How do Humans Determine Reflectance Properties under Unknown Illumination?

Roland W. Fleming, Ron O. Dror, Edward H. Adelson
Massachusetts Institute of Technology

1 Introduction

Humans often take for granted their ability to recognize materials such as metal, plastic and paper under a wide range of viewing conditions. Humans succeed at this task despite the fact that different combinations of illumination and surface reflectance can produce identical images. For example, a chrome sphere can be made to look like any other sphere with just the right illumination.

Under ordinary viewing conditions, an observer can draw on multiple sources of information to determine surface reflectance properties such as color, gloss, etc. Potentially useful cues include motion, stereo, knowledge of illumination conditions, and familiarity with the object. In order to test whether human subjects can judge reflectance properties in the absence of such cues, we measure their ability to estimate reflectance properties from isolated images of single spheres. We find that humans are in fact able to judge reflectance from one image taken out of its original context. To do this, they apparently take advantage of statistical regularities in real-world illumination.

Researchers in computer vision and graphics often assume that point source illumination simplifies the process of reflectance estimation. Figure 1 shows synthetic images of three identical spheres under different illuminations. Sphere A was rendered under point source illumination, while spheres B and C were rendered under photographically-captured real-world illumination. The impression of the material quality is clearer in B and C than in A. We show that humans in fact estimate reflectance more reliably under complex realistic illumination than under simple synthetic illumination.

Why might real-world illumination facilitate reflectance estimation? In the real world, light is typically incident on a surface from nearly every direction, in the form of direct illumination from luminous sources or indirect illumination reflected from other surfaces. The illumination at a given point in space can be described by the spherical image acquired by a camera that looks in every direction from that point. Recent work has shown that the spatial structure of such real-world illumination possesses statistical regularity similar to that of natural images [3]. A computer system can take advantage of this regularity in recognizing surface reflectances [2]. Humans might also exploit such regularities to solve this otherwise ill-posed problem.

We use a surface reflectance matching task in order to measure how good subjects are at estimating surface reflectance under various natural and artificial patterns of illumination. We not only quantify the accuracy with which humans estimate

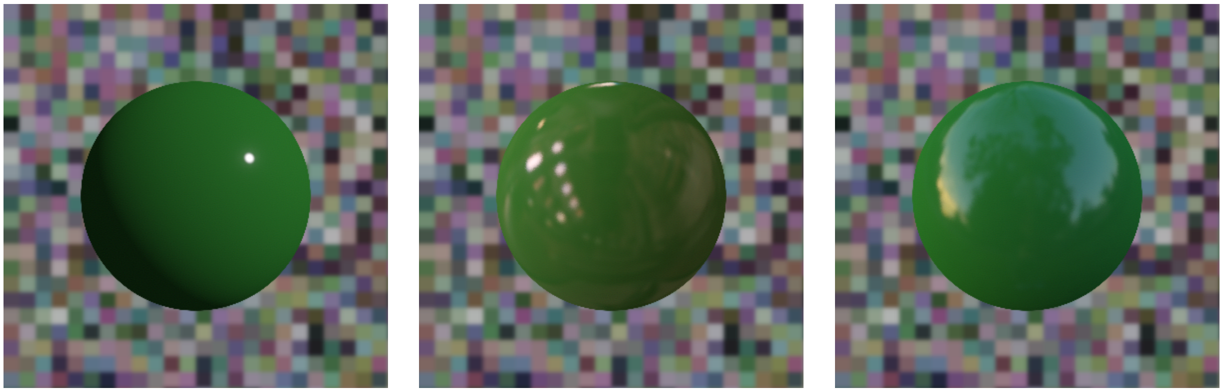

Figure 1: Three spheres with the same reflectance properties rendered under different illuminations. (A) single point source, (B) a real-world illumination used as the standard Match illumination in the experiments, (C) a real-world Test illumination.

surface reflectance from single isolated images, but also identify some of the image properties that subjects use to perform this task.

2 Methods

In our surface reflectance matching task, subjects were presented with two spheres that had been computer rendered under different illuminations. Their task was to adjust the surface reflectance of one sphere (the "Match") until it appeared to be made of the same material as the other sphere (the "Test") despite the difference in illumination. The spheres were viewed against a standard random-check background that was identical for all stimuli and thus provided no information about the illumination. Example matching stimuli are shown in Figure 1.

