Reflectance Estimation Under Natural Illumination

Neil Alldrin*
University of California, San Diego

Figure 1: Renderings using measured BRDFs.

Abstract

Knowledge of material reflectance properties is of central concern in computer graphics. Traditionally, BRDFs have been estimated with low dimensional parametric models while more recently data-driven techniques have become more common. While some progress has been made to ease the pain of direct BRDF measurements, the process is still far from ideal, typically requiring carefully calibrated aparatus and controlled lighting. The goal of this project is to ease this burden by enabling BRDF measurements under natural illumination using spherical harmonics.

CR Categories: I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Reflectance; I.2.10 [Artificial Intelligence]: Vision and Scene Understanding—Modeling and recovery of physical attributes

Keywords: brdf estimation, natural illumination

1 Introduction

Recently, data-driven techniques for modeling material reflectance properties have been gaining in popularity [Matusik et al. 2003a; Matusik et al. 2003b; Marschner et al. 2000]. This has in large part been spawned by image-based measurement devices which capture multiple samples of the 4D BRDF function in a single image. This drastically reduces the number of measurements required to fully sample a BRDF over more traditional devices like the gonioreflectometer. One problem that still persists with these setups is that the lighting must be carefully controlled.

This paper presents initial results of a novel technique for measuring isotropic BRDFs that enables measurements in uncontrolled lighting environments. The basic idea is to recover spherical harmonic representations of BRDFs by forming constraints from measured lighting and exitant radiance. First, the lighting environment is recovered with a light probe. Then a set of images of a material of known geometry are taken in the same lighting environment. Since the exitant radiance is computed in spherical harmonics as a linear sum of lighting and BRDF coefficients, we are able to form a system of linear constraints for the BRDF if we know the lighting and exitant radiance. The solution to such a linear system yields the spherical harmonic approximation to the BRDF.

*e-mail: nalldrin@cs.ucsd.edu

We begin with a survey of previous work, then describe our technique in more detail. Finally, we conclude with some initial results and future work.

2 Previous Work

The study of material reflectance properties has a long history in computer graphics. Initially, most reflectance models were represented by low-order parameterizations of the bi-directional reflecance function (BRDF). These models can be categorized as empirical and physically based. Empirical models such as the Lamertian model and the Phong model focus on creating "looking good" with less emphasis on being physically valid while physically based models like Torrence-Sparrow are derived from analysis of the physical properties of materials.

The gonioreflectometer was the first device to actually *measure* the reflectance properties of materials. However, the process is extremely slow as it only collects one sample at a time of a four-dimensional function. To overcome this limitation, various image based techniques were developed that collect multiple samples of the BRDF at a time.

Karner et al. [1996] measure a subset of a material's BRDF by placing a light source near a plane of the material and taking an image of the plane. This produces a dense set of samples since the incident and exitant angles vary over the surface of the plane. The collected data is then used to fit the material to some low-dimensional reflectance model. Some practical issues hindering the effectiveness of this technique are that nearby light sources often do not act as point light sources and multiple images are required to capture a full hemisphere of directions (as opposed to imaging a sphere which does this in a single image).

Marschner et al. [Marschner et al. 2000] measure the BRDF of a curved object of known geometry by capturing 2D slices of the 3D isotropic BRDF function using a camera as the sensor device. This drastically reduces the time required to measure the BRDF as a set of images over a 1D domain is all that's needed to collect all 3 dimensions of the isotropic BRDF.

A series of papers by Matusik et al. [2003a; 2003b] further the field of image-based BRDF measurement. In [Matusik et al. 2003a] BRDFs of isotropic materials are measured via images of spherical balls of that material under controlled lighting. Also addressed is determining the space of physically plausible BRDFs. Linear and non-linear models are used to reduce the dimensionality of the data to more compact representations. For the non-linear model, a 15

dimensional subspace is achievable. A method for moving along these subspaces was developed based on perceptually meaningful linear combinations of the 15 basis vectors constituting the 15 dimensional subspace. This enables users to create new BRDFs that still exhibit properties of real materials.

In [Matusik et al. 2003b] a library of 100 densely sampled BRDFs is analyzed to determine the best sampling stategy for measuring isotropic BRDFs. First a set of wavelet basis functions were computed that were best able to reproduce the original (densely-sampled) BRDFs. Under the assumption that novel BRDFs are similar to those in the BRDF library, new materials can be measured with many fewer images by recording fewer samples in low-frequency regions of the BRDF.

