Single Image Appearance Measurement

Neel S. Joshi Henrik Wann Jensen

University of California, San Diego

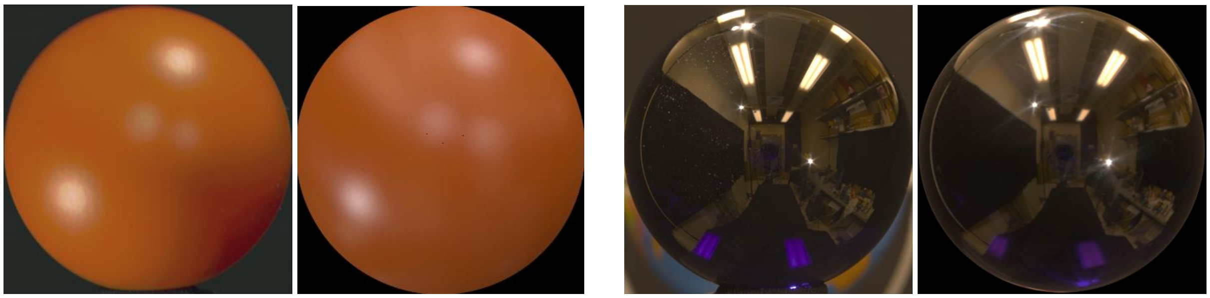

Figure 1: Rendering with recovered appearance model parameters. The image pair on the left is an image of a spray-painted plastic ball and a rendered version of a sphere using a Torrance-Sparrow BRDF with parameters recovered using our method. The right pair is if of a cobalt blue quartz sphere with the real image on the left and an image rendered with recovered parameters on the right.

Abstract

Creating realistic renderings in computer graphics depends on having knowledge of the reflectance properties of materials in the scene. For opaque materials the bi-directional reflectance function (BRDF) captures these properties. Full BRDFs can be measured using highly-accurate gonioreflectometers or more recently with techniques using digital cameras. These approaches sample the BRDF either densely or sparsely and using simple interpolation or fit a BRDF model respectively. Previous methods require images from a number of camera positions with a variety of known light positions. We present a technique to recover BRDF parameters from a single image under arbitrary lighting. Using a single high dynamic range image of material sample of known geometry, an environment map, and an iterative rendering and optimization approach we can recover parameters for a wide variety of materials in short amount of time with less stringent experimental constraints as compared to previous techniques. Our technique in not limited to recovering BRDF parameters for opaque materials. We can fit parameters for any low parameter appearance model that is used in a typical renderer. We have measured a number of opaque and transparent materials and we present results showing comparisons of an image of a sample material and the rendered counterpart using the recovered BRDF parameters. We further use our optimization approach to recover BRDF model parameters from a single render image of a number of materials rendered from a data-driven BRDF model.

1 Introduction

Previous BRDF measurement techniques have used a number of camera and light positions to cover the 4D space of the BRDF. Many of these techniques make assumptions about the isotropy of the BRDF to reduce the domain to three dimensions. While computation techniques often require numerous images and explicit knowledge of lighting and geometry, a human observer often has good intuition of a material's surface reflectance from a single image under arbitrary illumination. This somewhat unfair comparison brings up an interesting point. A single image contains a lot of information about a material's reflectance properties. This compels us to ask to what extent the BRDF of a material can be modeled using computational means on a single image under arbitrary lighting conditions.

The area of image based BRDF modeling and measurement has been addressed by numerous researchers. Several techniques have used single images for modeling BRDFs, but have placed significant restrictions or assumptions on their lighting. [Ikeuchi and Sato 1991] fit data from a single image to a Torrance-Sparrow model; however they assumed their light source was due to parallel-ray illumination. [Boivin and Gagalowicz 2001] also used a single image approach, but used a 3D model containing light positions and scene geometry to iterative compute BRDFs for surfaces in an image of a complex office scene. Most other techniques have used a number of images and a variety of approaches to fill in the missing data for rendering. [Ward 1992] in the development of his anistropic BRDF model, used multiple images and a data fitting processes to compute parameters for his model for a variety of materials. [Marschner et al. 1999] showed a method for quickly acquiring a dense set of BRDF measurements for isotropic materials. With this method rendering is accomplished by interpolating the BRDF by fitting a quadratic function to the sampled data. [Matusik et al. 2003] took this method even further and acquire samples more densely and render using the sampled BRDFs directly.

