in B. Girod, H. Niemann, H.-P. Seidel, Proc. *Vision, Modeling and Visualization '99*, infix, ISBN 3-89601-015-8

Image-Based BRDF Reconstruction

Hartmut Schirmacher[1], Wolfgang Heidrich[1], Martin Rubick[2], Detlef Schiron[2], and Hans-Peter Seidel[1]

[1] Max-Planck-Institut für Informatik, Computer Graphics Group
Im Stadtwald, D-66123 Saarbrücken, Germany
email: {schirmacher,heidrich,hpseidel}@mpi-sb.mpg.de

[2] University of Erlangen, Computer Graphics Group
Am Weichselgarten 9, D-91058 Erlangen-Tennenlohe, Germany
email: {mnrubick,dfschiro}@immd9.informatik.uni-erlangen.de

Abstract

This paper presents a method of reconstructing the parameters of a reflectance model from photographs of an object. We assume given object geometry and light source position, as well as a homogeneous reflectance function. We employ a semi-automatic camera calibration and reconstruct radiance values from several photographs using different exposure times. After choosing an appropriate set of samples from the radiance image, we compute the geometric parameters for each sample. The remaining free model parameters are determined through an iterative fitting process. The reconstructed model is validated by simulating the illumination of the test object with a ray tracer and comparing the resulting images with the real photographs. The method is applied to the Phong, Blinn-Phong, and cosine lobe model.

1 Introduction

Image based modeling and rendering has received a lot of attention during the last years. Using images as input often provides a very convenient way of modeling or redisplaying visually complex objects at very low computational and manual cost.

There are a lot of different models for simulating the reflection of light on a surface. The remaining problem, however, is to find the right model and the right parameters in order to imitate real materials. One way of doing this is measuring the materials with specialized sensors, which is of course a very expensive and exhaustive task.

Our approach, in contrast, aims at reconstructing the parameters of simple illumination models such as Phong, Blinn-Phong, and cosine lobe, from a series of photographs.

In Section 2, we will briefly review aspects of illumination modeling, reflectance measurement, and reconstruction of radiance values from images. Section 3 describes our experimental setup as well as the fitting process for the model parameters. Section 4, along with the images on pages 7 and 8, demonstrates and discusses the results of our experiments. In Section 5 we draw conclusions and point out some directions of future research.

2 Previous Work

This section describes the reflection of light off a surface, and the reflection models we have used in our approach. Then we review some aspects of classical reflectance measurement and describe the reconstruction of radiance values from photographs.

2.1 Reflection and BRDF

In the context of this paper we concentrate on a *local* model of light reflection. We assume a single unobtruded isotropic point light source, and we only consider the *direct* illumination falling from the light source onto the object and being reflected in the direction of the viewer.

To this end, we denote $\vec{v}$ the unit vector pointing from some point x on the object to the viewer, $\vec{l}$ the unit vector from x to the light source, and $\vec{n}$ the surface normal in x. The local illumination can be described as follows:

$$L_o(x, \vec{v}) \quad = \quad f_r(x, \vec{v} \leftarrow \vec{l}) g(x) I \cos(\vec{n}, \vec{l}).$$

The equation is formulated in units of *radiance* $L[W/(m^2 sr)]$, which is the radiant flux per projected unit area and solid angle arriving in or leaving a point on a surface. f_r, the so-called *bidirectional reflectance distribution function* or BRDF, controls how much of the incoming irradiance (incoming radiant flux per unit) from some direction $\vec{l}$ gets reflected in some other direction $\vec{v}$. The BRDF has units $[1/sr]$. I is the intensity of the light source (radiant flux per unit solid angle), and g is the source's geometry term in units $[sr/m^2]$. In the case of a point light source, $g = 1sr/r^2$ describes the quadratic falloff with the distance from the light source. The product $g(x)I$ is the incoming radiance at point x from direction $\vec{l}$: $g(x)I = L_i(x, \vec{l})$. For more information on BRDFs, reflectances, light sources, and similar issues, consult one of the many comprehensive textbooks on computer graphics, e.g. [3, 4].