Dror et al. [2001a; 2001b; 2001c;] have explored surface reflectance properties for materials of known geometry using single images. Their methodology is based on statistics of natural images and exploits various features to estimate the type of material being observed. This is useful for automated surface classification, but precise surface reflectance properties are not obtained, so it's usefulness for re-rendering the material are limited.

Perhaps the most related work to this project is that of Ramamoorthi. In his PhD dissertation [2002] he analyses forward and inverse rendering in the frequency domain using spherical harmonics. While he briefly touches on BRDF measurement in the frequency domain, the main point of his work is to form a mathematical framework for spherical harmonic representations of the rendering equations.

Other relevant work includes that of Boivin and Gagalowicz [2001], Yu et al. [1999], and Ikeuchi and Sato [1991].

3 Methodology

This section describes the BRDF measurement process in some detail, with the goal of enabling a reader to reproduce the results with a minimum of guesswork.

3.1 What Are We Measuring?

The first step in the process is to collect good material samples to work with. An ideal material sample for this project will be perfectly spherical, completely homogeneous, and isotropic. Moreover, since the BRDF is only an approximation of light transport, we require materials to have negligable subsurface scattering. To measure lighting, we also require a specular ball.

Baltec is a company that specializes in making precision balls. Their chrome plated steel balls work well for light probes and they have a variety of other materials that would work well for BRDF measurements. Other types of balls that are easily obtainable include billiard balls and street hockey balls.

3.2 Imaging

Once material samples are collected, we are ready to begin collecting measurements. This involves finding a good location for placing the spheres, calibrating the camera both geometrically and radiometrically, and finally taking images in high dynamic range.

Because we are not taking into account interreflections or spatially varying lighting effects, we must take care to place the spheres so they do not have nearby surfaces or light sources. Because the spheres cannot float in mid-air, this is impossible. A practical alternative is to place the spheres on a thin stand covered with black matte material.

Geometrically calibrating a camera is a well studied, but sometimes tedious task. Therefore, to simplify things, we assume the camera is orthographic. This assumption holds if the camera is far away from the object being imaged. The rule of thumb is to have the camera positioned at least ten times the distance of the size of the object being imaged. If the camera has an aspect ratio of 1 and no significant distorions (like radial distortion) then the image obtained directly from the camera will suffice.

Radiometric calibration involves finding the mapping between pixel intensity and light intensity arriving at the camera. Luckily, the Canon EOS camera used for this project has a linear response curve, so it doesn't require calibration.

A final requirement for measurements used in this project is that they be done in high dynamic range. Following the techniques of Paul Debevec, we take multiple images at different exposures and combine them into a single high-dynamic range image using HDR-Shop.

3.3 Measuring Lighting

3.3.1 Calibrating the Light Probe

A mirrored ball is never perfectly specular as it always absorbs some percentage of the light in practice. To account for this, the albedo of the mirrored ball needs to be measured and intensities adjusted accordingly. This can be done by placing a diffuse material in the scene so that it is directly visible by the camera and also visible in the reflection off the mirrored ball. The ratio of reflected intensity to direct intensity is the albedo of the mirrored ball.

One potential problem is that the light probe does not have uniform albedo. Rather, there are blemishes throughout the surface. This is not currently accounted for, but could potentially be a serious issue. One workaround would be to take multiple images of the light probe at different orientations and then average the intensities at a given angle. If enough images are taken the blemishes will have an equal effect in all directions.

3.4 Angles and Other Geometric Considerations

Figures 2 and 3 show the spherical coordinate system used throughout this paper. It will be convenient to distinguish between global and local coordinates. Global coordinates, denoted (θ, ϕ), are defined with respect to the viewing angle of a given image. θ and ϕ are the spherical coordinate equivalents to the Euclidean space centered at the center of the image of the sphere with z-axis pointing straight toward the viewer and x and y-axes defined as the x and y-axes of the image. Local coordinates, denoted (θ', ϕ'), are based on the surface normal at a given point on the surface of the imaged sphere. Local and global coordinates are related by a rotation in $SO(3)$.

3.4.1 Surface Normals

A common operation is to determine the surface normal at a given pixel in an image of a sphere. Because of our assumed orthographic projection model, the image is equivalent to the sphere projected

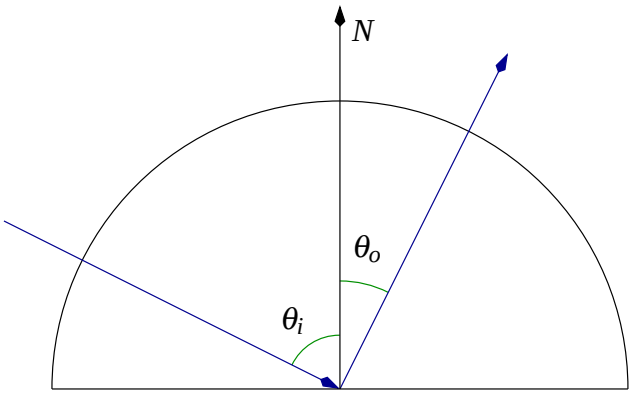

Figure 2: Elevation angles. θ_i is the incident angle and θ_o is the exitant angle.