While many images are needed to produce a purely data-driven BRDF model, with a high resolution photograph, a single image provides enough data for fitting a low-parameter model. Our approach follows along the lines of previous single image methods but relaxes the assumptions made by these methods to allow for BRDF acquisition with arbitrary lab setups. We use an iterative optimization scheme to fit the parameters of an appearance model. By using a rendering step in the optimization pipeline we directly minimize the error in the application domain.

The rest of this paper is organized as follows. First, we will give an overview of what data is needed to compute parameters to a BRDF and our approach to acquiring this data and then we will then discuss our optimization pipeline in detail. We will show results comparing acquired images of material samples to their rendered counterparts using the recovered model parameters. Lastly, we will show results using our approach to recover BRDF model parameters from a single render image of several materials rendered from data-driven BRDF models.

2 Reflectance Measurement

The BRDF of a material is a 4D function that gives the ratio of reflected radiance to the incoming irradiance for an incoming and outgoing ray of light. To compute a BRDF model from an image, for every point on an object it is necessary to known the incoming direction of the light, the outgoing direction of the reflected ray seen by the camera, the irradiance at that point and the reflected radiance. These values can be computed by measuring the scene lighting and scene geometry. Once these values are known, a BRDF can be calculated. While it is possible to use these values to directly fit the data to a BRDF equation and minimize the error of predicated ratio of radiance to irradiance to the captured ratio for image points, we choose to perform the data fitting using a ray-tracer and minimize the error in image space. As BRDF models are used primarily for rendering, it seems most logical to minimize the error in the most common application domain.

2.1 Measuring Scene Lighting

Image based lighting, [Debevec and Malik 1997], has shown that a single high dynamic range image of a mirrored ball or "light probe" fully records the incident radiance, from distant sources, at a single point at the location of the probe. By taking a high dynamic range image of a light probe in our measurement environment we can construct an illumination map. This illumination map is a record of all incident irradiance and incoming light directions at a single point.

While the light probe is omni-directional when imaged with an orthographic camera, it is well-known that with a perspective camera there are virtually no samples acquired from backwards facing direction of the light probe. The common strategy to handle this is to acquire two images of the probe, one from straight on and the other from 90 degrees from that view. It is then necessary to register the angular maps created from the light probe images. HDRShop, the program we used for creating HDR images and environment maps from light probes, provides some tools for performing the registration. When using environment maps for fitting parameters to translucent materials, HDRShop's registration was too inaccurate for our purposes. We instead implemented our own angular map registration, that given an initial guess computes a X,Y, and Z rotation for the angular map that minimizes the least squares error between hand-picked feature point correspondences in each image.

Now that we can construct an accurate environment map, we now know the properties of the incoming light at a particular point in our scene. If we replace the light probe with a sample of material to measure, we can take a single high dynamic range image of this material to record the reflected radiance for every point on the material as viewed from the camera.

2.2 Measuring Scene Geometry

The remaining parameters that are needed are the direction of the outgoing reflected rays from the sample as seen by the camera. To compute this data we need to know the scene geometry. This includes the needing to know the 3D world space coordinates of the sample being measured and the intrinsic and extrinsic properties of the camera. By using a sample of known geometry and by simulating the camera model used to capture the image data, we can implicitly compute the direction of the reflected rays. We use spherical samples of materials to both give a rich set of surface normals for better sampling of the BRDF and to provide simple and easily measurable scene geometry. Using a pinhole perspective model for the camera, and knowing the camera resolution, focal length, and diameter of the spherical sample, all parameters of the camera model can be computed up to an arbitrary scale.

2.3 Measuring Materials

Material acquisition is straightforward and quick. The sample material is placed at the same location as the light probe. Once the light probe has been captured, as long as the lighting levels and positions don't change, any number of samples may be captured simply by acquiring a high-dynamic image of the sample. Other than the image only the diameter of the sample sphere is needed. With these two inputs we can begin fitting parameters to an appearance model.

3 Appearance Model Fitting

Using the measurements of scene lighting and geometry, we can construct a physically accurate synthetic scene to render with a ray-tracer. The scene is constructed with a sphere located at the calculated 3D world coordinates and it is scaled to the known diameter, the camera is placed at the calculated 3D location with the computed field of view and uses a lens model to properly simulate the depth of field of the real camera given its aperture and focal length. The lighting in the scene is provided by using the light probe as a reflection map. Using this scene file and a physically accurate ray-tracer we can solve for the parameters of a appearance model for the captured image.

