L-168
LBL-31698
UC-350

Measuring and Modeling Anisotropic Reflection

Gregory J. Ward
Lighting Systems Research Group
Energy & Environment Division
Lawrence Berkeley Laboratory
Berkeley, California 94720

January 1992

This work was supported by the Assistant Secretary for Conservation and Renewable Energy, Office of Building Technologies, Building Equipment Division of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

Measuring and Modeling Anisotropic Reflection

Gregory J. Ward
Lighting Systems Research Group
Lawrence Berkeley Laboratory

ABSTRACT

A new device for measuring the spatial reflectance distributions of surfaces is introduced, along with a new mathematical model of anisotropic reflectance. The reflectance model is both simple and accurate, permitting very efficient reflectance data reduction and reproduction. The validity of the model is substantiated with comparisons to complete measurements of surface reflectance functions gathered with a new reflectometry device. This device uses imaging technology to capture the entire hemisphere of reflected directions instantly, which greatly accelerates the reflectance data gathering process and makes it possible to measure dozens of surfaces in the time that it used to take to do one. Example measurements and simulations are shown, and a table of fitted parameters for several surfaces is presented.

1. Introduction

Numerous empirical and theoretical models for the local reflection of light from surfaces have been introduced over the past 20 years. Empirical and theoretical models have the same goal of reproducing real reflectance functions, but the respective approaches are very different. An empirical model is simply a formula with adjustable parameters designed to fit a certain class of reflectance functions. Little attention is paid to the physical derivation of the model, or the physical significance of its parameters. A good example of an empirical model is the one developed by Sandford [Sandford85]. This is a four parameter model of isotropic reflection, where the parameters must be fit to a specific set of reflectance measurements. While two parameters correspond roughly to measurable quantities such as total reflectance and specularity, the other two parameters have no physical significance and are merely shape variables that make the specular lobe of the model more closely match the data. In contrast to an empirical model, a theoretical model attempts to get closer to the true distribution by starting from physical theory. A good example of a theoretical model is the one derived recently by He et al [He91]. This is also a four parameter isotropic model, but all four parameters have some physical meaning and can in principle be measured separately from the surface reflectance distribution. In practice, however, it is usually necessary to fit even a theoretical model to measurements of reflectance because the physical parameters involved are difficult to measure. This is the case in the He-Torrance model, since measurements of the requisite surface height variance and autocorrelation distance variables are impractical for most surfaces. Thus, the physical derivation of such a model serves primarily to inspire greater confidence, and is not necessarily a practical advantage when it comes to fitting measured data. As in all scientific disciplines, if the theory does not fit the data, then the theory must be discarded, not the data.

But where is the data? There is almost no published data on surface reflectance as a function of angle, and what little data is available is in the form of plane measurements of isotropic surfaces with no rotational variance in their reflectance functions. Thus, we have little to compare our reflectance models to, and no real assurance that they are valid. This means that we may once again be falling back on the "if it looks reasonable then it's OK" philosophy that has misdirected computer graphics so often in the past.

Why is the oldest specular model, the one introduced by Phong in 1975 [Phong75], still the most widely used to this day? This model is neither theoretically plausible nor empirically correct. Any renderings that use the straight Phong model are most likely wrong because the model is not physical, and more light may be emitted than is received for example. The sole virtue of the Phong model is its mathematical simplicity.

Simplicity is indispensable in computer graphics. Simplicity is what permits fast renderings and hardware implementations. Without it, a reflectance model is little more than a novelty. Even a relatively straightforward model such as the one developed by Torrance and Sparrow [Torrance67] and tailored for rendering applications by Blinn [Blinn77] and later Cook [Cook82] has been underutilized in computer graphics due to its moderately complex form. More recent introductions by Poulin and Fournier [Poulin90] as well as He and Torrance [He91] are likely to be even less popular. What is really needed for computer graphics is a simple reflectance model that works for most materials.

Our goal in this paper is not to present the ultimate mathematical model of reflectance, but to provide an elegant formula that is physically valid and fits the measured reflectance data. Here we will present both a new method for measuring isotropic and anisotropic reflectance distributions and a concise formula that fits these data with reasonable accuracy and simplicity.