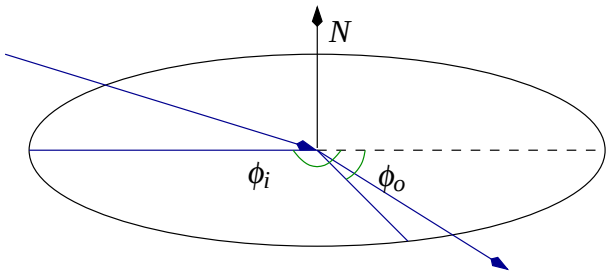

Figure 3: Azimuth angles. ϕ_i is the incident angle and ϕ_o is the exitant angle.

directly down onto a plane (see figure 4). This is a very simple projection model and the surface normal at a given pixel in the image is easily computed in terms of the sphere radius R and center point $x_0 = (x_0, y_0)$. Suppose the pixel coordinate is $x = (x, y)$, then the global angular coordinates of the surface normal are,

$$\begin{aligned} \theta_N &= \sin^{-1}\left(\frac{||x - x_0||}{R}\right) \\ \phi_N &= \tan^{-1}\left(\frac{y - y_0}{x - x_0}\right) \end{aligned}$$

where the global coordinate system has x and y axes aligned with the image, the z axis directed toward the image plane, and the image is centered about the center of the sphere.

3.4.2 Incident Lighting

When computing incident lighting angles from an image of a mirrored ball, the situation is slightly different. This is because what matters is not the surface normal, but the light reflected off the surface of the mirrored ball. Luckily, the incident lighting angle is very closely related to the surface normal angle. The only difference is that the elevation angle is doubled,

$$\begin{aligned} \theta_i &= 2\sin^{-1}\left(\frac{||x - R||}{R}\right) \\ \phi_i &= \tan^{-1}\left(\frac{y - R}{x - R}\right). \end{aligned}$$

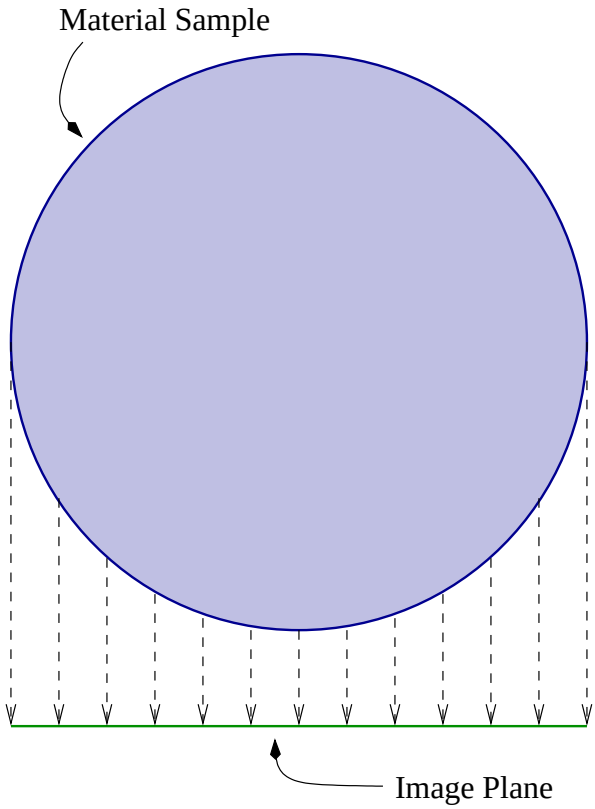

Figure 4: Orthographic projection model.