3.1 Non-Linear Optimization

We use an iterative non-linear optimization procedure to find the parameters of the appearance model that minimizes the error between the rendered image and acquired image in the least squares sense. We have used both a constrained Levenberg Marquardt method and an adaptive simulated annealing method for the non-linear optimization. Both methods minimize error as calculated by a user provided error function. Our error function renders a high dynamic range image of the constructed scene with the current appearance model parameters and the residual function is the sum of squared differences of the rendered and actual image of a sphere. The initial conditions depend on the particular model.

Depending on the material, a number of BRDF models could be appropriate. Our initial experiments have focused on capturing dielectric materials. For the rough dielectrics materials we fit a Torrance-Sparrow model. Our model has eight parameters, 3 for RGB diffuse reflectance, 3 for RGB specular reflectance, 1 for roughness, 1 for index of refraction. For the transparent dielectric there are 3 RGB parameters for transmission color and 1 for index of refraction. We've found setting the parameters to model an equally diffuse, specular, and rough object with a common index of refraction of 1.5, is a reasonable start, but we have also found that the naïve implementation of solving for all parameters at once is not always the most successful.

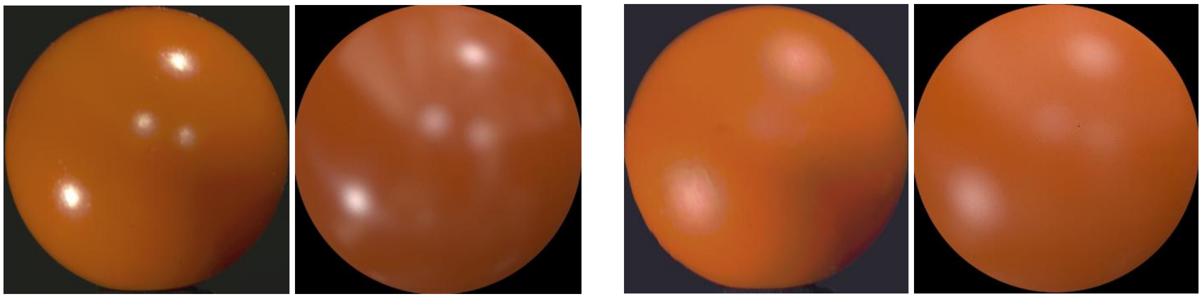

Figure 2: Painted plastic spheres. (a) The left image is of a ball with semi-gloss paint. The image on the right is rendered version of a sphere using recovered parameters. (b) The left image is a ball with matte paint. The image on the right is the rendered version.

3.2 Multipass Optimization

Instead of solving for all parameters at once we use a three pass method to recover the BRDF parameters for the Torrance-Sparrow model. The first pass solves for the diffuse reflectance, which we initialize to (0.5, 0.5, 0.5). The second pass solves for the specular component, index of refraction, and roughness, which we initialize to (0.5, 0.5, 0.5), 1.5, and 0.075. The final pass solves for all the parameters using the estimates from the first two passes as the initial parameters.

Furthermore the first two passes don't use all the pixels in the image. Instead each pass uses a mask image when computing the residual error. The mask for the first step is an estimate of which pixels are largely due to diffuse reflectance and the mask for the second step included the rest of the pixels which are assumed to be largely due to specular reflection.

The pixel masking in the first two passes plays an important part in allowing each pass to converge quickly on accurate parameters. Without a mask of some sort, the diffuse estimate will often overcompensate for intensity due to specular reflections and vice versa. For highly specular materials this can quite significant. Determining a good mask can be somewhat tricky since there is no one property that determines if a pixel intensity is largely due to diffuse or specular reflection. However since this mask is only used to help compute good initial parameters for the final optimization path, only a good approximation is needed.

One simple heuristic that works well is that very bright pixels are often due to specular reflection. Thus a simple threshold on pixel intensity could be used to generate a reasonable mask. Picking a good threshold is important and it will vary depending on the material and scene. Instead of using a simple constant threshold, we use a slightly more sophisticated approach. We render our scene with a completely Lambertian, 100% reflective sphere. This image represents the maximum diffuse contribution to the pixel intensity for every image pixel on the sphere. To create the mask, we use a threshold on the difference of this image and the actual image of the sample. Thus any pixels in the actual target image that have an intensity greater than that of the maximum possible diffuse contribution, must have some specular component to them. The inverse of the diffuse mask is used for the second pass specular, roughness, and index of refraction step. Both masks are also adjusted to mask out pixels that don't fall on the sphere.