2.2 Reflection Models

Lafortune et al. [5] propose a way of approximating reflectance functions through the so-called *generalized cosine lobe model*. This allows for such effects as off-specular reflection, and retro-reflection. They demonstrate the usefulness of their method by approximating several physically-based models and showing rendered results.

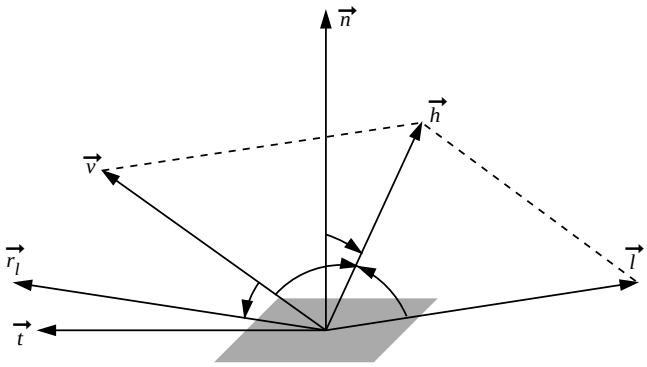

Figure 1: Geometric situation for reflection models. $\vec{h}$ is the so-called halfway vector between $\vec{v}$ and $\vec{l}$: $\vec{h} = (\vec{v} + \vec{l})/|\vec{v} + \vec{l}|$.

In this work, however, we concentrate on very simple BRDFs f_r which can be controlled analytically by only a few parameters. This approach allows us better control over our experiments, and of course it can be extended to more sophisticated reflection models in future. Furthermore, we assume that the BRDF does not change with varying x.

Phong model. The most commonly used lighting model is the that of Phong [7]. It can be computed quite efficiently, but is not physically justified. It describes the reflection of light as a linear combination of perfect diffuse reflection and specular reflection, the latter characterized as a power of the cosine between reflected light vector and viewing direction. Assuming unit vectors, it can be written as

$$L_o(\vec{v}) \quad = \quad \left[k_d < \vec{n}, \vec{l} > + k_s < \vec{r_l}, \vec{v} >^n\right] L_i(\vec{l})$$

where $\vec{n}$ is the surface normal, $\vec{l}$ the vector to the light source, $\vec{v}$ is vector pointing to the viewer, and $\vec{r_l}$ the light vector, reflected at the surface. All these vectors are functions of the position x on the surface. Figure 1 shows the geometric situation on the surface.

If we write down the BRDF f_r corresponding to the Phong model (cf. Sec. 2.1), we get:

$$f_r(\vec{v} \leftarrow \vec{l}) \quad = \quad k_d + k_s \frac{< \vec{r_l}, \vec{v} >^n}{< \vec{n}, \vec{l} >}.$$

Blinn-Phong model. Instead of using the cosine between reflected light vec-

tor and viewing vector, the so-called *Blinn-Phong model* [1] uses the the cosine between the halfway vector and the surface normal:

$$f_r(\vec{v} \leftarrow \vec{l}) \;=\; k_d + k_s \frac{<\vec{h}, \vec{n}>^n}{<\vec{n}, \vec{l}>}.$$

This can be imagined as describing a large number of very small, randomly distributed perfect mirror facets. The halfway vector $\vec{h}$ corresponds to the facet normal direction which would cause perfect mirror reflection from $\vec{l}$ in direction of $\vec{v}$. The powered cosine is a distribution function for the facet normals that falls off with increasing deviation from $\vec{h}$.

Cosine lobe model. Since the Phong model does not conserve energy (e.g. the denominator of the Phong BRDF causes singularities) and is furthermore not symmetric in $\vec{l}$ and $\vec{v}$, Lewis [6] proposed to use the Phong model directly as BRDF, and to add an energy conserving factor:

$$f_r(\vec{v} \leftarrow \vec{l}) \;=\; k_d \;+\; k_s \, \frac{n+2}{2\pi} \; <\vec{r}_l, \vec{v}>^n .$$

2.3 Reflectance Measurement

Many researchers have worked on measuring reflection characteristics of real objects. The classic approach is to use a so-called *gonioreflectometer*, which is an opto-mechanical device with a light sensor, a movable light source and a rotating plate for placing the object. Ward [12, 11] improves on this in terms of cost and ease-of-use, by using a fish eye lens, a CCD camera, and a hemispherical mirror in combination with a movable light source.