The reflectances of all rendered spheres were represented by the Ward model [4], a physically realizable variant of the Phong shading model. This model approximates the physical reflectance properties of a wide range of surfaces with a few free parameters. The diffuse reflectance of each sphere was fixed. Subjects simultaneously adjusted two parameters: the amount of specular reflection, and the spatial spread or blur of the specular component.

The Match sphere that the subject adjusted was viewed under the same standard illumination in all experiments. This allowed us to examine systematic effects of illumination upon matching performance. The standard Match illumination and eight real-world Test illuminations were based on spherical high dynamic range light probe images acquired photographically by Debevec *et al.* [1]. In addition, we used three artificial Test illuminations corresponding to a single point source, a collection of point sources, and a single extended rectangular source. This allowed us to compare performance between real-world and artificial conditions of illumination.

A control condition was also included to eliminate the possibility that subjects performed the task by simply matching low-level image cues rather than attending to the material qualities of the spheres. In this control condition, subjects matched

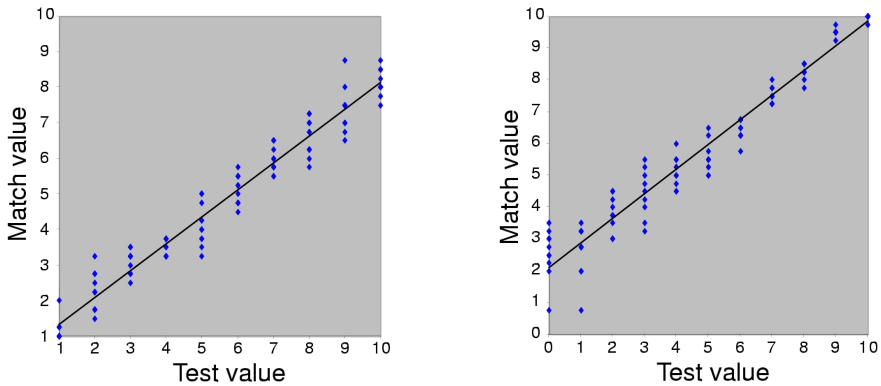

A B

Figure 2: Match values of the parameters as a function of Test values, averaged across four subjects, for one of the real-world illuminations. Each point corresponds to one Test stimulus. The line represents the best-fit linear model. (A) Matches for the parameter controlling the amount of specular reflection. (B) Matches for the parameter controlling width of the specular lobe, which corresponds to the roughness of the surface. Units are arbitrary.

photographic negatives of the original stimuli. Negatives contain many of the same spatial properties as the original stimuli but do not give rise to a coherent, uniform impression of surface reflectance.

3 Results

For surfaces viewed under real-world illumination conditions, subjects could perform the surface reflectance matching task despite the fact that the problem is ill posed. Figure 2 shows example matching functions for one of the real-world illuminations, averaged across four subjects.

Subjects perform better with spheres illuminated by real-world illumination maps than those illuminated by artificial sources. This effect was particularly noticeable under illumination by point sources. Although previous work on reflectance estimation has suggested that a point source provides the most information about reflectance properties, humans perform better under complex illumination more typical of the real world, where light is incident on a surface from nearly every direction.

Humans do make systematic errors in reflectance estimation even when viewing surfaces rendered under complex real-world illumination. Some of these biases depend on illumination conditions. We will discuss the relationship between these systematic biases and certain image properties, such as the maximum observed brightness and the sharpest visible edges. The relationships suggest that certain image cues play an important role in human reflectance estimation.

Performance was considerably worse for the negative control images than for their positive counterparts. This implies that subjects do not simply match low-level image cues, but exploit their percept of the material properties to perform the task.

References

[1] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. *Computer Graphics (SIGGRAPH)*, 2000.

[2] R. O. Dror, E. H. Adelson, and A. S. Willsky. Surface reflectance estimation and natural illumination statistics. In *Proceedings of the Second International Workshop on Statistical and Computational Theories of Vision*, Vancouver, Canada, July 2001.

[3] R. O. Dror, T. Leung, A. S. Willsky, and E. H. Adelson. Statistics of real-world illumination. Submitted.

[4] G. J. Ward. Measuring and modeling anisotropic reflection. *Computer Graphics (SIGGRAPH)*, 26(2):265–72, 1992.