For the transparent materials we acquired, we found simulated annealing necessary as we could not easily segment and solve for separate parameters in multiple passes to compute good initial guesses as we did for the other materials. For the transparent dielectrics we initialize the transmission color to (0.5, 0.5, 0.5) and the index of refraction to 1.5 before performing simulated annealing.

4 Results

In our experiments we have acquired images of several spray-painted plastic balls, each pained with satin, semi-glossy, and matte paint. We also captured two color quartz crystal spheres and fit a dielectric shader model to them by solving for index of refraction and transmission color. We have also rendered spheres using data driven BRDF models for several dielectrics and metals measured by [Matusik et al. 2003] and used our rendering and optimization approach to fit Torrance-Sparrow models to single images rendered from there data.

Figure 1 shows an image of a plastic ball painted with satin paint and the rendered result using the recovered BRDF parameters. Figure 2 shows two other painted balls painted with matte and semi-gloss paint. Our method recovers and reproduces the spread of the highlight well for the three painted balls. For the slightly less glossy balls the overall appearance is quite similar. For the semi-glossy ball in Figure 2 the rest of the surface other than the highlight is less convincing.

Figure 1 and Figure 3 show captured and rendered images of deep blue and light purple crystal balls. For the deep blue ball where the transmission component was low, but colorful, the optimization process recovered the color nicely. For the light purple sphere, the transmission color is subtle and as a result a mismatch in color produced very little error in the error function, thus affected the optimization process very little. Thus the recovered color is slightly incorrect.

Figure 4 shows several materials rendered from data-driven BRDFs, [Matusik et al. 2003], under arbitrary artificial lights compared to spheres rendered using a Torrance-Sparrow BRDF model with the recovered parameters.

Table 1 lists the recovered parameters for the appearance models used to render the images in Figures 1 through 4.

5 Discussion and Future Work

We have shown how to efficiently and accurate fit parameters of appearance models to materials by matching a rendered image to an actual image of a material sample using an iterative optimization approach. This method is fast and requires little overhead. It allows BRDFs to be computed under less stringent experimental conditions with convincing results.

This method is limited by the power of the appearance model used. As it is a method for recovering model parameters, it falls

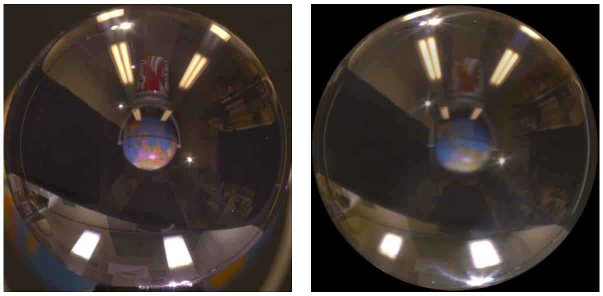

Figure 3: Light purple quartz sphere. The captured image is on the left and the rendered image is on the right. The rendered image is slightly blurry due to limited resolution in the environment map.

short by design, where the model fails to capture the reflectance properties of a material. Many materials observe much more complicated reflectance functions than what can be modeled by a low parameter model. As our results show certain materials do not fit these models well and thus produce less convincing renderings. For these cases, if the inaccuracy is undesirable, using a data-driven model may be the best solution. Our approach can help identify these situations with relative ease.

There are several areas of future work. We plan to acquire anisotropic metals and fit a Ward reflectance model to them. We also plan to acquire more transparent and translucent materials.

It is our hope that this method will increase realism in current graphics renderings that use appearance models by allowing users to choose more accurate parameters for the models they use instead of the current approach which is often for an artist to hand tweak the parameters. We believe that this approach, as it is fast relatively painless, will also allow the computer graphics community to more easily evaluate current BRDF models and their appropriateness for modeling real world materials.

References

BOIVIN S. AND GAGALOWICZ A. "Image-Based Rendering of Diffuse, Specular and Glossy Surfaces from a Single Image." Proceedings of SIGGRAPH 2001.