However, conventional measurement approaches still require very special devices, and aim at capturing a general BRDF in all four dimensions. Unfortunately, the error rates reported for these instruments are still quite high when compared to the great effort they require. In this paper, we want to go another way, by assuming a certain simple reflection model to begin with, and exploiting knowledge about the model in order to facilitate the imaging, sampling, and fitting process.

2.4 Radiance Reconstruction

Lighting computations are normally done in terms of radiance. However, simple sensors like a digital camera return byte-valued pixel colors, which are, in general, a non-linear function of the true radiance. To bridge this gap, Debevec and Malik [2] have proposed a method of computing the *response curve* of a camera/film combination. By taking photographs from the same scene, but with different exposure times, they reconstruct the response of the sensor to the incident irradiance. Since it is known that doubling the exposure time doubles the amount of energy received at the sensor, one can compute the relation between radiance and sensor response through a least squares approximation. This makes it possible to store the information from multiple shots of the same object as a single *high dynamic range image* containing floating point radiance values instead of byte-valued pixel colors.

In addition to converting the sensor response to radiance values, this method can also be applied inversely to convert back radiance values to pixel colors for displaying simulation results. This way, we can effectively compare the results of our simulation to the photographs in a common color space [2, 9].

3 Reflectance Fitting

Given a simple illumination model (Phong, Blinn-Phong, or cosine lobe), and some photographs of a known object illuminated by a known light source, we want to reconstruct the free parameters of the illumination model, characterizing the object's reflection properties with respect to the given model. In what follows, we describe our experimental setup as well as the computational methods we have applied.

3.1 Light, Scene, and Camera

Our measuring setup (cf. Fig. 2) consists of a fixed light source (a small halogen bulb imi-

Figure 2: Experimental setup with camera (front), light bulb (upper left), and calibration grid.

tating a point light), a consumer quality digital camera on a tripod, and a calibration grid defining the floor of our coordinate system. We measure the position of the light relative to the the grid. Then, we take a photograph of the calibration grid and determine the camera parameters by means of Tsai's calibration method [10]. This way we obtain a common coordinate system for the light, camera, and objects.

After having calibrated the system, we place flat test objects (in practice, tiles of different materials) on the grid. Then, we take a photograph of the tile from the known camera position (with no other light source than our light bulb). For any given pixel of the resulting image, we can determine the world position of the corresponding point on the top of our tile. For each pixel, we can compute the camera and light vector. The surface normal is orthogonal to the tile. So each pixel gives us the resulting radiance as well as all geometric parameters for the illumination function.

3.2 Radiance Reconstruction

The camera is placed in such way that we can observe the highlight caused by the reflection on the tile. We take a series of photographs of each test object with fixed position and varying exposure times, and combine the information from all images into a single high dynamic range image along the lines of [2]. Due to the limited facilities of the simple camera, we were not able to eliminate saturated samples completely.

Since we cannot determine neither the absolute radiance values, nor the absolute power of the light source, these two values result in a single constant scaling factor which remains unknown throughout our experiments. Fortunately, this factor can be neglected since the input and the simulated output are transformed by the same (real and simulated) sensor.

3.3 Parameter Fitting

Each illumination model can be seen as a radiance function in several parameters:

$$f_r \;=\; f_r(k_d, k_s, n,\; x, \vec{l}(x), \vec{v}(x), \vec{r}_l(x), \vec{h}(x))$$