3.4.3 Differential Solid Angle

To evaluate integrals over the sphere, we need to compute the differential solid angle, $\sin\theta d\theta d\phi$, corresponding to a pixel in the image of a light probe. We can compute $d\theta$ and $d\phi$ separately and then multiply them together. We start with $d\phi$. First note that we should have a total of 2π radians if we integrate the differential azimuthal angles about a fixed radius r. Also note that we have a total of $2\pi r$ pixels to divide the 2π radians among (this is just the circumference in pixel units for a circle of radius r). Therefore we cover $d\phi = 2\pi\frac{1}{2\pi r} = \frac{1}{r}$ radians for a pixel at radius r. Noting that $r = \frac{\theta R}{\pi}$, we can rewrite this as,

$$d\phi = \frac{\pi}{\theta R}.$$

The formulation for $d\theta$ is only slightly more difficult. We can rewrite the formula for θ as

$$\theta = 2\sin^{-1}\left(\frac{r}{R}\right)$$

where $r = x - R$. Now we can take the derivative with respect to r,

$$\frac{d\theta}{dr} = \frac{1}{R}\frac{1}{\sqrt{1 - (r/R)^2}}.$$

Since each pixel occupies distance $dr = 1$ we get

$$\begin{aligned} d\theta &= \frac{1}{R}\frac{1}{\sqrt{1 - (r/R)^2}}dr \\ d\theta &= \frac{1}{R\sqrt{1 - (r/R)^2}}. \end{aligned}$$

Putting it all together, we arrive at a differential surface area of

$$\begin{aligned}\sin\theta d\theta d\phi &= \sin\theta \frac{1}{R\sqrt{1-(r/R)^2}} \frac{\pi}{\theta R} \\ &= \frac{\pi sin\theta}{R\theta\sqrt{1-(r/R)^2}}.\end{aligned}$$

3.5 Spherical Harmonics

The spherical harmonics[1] are a set of orthonormal basis functions defined over the sphere. Given a function over the sphere $f(\theta,\phi)$, the function can be reformulated as

$$f(\theta,\phi) = \sum_{l=0}^{\infty}\sum_{m=-l}^{l} f_{lm}Y_{lm}(\theta,\phi)$$

where Y_{lm} are the orthonormal spherical harmonic basis functions. To compute the spherical harmonic coefficients of f we can integrate over the sphere,

$$f_{lm} = \int_{\phi=0}^{2\pi}\int_{\theta=0}^{\pi} f(\theta,\phi)Y^*_{lm}(\theta,\phi)\sin(\theta)d\theta d\phi.$$

3.5.1 Lighting

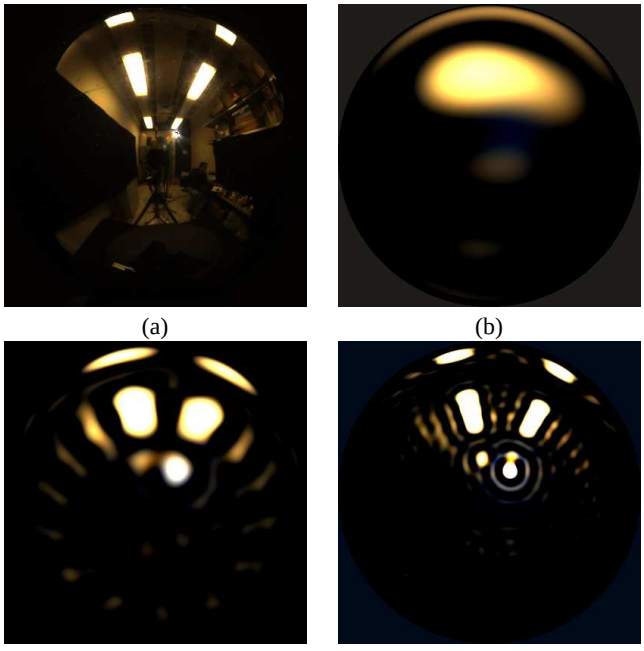

(a) (b) (c) (d)

Figure 5: Comparison of original light probe image (a) with ones generated from a spherical harmonic decomposition using terms up to (b) $l = 5$, (b) $l = 15$, (c) $l = 35$. Notice the distinct ringing effects due to high frequency aliasing.

The incident lighting can be represented as a function over the sphere (if we assume distant lighting). The spherical harmonic decomposition of the lighting is

$$L(\theta_i,\phi_i) = \sum_{l=0}^{\infty}\sum_{m=-l}^{l} L_{lm}Y_{lm}(\theta_i,\phi_i).$$

Figure 5 shows approximations to a lighting environment using spherical harmonics.

