Estimating Surface Reflectance from Images

Ron O. Dror & Edward H. 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, matteness, and roughness, given a single image under unknown, real-world lighting.

Motivation: An ability to recognize surface reflectance under uncontrolled conditions provides a valuable cue for inferring the properties or identity of a material, and also allows realistic surface rendering in computer graphics. Figure 1 shows images of nine spheres, each photographed in two real-world settings. The two images of each sphere are completely different at the pixel level because illumination varies from one location to another. Yet, a human observer easily recognizes that the images in each column represent spheres of similar materials, while the images in different columns represent spheres of different materials. A human could classify a sphere photographed in a third setting into one of these nine categories according to its apparent material properties. We wish to develop a computer algorithm with a similar ability to recognize surface reflectance.

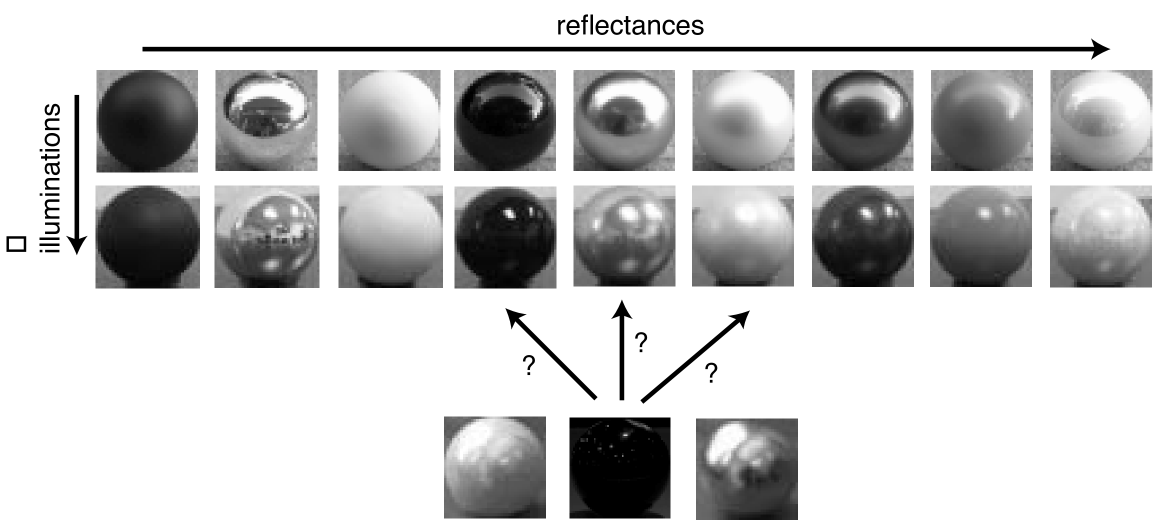

Figure 1: A reflectance classification problem. Each of nine spheres was photographed under multiple real-world illumination conditions. We trained a nine-way reflectance classifier using the images corresponding to six illuminations, and then used it to classify individual images under the seventh illumination.

Previous Work: The importance of reflectance models in computer graphics has motivated several researchers to develop image-based reflectance estimation techniques [4, 3, 2, 5, 1]. These techniques have either used controlled laboratory lighting [4, 3], or assumed that illumination, if not known in advance, could be estimated explicitly from the input images [2, 5, 1].

Approach: The image of a surface such as a sphere is determined by its optical reflectance properties combined with the distribution of surrounding light. A chrome-plated sphere simply presents a distorted picture of the environment. Yet, chrome spheres in different environments all look like chrome. 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 surface reflectance and statistics of observed images. Using photographs or computer

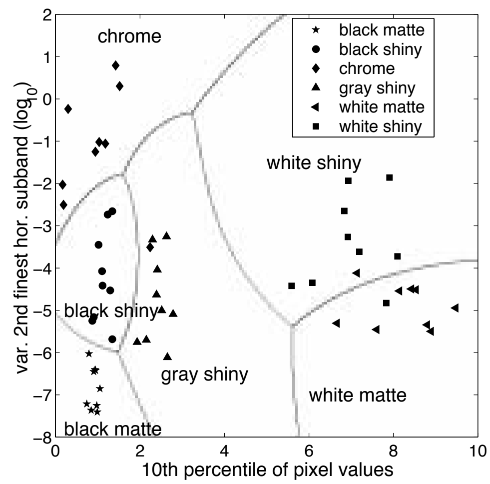

Figure 2: Solid symbols indicate locations in a two-dimensional feature space of images of spheres of six reflectances, each rendered under nine different real-world illuminations. Spheres with the same reflectance tend to cluster in the feature space. Lines separate the regions which our classifier assigns to different reflectances. We can create a more accurate classifier by using additional features.

graphics renderings of various materials under real-world illuminations, we compute image features in the image and wavelet domains. We use these features to classify surface reflectances (Figure 2) or to estimate reflectance parameters. We must extract a small set of informative features from each image to achieve acceptable performance.

We are also performing a series of psychophysical experiments to better understand how humans recognize surface reflectances and other characteristic properties of materials.

Difficulty: This problem presents a major challenge because it is highly ill posed. Different combinations of illumination and reflectance can produce the same image, so we must rely on real-world priors that capture the characteristic statistical structure of illumination.

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 plan to generalize our estimation techniques to the case where geometry is not known in advance, either by estimating reflectance and geometry simultaneously or by identifying image statistics which are invariant to geometry. We also hope to address the reflectance estimation problem when reflectance varies across the surface.

Research Support: This research is supported by NTT, Unilever, an NIH Grant EY11005-04 to E.H.A., and an NDSEG Fellowship to R.O.D.

References:

[1] S. Boivin and A. Gagalowicz. Image-based rendering of diffuse, specular and glossy surfaces from a single image. *Computer Graphics (SIGGRAPH)*, 2001.

[2] S. R. Marschner. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University, Ithaca, NY, 1998.

[3] Y. Sato, M. D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *Computer Graphics (SIGGRAPH)*, 1997.

[4] S. Tominaga and N. Tanaka. Estimating reflection parameters from a single color image. *IEEE Computer Graphics and Applications*, 20:58–66, September / October 2000.

[5] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *Computer Graphics (SIGGRAPH)*, pages 215–24, Los Angeles, 1999.