2. Definition of the BRDF

The interaction of light with a surface can be expressed as a single function, called the *bidirectional reflectance distribution function*, or BRDF for short [Nicodemus77]. This is a function of four angles, two incident and two reflected, as well as the wavelength and polarization of the incident radiation.
For the sake of simplicity, we will leave wavelength and polarization out of our equations, but keep in mind that they are contained implicitly in the function f, which is defined in terms of incident and reflected radiance by the following integral:

$$L(\theta_r,\phi_r) = \int_0^{2\pi} \int_0^{pi/2} L(\theta_i,\phi_i)\, f(\theta_i,\phi_i;\theta_r,\phi_r) \cos\theta_i \sin\theta_i \, d\theta_i \, d\phi_i \qquad (1)$$

where:

θ is the polar angle measured from the surface normal

ϕ is the azimuthal angle measured about the surface normal

$L(\theta_r,\phi_r)$ is the reflected radiance (watts/steradian in SI units)

$L(\theta_r,\phi_r)$ is the incident radiance

$f(\theta_i,\phi_i;\theta_r,\phi_r)$ is the BRDF (steradian^{-1})

notes:

The limits of the θ integral can be changed to $\int_0^{\pi}$ for translucent surfaces.

The function f is *bidirectional* because the incident and reflected directions can be reversed and the function will return the same value. This arises from the fact that the physics of light is the same run backwards as forwards, which is why light-backwards ray tracing works [Whitted80].

3. Measuring the BRDF of a Surface

A device for measuring BRDFs is called a *gonioreflectometer*. The usual design for such a device incorporates a single photometer that is made to move in relation to a surface sample, which itself moves in relation to a light source, all under the control of a computer. Because BRDFs are in general a function of four angles, two incident and two reflected, such a device must have four degrees of mechanical freedom to measure the complete function. This requires substantial complexity in the apparatus and long periods of time to measure a single surface. A typical gonioreflectometer arrangement, designed by Murray-Coleman and Smith [Murray-Coleman90], is shown in Figure 1.

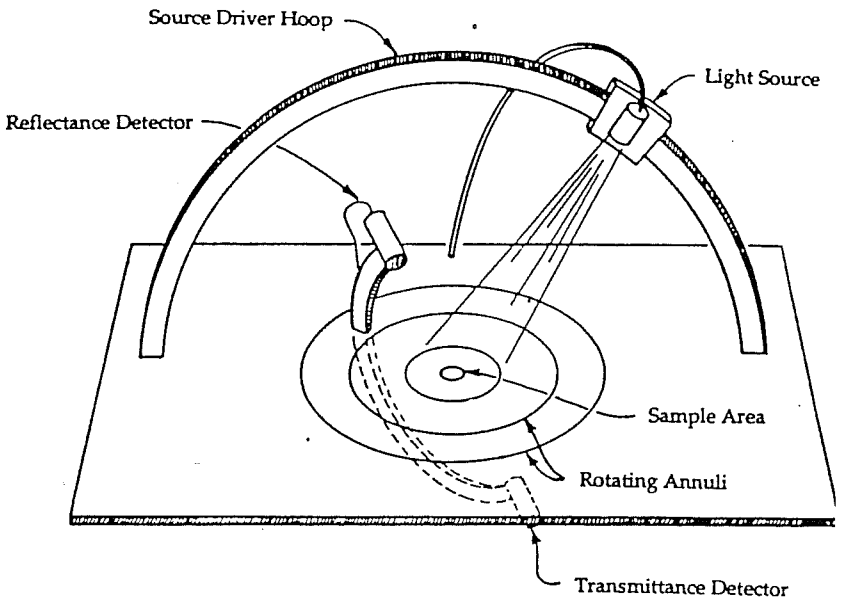

Figure 1. A conventional gonioreflectometer with movable light source and photometer.

There are several labs in North America where one can send a surface sample for BRDF characterization. For a few hundred dollars, one can get a three plane measurement of an isotropic material at four or five angles of incidence. (An isotropic material has a BRDF that is independent of rotation about the normal. Therefore, only one ϕ_i direction is sampled.) Unfortunately, a comprehensive BRDF measurement of an anisotropic surface typically costs a few *thousand* dollars. (An anisotropic material reflects light differently at different angles of rotation, thus multiple ϕ_i directions must be sampled.) Because of the difficulty and expense of the BRDF measurements themselves, only the very richest research programs can afford their own data. This data is essential, however, for the correct modeling of surface reflectance.