We distinguish the *geometric* parameters $x, \vec{l}, \vec{v}, \vec{h}$, and $\vec{r}_l$, which are known per pixel, from the photometric, or *free*, parameters k_d, k_s, and n. The free parameters have to be determined by an optimization process. To this end, we choose a set of samples (pixels) s from our radiance image. Then we estimate some initial values $\tilde{k}_d, \tilde{k}_s$, and $\tilde{n}$, and compute the *simulated* radiance $\tilde{L}_o$ for each pixel s:

$$\tilde{L}_o(s) \;=\; f_r(s, \tilde{k}_d, \tilde{k}_s, \tilde{n}) \cdot L_i(s) \cdot \cos(\vec{n}, \vec{l})$$

Now we can compute the L_2 error over the chosen set of samples:

$$E(\tilde{k}_d, \tilde{k}_s, \tilde{n}) \;=\; \sum_s ||\tilde{L}_o(s) - L_o(s)||_2$$

Through optimization, we can minimize this error function by modifying the vector of free parameters accordingly. Assuming an isotropic point light source, the gradient of the radiance can be computed analytically for the illumination models presented in Section 2. For details on the optimization process, please refer to [8].

3.4 Sampling

The remaining task is to choose a representative set of samples, since optimizing the parameters using all pixels is bound to be intractable. Trying to use uniformly distributed

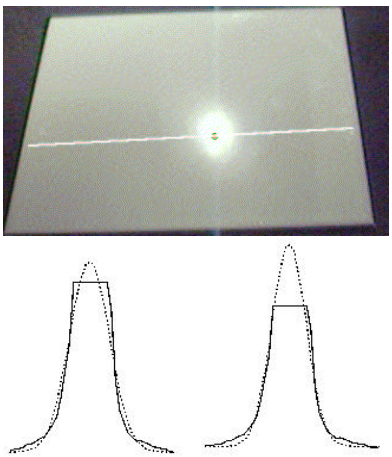

Figure 3: Top: sample line through the highlight on the tile. B.l.: reconstructed (solid line) and simulated (dotted) radiance on the sample line; b.r.: using only unsaturated samples.

random samples has turned out to be numerically fragile, and often the global minimum of of the error cannot be found. This is due to the fact that the image contains a lot of diffuse samples far away from the highlight. In order to robustly reconstruct the Phong exponent n and the specular weight k_s, it is necessary to sample the highlight very densely in order to capture its gradient.

So instead of random sampling, we let the user draw a line through the region of interest. Regular samples along that line are used for the optimization process. We also experimented with 2D sample areas (by letting the user mark a rectangle in the image), but apart from resulting in many more samples and drastically increasing the computation time, the area-based method did not produce very different results. It should be noted that the line-based method can only be applied to isotropic reflection models. If the reflection is not invariant under rotation around the normal, a single line cannot capture all aspects of the model.

Another important issue is the treatment (detection and exclusion) of saturated samples. Figure 3 shows a sample line drawn by the user, the measured radiance along that line (solid lines), and the simulated radiance using the reconstructed Phong model (dotted). The curve reconstructed from all samples gives a much worse radiance approximation than the curve reconstructed from unsaturated samples only.

4 Results

In order to validate our approach, we have first reconstructed reflectance parameters from ray traced images with known reflectance models. As expected, the fitting method was able to find the correct parameters with only a few optimization iterations (in general, less than 15). We always used sample lines through the center of the highlight (which could be detected automatically by finding the maximum radiance value from unsaturated images). We varied the number of samples (100 to 800), as well as the sample line direction. In nearly call cases, the parameters have been reconstructed with an error less than 1/1000. The only limitation of the optimization approach is that it requires reasonable start values for the fitting process in order to converge quickly.

For validating the measurement of real objects, we have created a ray tracing scene description of our experiment (camera, polygon, point light) and simulated the illumination of the polygon using the reconstructed model parameters. We used the inverse response function of our camera in order to remap the radiance values returned by the ray tracer to pixel color values. The original photographs as well as the simulated images (using Phong, Blinn-Phong, and cosine lobe) for four test objects are shown in Figures 4 – 7. In contrast to our experiments with rendered images, the fitting from real images requires a minimum of 500 samples in order to converge robustly. The variation of the coefficient and exponent values lies below 5% for line samples, which seems to be in the same order of magnitude as the noise present in the original images.