3.5.2 Rotation

To represent lighting in local coordinates, the lighting environment needs to be rotated into the local tangent frame. We represent rotations in terms of Euler angles (α,β,γ) with rotations about the Z,Y, and Z axes. In cartesian coordinates the rotation can be written as a product of three rotation matrices R_z^{γ}, R_y^{α}, R_z^{β},

$$X_{\text{rot}} = R_z^{\beta}R_y^{\alpha}R_z^{\gamma}X$$

In the spherical harmonic representation, the rotated lighting coefficient $L_{lm'}$ is a linear combination of the unrotated lighting coefficients $\{L_{lm}|-l\le m\le l\}$ at level l,

$$L(\theta_i,\phi_i) = \sum_{l=0}^{\infty}\sum_{m=-l}^{l}\sum_{m'=-l}^{l} L_{lm}D^l_{mm'}(\alpha,\beta,\gamma)Y_{lm'}(\theta'_i,\phi'_i)$$

where D^l is a matrix satisfying

$$D^l_{mm'}(\alpha,\beta,\gamma) = \int_{\phi=0}^{2\pi}\int_{\theta=0}^{\pi} Y_{lm}(R_{\alpha,\beta,\gamma}(\theta,\phi))Y^*_{lm'}(\theta,\phi)\sin\theta d\theta d\phi.$$

A simpler form for D^l is

$$\begin{aligned}D^l_{mm'}(\alpha,\beta,\gamma) &= d^l_{mm'}(\alpha)e^{Im\beta}e^{Im'\gamma} \\ d^l_{mm'}(\alpha) &= \int_{\phi=0}^{2\pi}\int_{\theta=0}^{\pi} Y_{lm}(R_y(\alpha)(\theta,\phi))Y^*_{lm'}(\theta,\phi)\sin\theta d\theta d\phi.\end{aligned}$$

3.5.3 BRDF

The BRDF, $\rho(\theta'_i,\phi'_i,\theta'_o,\phi'_o)$, is a 4D-function over incident and exitant angles. Since the BRDF is always multiplied by $cos(\theta'_i)$ we create a new function $\hat{\rho}(\theta'_i,\phi'_i,\theta'_o,\phi'_o) = \rho(\theta'_i,\phi'_i,\theta'_o,\phi'_o)\cos\theta'_i$. To enforce reciprocity, we can further multiply by $\cos\theta'_o$, so that $\tilde{\rho} = \hat{\rho}\cos\theta'_o$. Now $\tilde{\rho}$ can be represented using two spherical harmonic expansions, one for the incident lighting and one for exitant radiance,

$$\tilde{\rho}(\theta'_i,\phi'_i,\theta'_o,\phi'_o) = \sum_{l=0}^{\infty}\sum_{n=-l}^{l}\sum_{p=0}^{\infty}\sum_{q=-p}^{p} \tilde{\rho}_{lm,pq}Y^*_{ln}(\theta_i,\phi_i)Y_{pq}(\theta_o,\phi_o).$$

where the expansion about the incident directions are in terms of the complex conjugate of the spherical harmonic coefficients which allows for later simplifications.

Isotropy reduces the BRDF from a 4D function $\rho(\theta'_i,\phi'_i,\theta'_o,\phi'_o)$ to a 3D function $\rho(\theta'_i,,\theta'_o,|\phi'_o-\phi'_i|)$. In spherical harmonics, the coefficients simplify to

$$\hat{\rho}_{lpq} = \hat{\rho}_{lq,pq} = \hat{\rho}_{l(-q),p(-q)}.$$

[1] Most of the results presented here are derived by Ramamoorthi in his PhD dissertation [Ramamoorthi 2002].

3.5.4 Reflected Light Field

The reflected light field, $B(\alpha,\beta,\gamma,\theta_o',\phi_o')$, is a function of surface orientation and exitant radiance direction. If we assume isotropy, γ drops out and we are left with a 4D function $B(\alpha,\beta,\theta_o',\phi_o')$, which can be represented in the frequency domain as

$$B(\alpha,\beta,\theta_o',\phi_o') = \sum_{l=0}^{\infty}\sum_{m=-l}^{l}\sum_{p=0}^{\infty}\sum_{q=-min(l,p)}^{min(l,p)} B_{lmpq}C_{lmpq}(\alpha,\beta,\theta_o',\phi_o')$$

where $C_{lmpq}(\alpha,\beta,\theta_o',\phi_o') = \Lambda_l^{-1}D^l_{mq}(\alpha,\beta)Y_{pq}(\theta_o',\phi_o')$, $\Lambda_l = \sqrt{\frac{4\pi}{2l+1}}$.

Since the reflected light field coefficients can be expressed in terms of the incident lighting coefficients and the BRDF coefficients as $B_{lmpq} = \Lambda_l L_{lm}\hat{\rho}_{lpq}$, we can substitute this into the previous formula,

$$B(\alpha,\beta,\theta_o',\phi_o') = \sum_{l=0}^{\infty}\sum_{m=-l}^{l}\sum_{p=0}^{\infty}\sum_{q=-min(l,p)}^{min(l,p)} L_{lm}D^l_{mq}(\alpha,\beta)Y_{pq}(\theta_o',\phi_o')\hat{\rho}_{lpq}$$

3.6 Recovering the BRDF Coefficients

Ramamoorthi recovers the BRDF coefficients by measuring the incident lighting and the entire reflected light field from which the BRDF coefficients can be directly recovered using the relation $B_{lmpq} = \Lambda_l L_{lm}\hat{\rho}_{lpq}$. The problem with this technique is that it requires taking a dense set of images to recover the reflected light field, but it should be possible to get by with many fewer measurements for low frequency approximations to the BRDF.