3.1. An Imaging Gonioreflectometer

The Lighting Systems Research Group at Lawrence Berkeley Laboratory has developed a relatively simple device for measuring BRDFs that uses imaging technology to obtain results more quickly and at a lower cost than conventional methods. This *imaging gonioreflectometer* has been developed over the past three years and represents an important advance towards more practical characterization of BRDFs for lighting simulation and computer graphics. It is our hope that other laboratories and research institutions will construct their own versions of this apparatus and thereby make BRDF measurement a more common and economical practice.

The basic arrangement of the LBL imaging gonioreflectometer is shown in Figure 2†. The key optical elements are a half-silvered hemisphere or hemi-ellipsoid and a charge-coupled device (CCD) camera with a fisheye lens. Combined, these elements take care of the two degrees of freedom handled by a mechanically controlled photometer in a conventional gonioreflectometer. Light reflected off the sample surface in holder A is collected by the hemispherical mirror and reflected back into the fisheye lens and onto the CCD array B. By focusing the lens at one half the hemisphere radius, a near perfect imaging of the reflected angles takes place. (See ray diagram in Figure 3.) Because of this highly efficient collector arrangement,

†A U.S. patent is pending on the imaging gonioreflectometer. If granted, the patent will restrict other patents on similar devices, but will not otherwise limit the free availability of the invention, since it was developed under Department of Energy funding.