The visual similarity of simulated and real images is quite good, considering the shortcomings of our sensor (saturated images, blooming, noise), of the light source (not

point-like), and of the object surface (surface variations not covered in our approach, cf. Fig. 7). It can be noted that the Blinn-Phong model does the best job in reconstructing the shape of the highlight. On the other hand, the cosine lobe model seems to fit the gradient of the real highlight much better than the other two models, which is obviously due to the physically correct cosine weighting (cf. Sec. 2.2).

5 Conclusions

We have presented a method for reconstructing parameters of simple reflection models from images, and demonstrated the validity of our approach by applying it to rendered and real images using the Phong, Blinn-Phong, and cosine lobe model, and by visually comparing simulated images and photographs. We exploited several constraints (known geometry and light source, calibration grid, homogeneous reflectance function) in order to get the desired results, and our method turned out to work quite well for the given setup.

However, the general goal of image based modeling and rendering is to remove most of the above mentioned constraints, so that in future it will be possible to reconstruct many more parameters from a series of images (e.g. approximate geometric model, surface bumps, incoming light field, etc.). There is still much research and engineering work to be done to this end. Results from the different areas of computer vision and computer graphics have to be applied and combined in order to get the maximal range of information out of the images. And, besides that, image based modeling and rendering could profit very much from any improvements in commonly available sensors.

Acknowledgments

Part of the authors' work was funded by the *Deutsche Forschungsgemeinschaft* under grant SFB 603/C2. Thanks to Benno Heigl for passing his nights helping us with the camera calibration and experimental setup.

References

[1] J.F. Blinn and M.E. Newell. Texture and reflection in computer generated images. *Communications of the ACM*, 19:542–546, 1976.

[2] P.E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Proc. SIGGRAPH '97*, pages 369–378, 1997.

[3] J.D. Foley, A. van Dam, S.K. Feiner, and J.F. Hughes. *Computer Graphics, Principles and Practice, Second Edition*. Addison-Wesley, Reading, Massachusetts, 1990.

[4] Andrew Glassner. *Principles of Digital Image Synthesis*. Morgan Kaufmann, 1995.

[5] E.P.F. Lafortune, S.-C. Foo, K.E. Torrance, and D.P. Greenberg. Non-linear approximation of reflectance functions. In *Proc. SIGGRAPH 97*, pages 117–126, 1997.

[6] R.R. Lewis. Making shaders more physically plausible. In *Fourth Eurographics Workshop on Rendering*, pages 47–62, June 1993.

[7] B.T. Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, June 1975.

[8] M. Rubick. Rekonstruktion von Reflexionsdaten aus Bildern. Diplomarbeit, IMMD IX, Universität Erlangen-Nürnberg, 1999.

[9] D. Schiron. Vergleich und Bewertung rekonstruierter photometrischer Daten. Diplomarbeit, IMMD IX, Universität Erlangen-Nürnberg, 1999.

[10] R.Y. Tsai. An efficient and accurate camera calibration technique for 3D machine vision. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, pages 364–374, 1986.

[11] G.J. Ward. Towards more practical reflectance measurements and models. In *Graphics Interface '92 Workshop on Local Illumination*, pages 15–21, May 1992.

[12] G.J. Ward. Measuring and modeling anisotropic reflection. In *Proc. SIGGRAPH '92*, pages 265–272, 1992.