Instead of directly solving for the BRDF coefficients, we form a linear system of constraints on the BRDF coefficients and solve it using least squares minimization. The constraints are obtained from the formula for $B(\alpha,\beta,\theta_o',\phi_o')$ in the previous section. In this formula, everything but the BRDF coefficients, $\hat{\rho}_{lpq}$ are known: $B(\alpha,\beta,\theta_o',\phi_o')$ is the intensity measured at a single pixel at a given surface normal, L_{lm} is the lighting which we measure with a mirrored ball, and $D^l_{mq}(\alpha,\beta)$ and $Y_{pq}(\theta_o',\phi_o')$ can be directly computed.

By stringing out the constraints we can form a linear system of the form

$$Ax = b$$

where A is composed of the $L_{lm}D^l_{mq}(\alpha,\beta)Y_{pq}(\theta_o',\phi_o')$ terms, b is the exitant light field, and x is the BRDF coefficients which we are solving for. A standard technique for solving inhomogeneous linear equations of this form is to find the least squares solution,

$$x = (A^TA)^{-1}A^Tb.$$

While this works well in theory, in practice the number of unknowns grows very rapidly with increasing frequency making computations unwieldy. To account for this, we solve for each level l sequentially. The algorithm consists of

- For $l = 1$ to $l_{\max}$
- Solve the system $A_l x_l = b$ for x
- Let $b = b - A_l x_l$
- Repeat

This method yields good results while keeping the linear systems managable.

4 Results

Figures 6 to 8 show some results of the BRDF recovery technique. One readily apparent shortcoming is that higher frequency spherical harmonic representations exhibit a lot of undesired artifacts (see figure 7. As seen in the second row of figure 6, specularities are also not handled very well with the current technique. However, diffuse materials look pretty good as they can be represented well with low-frequency spherical harmonics.

Figure 9 shows the lighting environment used to capture the yellow sphere. Note the computation works despite a quite complex lighting environment.

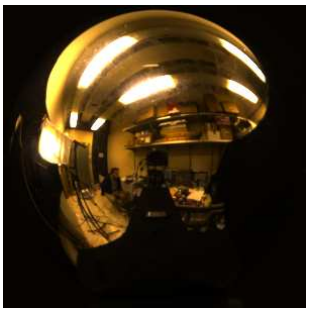

Figure 9: Lighting used to capture the yellow sphere

5 Conclusion

This paper has presented a new way to capture BRDFs under natural illumination. While the results are still somewhat preliminary, the method is capable of capturing the reflectance of diffuse materials. While using higher frequency terms was not particularly successful thus far, I believe that more experimentation will yield better results on more diverse BRDFs. Also, other extensions such as the use of spherical wavelets instead of spherical harmonics could further improve performance.

References

BOIVIN, S., AND GAGALOWICZ, A. 2001. Image-based rendering of diffuse, specular and glossy surfaces from a single image. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, ACM Press, 107–116.

DEBEVEC, P. E., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., 369–378.

DROR, R., LEUNG, T., WILLSKY, A., AND ADELSON, E. Statistics of real-world illumination.

DROR, R., ADELSON, E., AND WILLSKY, A., 2001. Recognition of surface reflectance properties from a single image under unknown real-world illumination.

DROR, R., ADELSON, E., AND WILLSKY, A., 2001. Surface reflectance estimation and natural illumination statistics.

DROR, R. O., ADELSON, E. H., AND WILLSKY, A. S. 2001. Estimating surface reflectance properties from images under unknown illumination. In *Proceedings of the SPIE 4299: Human Vision and Electronic Imaging IV*, SPIE.

HERTZMANN, A., AND SEITZ, S. 2003. Shape and materials by example: A photometric stereo approach. In *IEEE Conference on Computer Vision and Pattern Recognition*. Madison, Wisconsin.

IKEUCHI, K., AND SATO, K. 1991. Determining reflectance properties of an object using range and brightness images. *IEEE Trans. Pattern Anal. Mach. Intell. 13*, 11, 1139–1153.