DEBEVEC P AND MALIK J. "Recovering High Dynamic Range Radiance Maps from Photographs." Proceedings of SIGGRAPH 1997.

IKEUCHI K. AND SATO K. "Determining reflectance properties of an object using range and brightness images." IEEE Transactions on Pattern Analysis and Machine Intelligence, November, 1991.

MARSCHNER S., WESTIN S., LAFORTUNE E., TORRANCE K., AND GREENBERG D. "Image-based BRDF Measurement Including Human Skin." Eurographics Workshop on Rendering 1999.

MATUSIK, W., PFISTER H., BRAND M., AND MCMILLAN L. "A Data-Driven Reflectance Model." Proceedings of SIGGRAPH 2003.

WARD G. "Measuring and Modeling Anisotropic Reflection." Proceedings of SIGGRAPH 1992.

Material	Model	Parameters
Satin paint	Torrance-Sparrow	diffuse 0.124 0.026 0.00, specular 0.104 0.104 0.104 roughness 0.079 ior 2.34
Semi-Gloss Paint	Torrance-Sparrow	diffuse 0.113 0.021 0.000 specular 0.117 0.117 0.117 roughness 0.0541 ior 2.310
Matte Paint	Torrance-Sparrow	diffuse 0.169 0.033 0.000 specular 0.153 0.153 0.153 roughness 0.106 ior 1.96
Yellow plastic	Torrance-Sparrow	diffuse 0.292 0.117 0.027, specular 0.292 0.267 0.239 roughness 0.039 ior 1.43
Teflon	Torrance-Sparrow	diffuse 0.295 0.292 0.294 specular 0.446 0.482 0.527 roughness 0.220 ior 1.637
Nylon	Torrance-Sparrow	diffuse 0.158 0.178 0.193 specular 0.841 0.818 0.795 roughness 0.067 ior 1.150
Silver	Torrance-Sparrow	diffuse 0.124 0.127 0.133 specular 0.676 0.606 0.533 roughness 0.138 ior 3.23
Green acrylic	Torrance-Sparrow	diffuse 0.022 0.080 0.050 specular 0.253 0.247 0.240 roughness 0.005 ior 1.44
Aluminum-bronze	Torrance-Sparrow	diffuse 0.095 0.076 0.059 specular 0.510 0.316 0.175 roughness 0.023 ior 1.569
Copper	Torrance-Sparrow	diffuse 0.063 0.043 0.034 specular 0.713 0.510 0.388 roughness 0.005 ior 2.603
Nickel	Torrance-Sparrow	diffuse 0.047 0.047 0.049 specular 0.909 0.767 0.650 roughness 0.033 ior 2.402
Blue quartz	Dielectirc	ior 1.517 transmission 0.039 0.007 0.163
Purple quartz	Dielectirc	ior 1.569 transmission 0.611 0.693 0.936