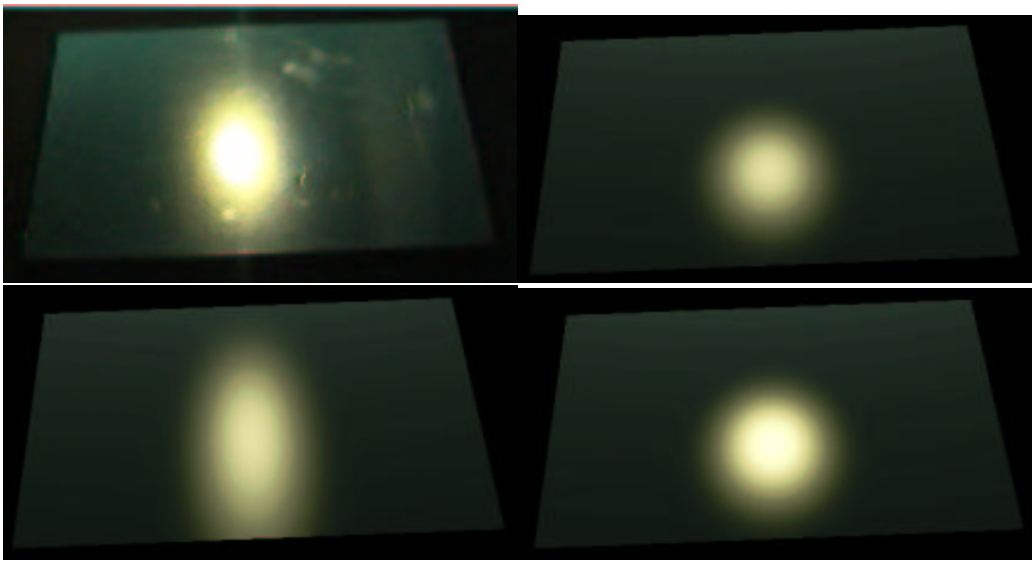

Figure 4: Real and rendered images of a plastic coating. From left to right, top to bottom: photograph; Phong model, Blinn-Phong model, cosine lobe model.

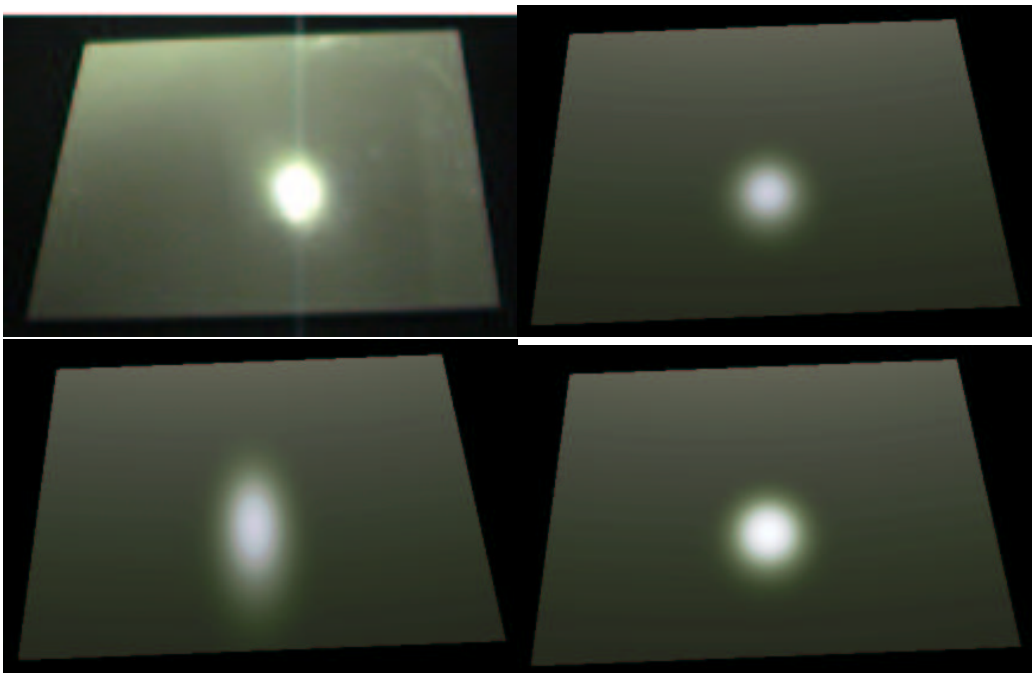

Figure 5: Tile with very specular coating. From left to right, top to bottom: photograph; Phong model, Blinn-Phong model, cosine lobe model. Note the blooming effect on the Blinn-Phong highlight.