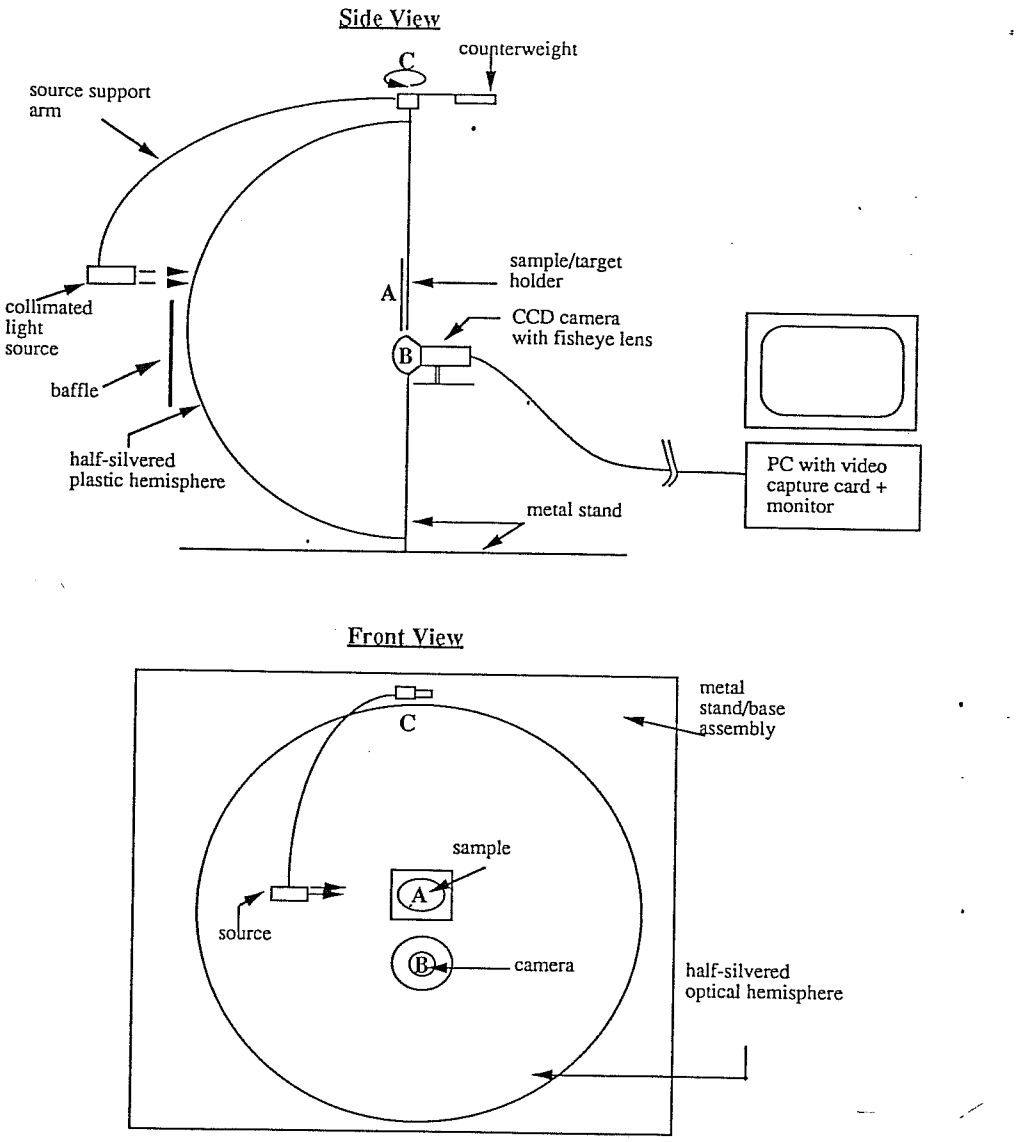

Figure 2. The LBL imaging gonioreflectometer.

the light source does not have to be very bright to obtain a good measurement, and can thus be optimized for collimation to get the best possible angular resolution. In our device, a 3-watt quartz-halogen lamp is used with an optically precise parabolic reflector to produce a well collimated beam. White light is preferable for this type of measurement, although an array of colored filters may be used to measure the spectral dependence of the BRDF. The hemisphere is half-silvered to allow the light beam to illuminate the sample, and an exterior baffle shields the camera from stray radiation. This unique arrangement of light source and optics allows retroreflection (light reflected back towards the light source) and transmission to be measured as well.

The incident θ_i and ϕ_i angles are controlled mechanically by pivoting the light source arm at point C and the sample holder at point A, respectively. In our current prototype, the light source is moved by a computer-controlled motor during data collection, and the sample is rotated manually. Because the hemisphere of reflected directions is captured in a single image, data collection proceeds quite rapidly and a complete BRDF can be recorded in a few minutes, including time for manual rotation of the sample.

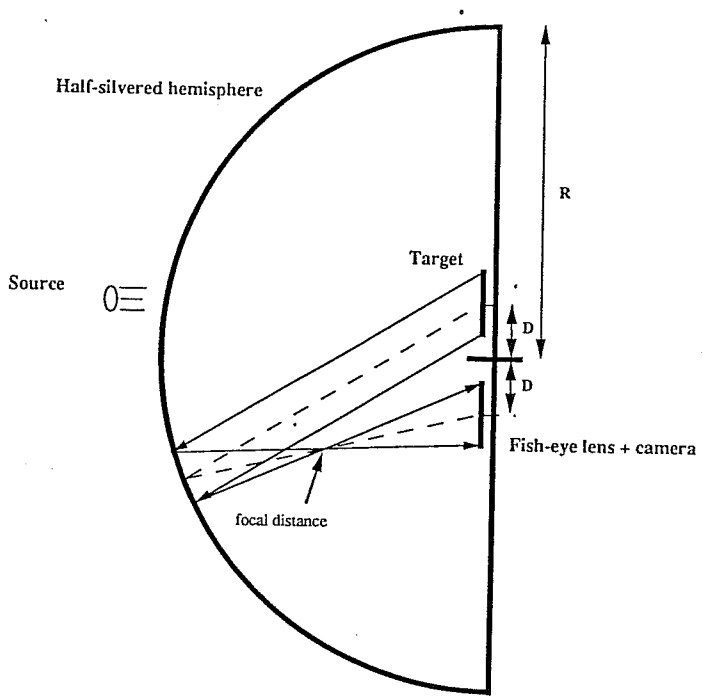

Figure 3. Imaging gonioreflectometer geometry. Light reflected by the sample in a specific direction is focused by the hemisphere or hemi-ellipsoid through a fisheye lens onto a CCD imaging array.