KARNER, K. F., MAYER, H., AND GERVAUTZ, M. 1996. An image based measurement system for anisotropic reflectance. In *Proceedings of Eurographics*, Eurographics Association, 119–128.

MARSCHNER, S., WESTIN, S., LAFORTUNE, E., AND TORRANCE, K., 2000. Image-based brdf measurement.

MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics 22*, 3 (July), 759–769.

MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. Efficient isotropic brdf measurement. In *Proceedings of the 14th Eurographics workshop on Rendering*, Eurographics Association, 241–247.

RAMAMOORTHI, R., AND HANRAHAN, P. 2001. A signal-processing framework for inverse rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, ACM Press, 117–128.

RAMAMOORTHI, R., 2002. A signal processing framework for forward and inverse rendering, August.

YU, Y., DEBEVEC, P., MALIK, J., AND HAWKINS, T. 1999. Inverse global illumination: recovering reflectance models of real scenes from photographs. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., 215–224.

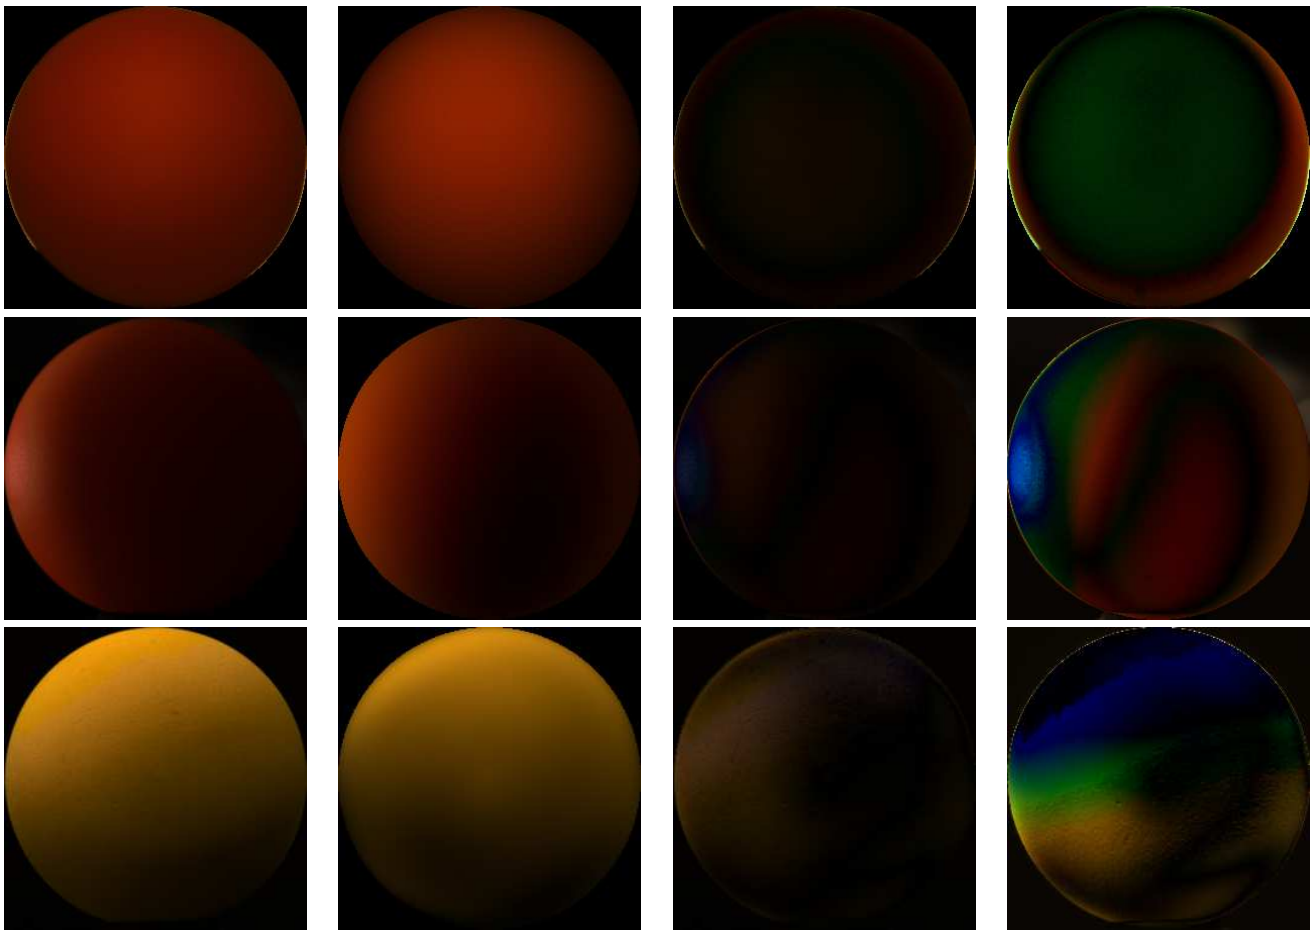

(a) (b) (c) (d)

Figure 6: Rendering with measured BRDFs. (a) Actual image, (b) rendered image, (c) difference, (d) difference magnified by 2-fstops. The top two rows are rendered using spherical harmonic frequencies up to $l = 2$ while the bottom row is rendered using up to $l = 5$.