Table 1: Recovered model parameters for various materials

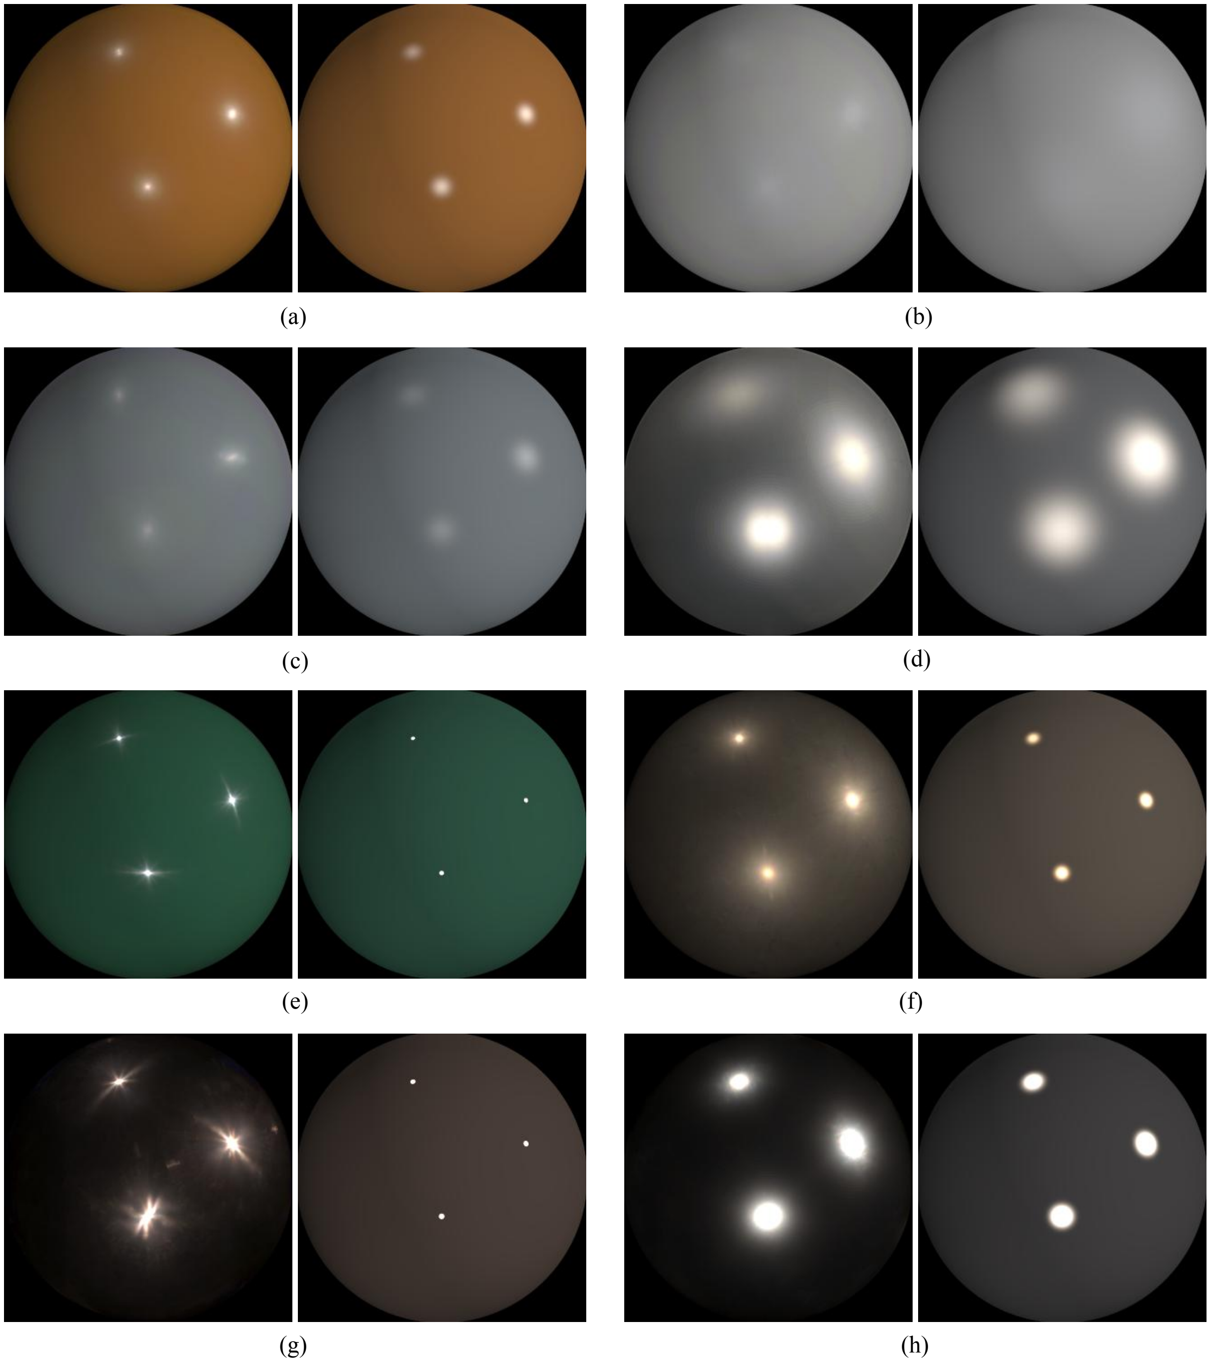

Figure 4: Model fitting to data-driven BRDFs. For each image pair the left image was rendered using a data-driven BRDF model and the right image was rendered using a Torrance-Sparrow model with parameters recovered using our method. (a) Yellow plastic. (b) Teflon. (c) Nylon. (d) Silver. (e) Green acrylic. (f) Aluminum-bronze. (g) Copper. (h) Nickel.