3.2. Calibration and Data Reduction

All measurements are made relative to a standard diffuse sample and a background measurement. The background measurement is made without any sample at all (using the dark room to simulate a black body) and is subtracted from the other measurements to reduce the effects of stray and ambient light.
The standard sample measurement is used as a basis for obtaining absolute reflectance values using the following simple equation at each image point:

$$f_{sample} = \frac{f_{measured} - f_{background}}{f_{standard} - f_{background}} \cdot \frac{\rho_{standard}}{\pi} \tag{2}$$

where:

$\rho_{standard}$ is the total reflectance of the standard sample

The ability to measure absolute BRDF values directly is an important feature of the imaging gonioreflectometer. Most other devices rely on auxiliary measurements of directional reflectance (ie. total reflectance for light incident at some θ_i, ϕ_i) and numerical integration to arrive at absolute quantities.

Recovering the reflected angles from pixel locations in the captured image is accomplished in two steps. The first step is to determine the mapping from image point locations to the lens incident direction. This is a function of the particular fisheye lens used, the camera, and the video capture board. Since it varies so much from one implementation to the next and the mapping is easily measured, we will not discuss it any further here. The second step is to compute the target reflection angles from these camera incident angles.

Figure 3 shows the geometry involved, and after a bit of trigonometry one can derive the following approximation:

$$r_c = D\,\sin\phi_c\,\sin\theta_c + \sqrt{D^2\sin^2\phi_c\,\sin^2\theta_c + R^2 - D^2}$$

$$\theta_r = \cos^{-1}\left[\frac{r_c\,\cos\theta_c}{\sqrt{r_c^2\cos^2\phi_c\,\sin^2\theta_c + (r_c\,\sin\phi_c\,\sin\theta_c - 2D)^2 + r_c^2\cos^2\theta_c}}\right] \tag{3}$$

$$\phi_r = \tan^{-1}\left[r_c\,\sin\phi_c\,\sin\theta_c - 2Dr_c\,\cos\phi_c\,\sin\theta_c\right]$$

where:

θ_r is polar angle relative to target

ϕ_r is azimuthal angle relative to target, right is 0°

θ_c is polar angle relative to camera

ϕ_c is azimuthal camera angle, right is 0°

R is radius of sphere or approximate radius of ellipsoid

D is one half the separation between target and camera centers

r_c is an intermediate result which is the distance from camera to reflector

notes:

The arctangent in the above equation should be computed using the signs of the numerator and denominator to get a range of 360°. Many math libraries provide a function named atan2 for this purpose.

The above equations are a good approximation both for hemispherical and hemi-ellipsoidal reflectors as long as D is small in relation to R.