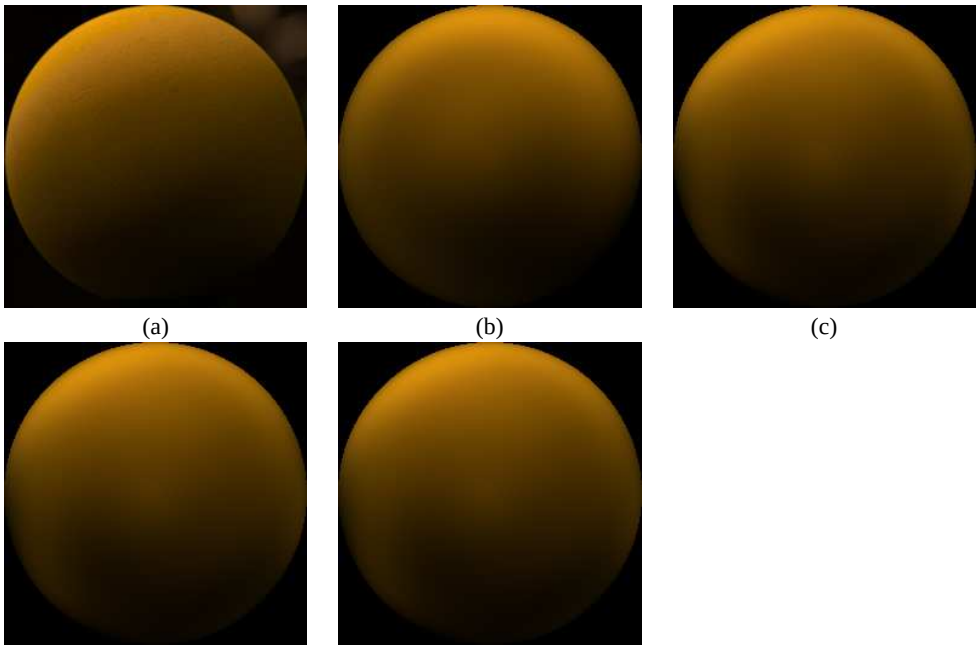

(a) (b) (c)

(d) (e)

Figure 7: Higher order frequency terms. (a) Actual image, (b) $l = 2$, (c) $l = 5$, (d) $l = 10$, (e) $l = 15$.

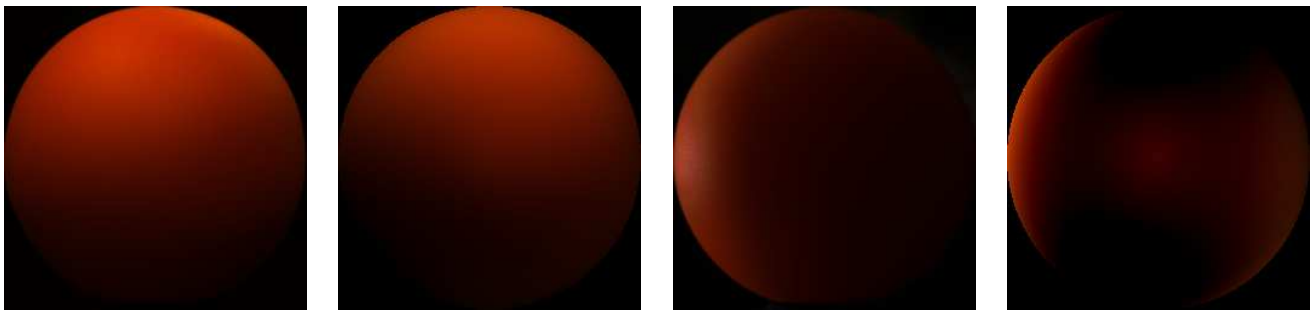

(a) (b) (c) (d)

Figure 8: BRDF estimated from only four images. (a) Actual image A, (b) rendered image A, (c) Actual image B, (d) rendered image B. While it looks ok under some lighting conditions, it doesn't on others.