and reflectance can produce the same image, so we must rely heavily on real-world priors that capture the characteristic statistical structure of illumination and the fact that some types of reflectance are more likely than others. Second, we wish to estimate a small set of reflectance parameters as efficiently as possible while discounting the effects of a much larger set of lighting parameters.

Impact: Image-based reflectance estimates could play a fundamental role in machine vision, as they do in human vision. Most directly, they provide a powerful tool for identifying materials or inferring properties of unknown materials. Indirectly, reflectance estimates could improve a wide variety of machine vision algorithms for motion estimation, three-dimensional reconstruction, and object recognition, which currently suffer from errors due to specularities and other phenomena associated with non-Lambertian reflectance. Improved techniques for reflectance estimation and statistical models for illumination will also facilitate modeling of real-world objects and scenes for computer graphics.

Future Work: We hope to generalize our estimation techniques to larger sets of object shapes and reflectance models. We also plan to perform psychophysical experiments comparing human performance in reflectance estimation to that of our computational method, and to tackle the challenge of determining geometry and reflectance simultaneously.




Figure 1: A chrome sphere photographed under two different illumination conditions. At a pixel level, the images are completely different, but both are recognizably chrome.

Research Support: DOD Fellowship (Dror) and Nippon Telegraph and Telephone Corp.

References:

- [1] Y. Sato, M.D. Wheeler, and K. Ikeuchi. Object Shape and Reflectance Modeling from Observation. *SIGGRAPH*, 1997.
- [2] S. Tominaga and N. Tanaka. Estimating Reflection Parameters from a Single Color Image. *IEEE Computer Graphics and Applications* 20:58-66, 2000.
- [3] G.J. Ward. Measuring and Modeling Anisotropic Reflection. *SIGGRAPH*, 1992.
- [4] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse Global Illumination: Recovering Reflectance Models of Real Scenes from Photographs. *SIGGRAPH*, 1999.