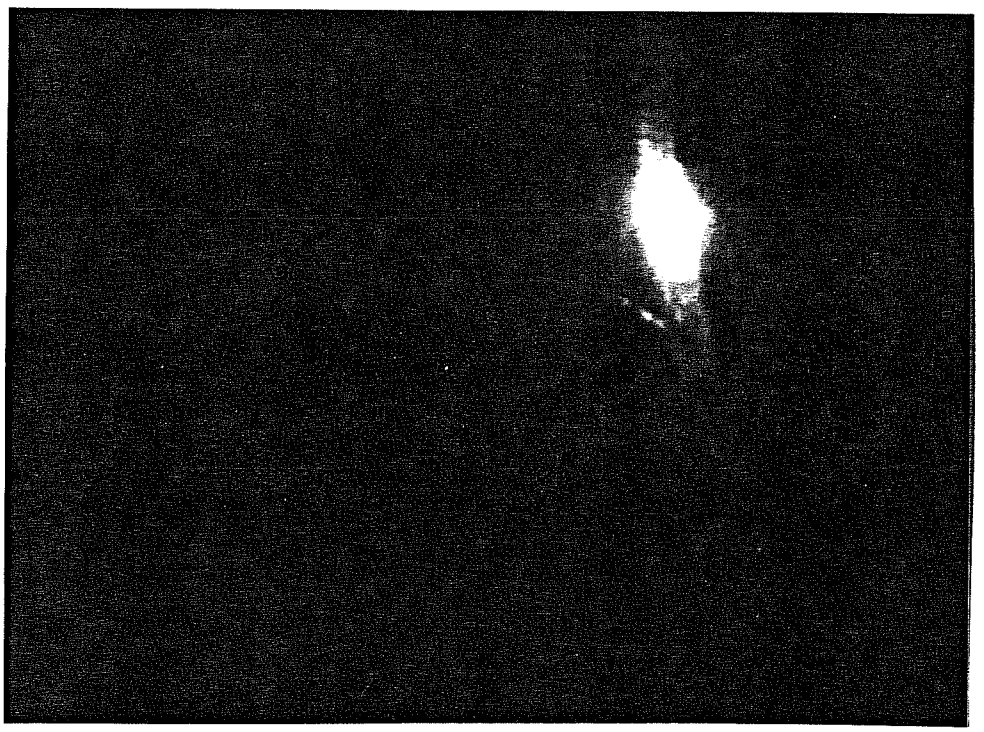

Figure 4. An image captured by the gonioreflectometer from an unfinished aluminum sample.

The image captured by our gonioreflectometer for a piece of unfinished aluminum illuminated at $(\theta_i,\phi_i)=(30°,0°)$ is shown in Figure 4. Although the image was reduced before data reduction to a resolution of 108 by 80 pixels, there is still much more information than is needed for an accurate lighting simulation. Also, since two or more f-stops may be used to capture the full dynamic range of the BRDF, there is often redundant information where the useful ranges of exposures overlap. We therefore apply a program to eliminate crowding of data points and insure that the peak is recorded at a high enough angular resolution while the rest of the usable distribution is recorded at a uniform density. The resulting point plot for this surface is shown in Figure 5. A data fitting program can then be used to match the reduced data set to a specific reflectance model.

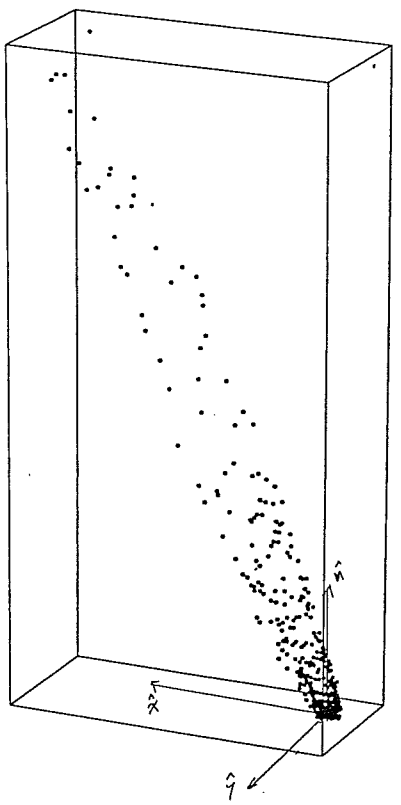

Figure 5. A 3-dimensional scatter plot of BRDF values corresponding to the image shown in Figure 4. Visualization of such data is improved by animation or stereograms.

4. Modeling Anisotropic Reflectance

Armed with a device that can measure anisotropic reflectance economically, we need a mathematical model that can be fit to our newfound data. Using the data directly is impractical because it requires too much memory, and oftentimes the data is noisy and not complete enough to cover the entire domain of the BRDF. We could represent the BRDF as a sum of terms in a spherical harmonic series, but this would still be expensive in terms of CPU time and of memory [Cabral87][Sillion91]. We would prefer a model that fits the data with as few parameters as possible. Ideally, these parameters would be either physically derived or meaningful so that they could be set manually in the absence of any data at all.

Many models have been suggested for isotropic reflection, but only a few models have been published for the more general anisotropic case. Kajiya published a fairly robust method for deriving BRDFs of metals from surface microstructure [Kajiya85]. However, his approach is not amenable to fitting measured reflectance data because the parameter space is too large (ie. all possible surface microstructures) and the BRDFs take too long to compute. Poulin and Fournier developed a model based on cylindrical scratches that is better suited [Poulin90], but their model restricted to a specific microstructure with cross-sectional uniformity, and its evaluation still seems unduly expensive for the results obtained.