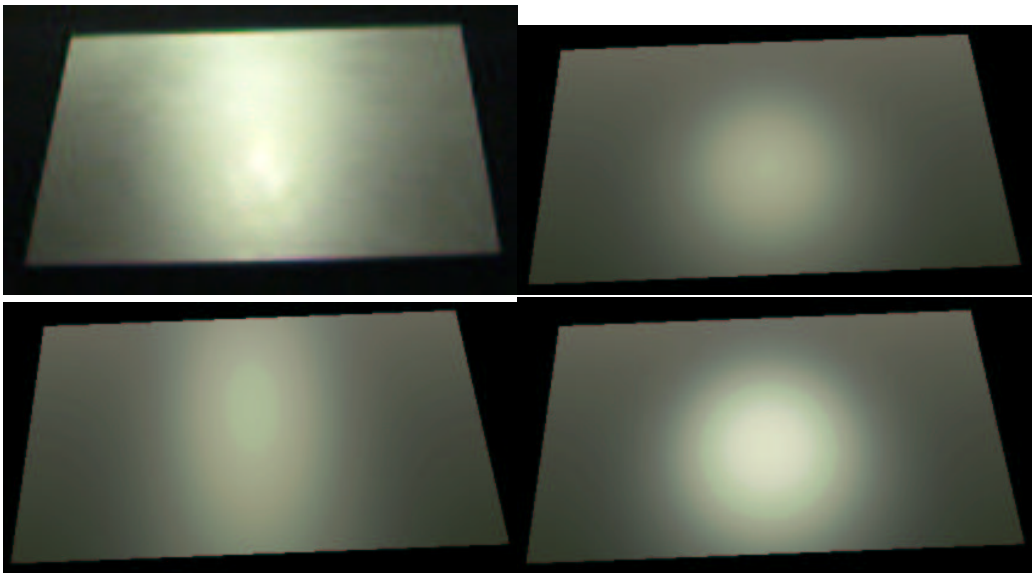

Figure 6: Wood tile. From left to right, top to bottom: photograph, Phong model, Blinn-Phong model, cosine lobe model.

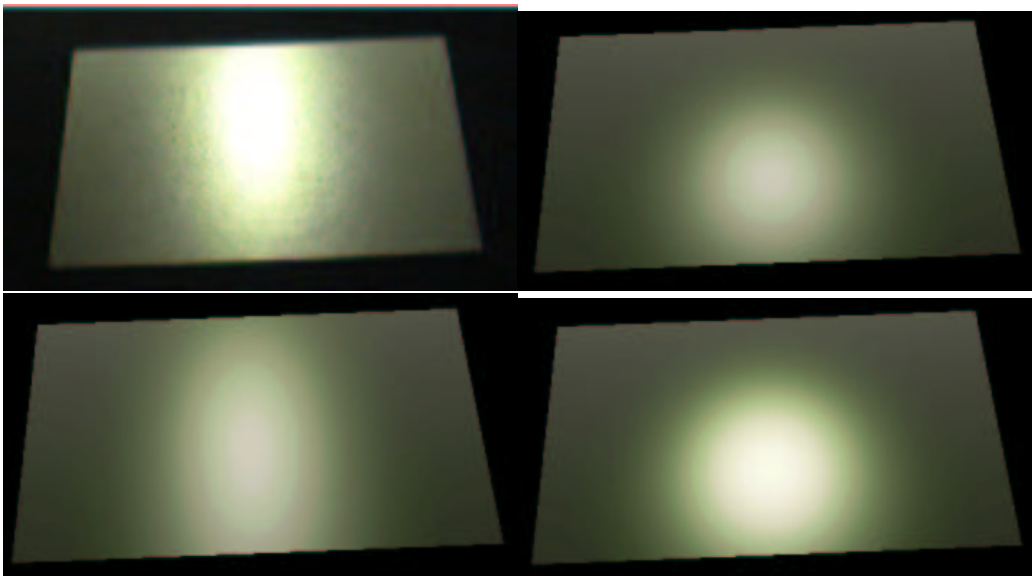

Figure 7: Another tile. From left to right, top to bottom: photograph, Phong model, Blinn-Phong model, cosine lobe model. Note that the visible surface perturbations in the photograph do not disturb the BRDF reconstruction too much.