Our goal is to fit our measured reflectance data with the simplest empirical formula that will do the job. If we can develop a model with physically meaningful parameters without adding undue complexity, so much the better.

4.1. The Isotropic Gaussian Model

The Gaussian distribution has shown up repeatedly in theoretical formulations of reflectance [Beckmann63][Torrance67][Cook82], and arises from certain minimal assumptions about the statistics of a surface height function. It is usually preceded by a Fresnel factor and geometrical attenuation factors, and often by an arbitrary constant. Since the geometric attenuation factors are typically difficult to integrate and tend to counteract the Fresnel factor anyway, we have replaced all of these coefficients with a single normalization factor that simply insures the distribution will integrate easily and predictably over the hemisphere.

$$f_{iso}(\theta_i,\phi_i;\theta_r,\phi_r) = \frac{\rho_d}{\pi} + \rho_s \cdot \frac{1}{\cos\theta_i \cos\theta_r} \cdot \frac{\exp[-\tan^2\delta/\alpha^2]}{4\pi\alpha^2} \tag{4}$$

where:

ρ_d is the diffuse reflectance

ρ_s is the specular reflectance

δ is the angle between vectors $\hat{n}$ and $\hat{h}$ shown in Figure 6

α is the standard deviation (RMS) of the surface slope

notes:

The ρ values may have some spectral dependence, and this dependence may vary as a function of angle so long as $\rho_d + \rho_s < 1$.

The normalization factor $\frac{1}{4\pi\alpha^2}$ is accurate as long as α is not much greater than 1, when the surface becomes mostly diffuse anyway.

The main difference between this isotropic Gaussian reflectance model and that of Phong is its physical validity. For example, most Phong implementations do not have the necessary bidirectional characteristics that constitute a valid BRDF model. It is clear by inspection that the above formula is symmetric with respect to its incident and reflected angles. Without this symmetry, a BRDF model cannot possibly be physical because the simulated surface reflect light differently in one direction than the other, which is forbidden by natural law. Also, without proper normalization, a reflectance model does not yield correct energy balance and thus cannot produce physically meaningful results. Even the model introduced recently by He

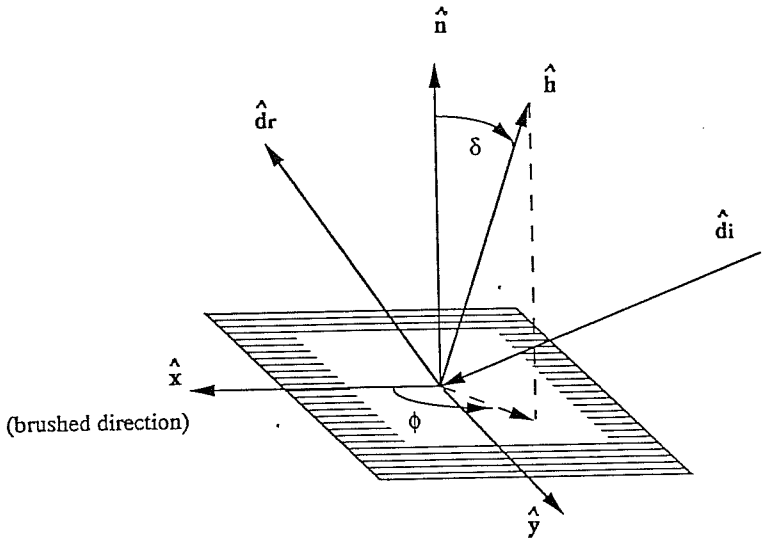

Figure 6. Angles and vectors used in reflection equations. The incident light arrives along vector $\hat{d}_i$ and is measured or simulated in direction $\hat{d}_r$. The polar angle between the half vector $\hat{h}$ and the surface normal $\hat{n}$ is δ. The azimuthal angle of $\hat{h}$ from the brushed direction $\hat{x}$ is ϕ.

et al [He91] with its rigorous physical derivation does not seem to pay close enough attention to normalization. Specifically, the so-called ambient term in the He-Torrance model is added without regard to the overall reflectance of the material, which by nature of the model is very difficult to compute. Comparisons were not made in He's paper between the reflectance model and absolute BRDF measurements (the data was scaled to match the function), thus normalization was not even demonstrated empirically. The fact that normalization was not adequately treated in He's otherwise impeccable derivation shows just how much normalization is overlooked and undervalued in reflectance modeling. The simplicity of the model presented here is what allows us to incorporate built-in normalization and other desirable features, such as fast evaluation and uniformly weighted sampling.

4.2. The Anisotropic (Elliptical) Gaussian Model

It is relatively simple to extend the Gaussian reflectance model to surfaces with two perpendicular (uncorrelated) slope distributions, α_x and α_y.
The normalized distribution is as follows:

$$f(\theta_i,\phi_i;\theta_r,\phi_r) = \frac{\rho_d}{\pi} + \rho_s \cdot \frac{1}{\cos\theta_i \cos\theta_r} \cdot \frac{\exp[-\tan^2\delta\,(\cos^2\phi/\alpha_x^2 + \sin^2\phi/\alpha_y^2)]}{4\pi\alpha_x\alpha_y} \tag{5a}$$

where:

ρ_d is the diffuse reflectance

ρ_s is the specular reflectance

α_x is the standard deviation of the surface slope in the brushed direction

α_y is the standard deviation of the surface slope perpendicular to the brushed direction

δ is the angle between the half vector, $\hat{h}$ and the surface normal, $\hat{n}$.

ϕ is the azimuth angle of the half vector projected into the surface plane.

A computationally convenient approximation for f is:

$$f(\theta_i,\phi_i;\theta_r,\phi_r) = \frac{\rho_d}{\pi} + \rho_s \cdot \frac{1}{\cos\theta_i \cos\theta_r} \cdot \frac{1}{4\pi\alpha_x\alpha_y} \exp\left[-2\frac{\left[\frac{\hat{h}\cdot\hat{x}}{\alpha_x}\right]^2 + \left[\frac{\hat{h}\cdot\hat{y}}{\alpha_y}\right]^2}{1+\hat{h}\cdot\hat{n}}\right] \quad (5b)$$

where:

$$\hat{h}\cdot\hat{x} = \frac{\sin\theta_r\cos\phi_r + \sin\theta_i\cos\phi_i}{||\vec{h}||}$$

$$\hat{h}\cdot\hat{y} = \frac{\sin\theta_r\sin\phi_r + \sin\theta_i\sin\phi_i}{||\vec{h}||}$$

$$\hat{h}\cdot\hat{n} = \frac{\cos\theta_r + \cos\theta_i}{||\vec{h}||}$$

$$||\vec{h}|| = \left[2 + 2\sin\theta_r\sin\theta_i(\cos\phi_r\cos\phi_i + \sin\phi_r\sin\phi_i) + 2\cos\theta_r\cos\theta_i\right]^{1/2}$$

For vector calculations, the following substitutions are used:

$$\vec{h} = \hat{d}_r + \hat{d}_i$$

$$\hat{h} = \frac{\vec{h}}{||\vec{h}||}$$

$$\cos\theta_r = \hat{d}_r\cdot\hat{n}$$

$$\cos\theta_i = \hat{d}_i\cdot\hat{n}$$

$\hat{d}_r$ is the reflected ray direction (away from surface)

$\hat{d}_i$ is the incident ray direction (away from surface)

$\hat{x}$ is a unit vector in the surface plane in brushed direction

$\hat{y}$ is a unit vector in the surface plane perpendicular to the brushed direction

The orientation vector $\hat{x}$ can be computed by:

$$\hat{x} = \frac{\hat{n}\times\vec{\psi}}{||\hat{n}\times\vec{\psi}||}$$

where:

$\vec{\psi}$ is any vector in the plane perpendicular to the brushed direction

As in the isotropic case, the normalization of the above anisotropic model is such that the total surface reflectance will equal the diffuse reflectance coefficient, ρ_d, plus the "semispecular" or "directional-diffuse" coefficient, ρ_s. The two other model parameters, α_x and α_y, represent the standard deviation of the surface slope in each of two perpendicular directions. Thus, all four of the model's parameters have physical meaning and can be set independently of measured data to produce a valid reflectance function. As long as the total reflectance $\rho_d + \rho_s$ is less than 1 and the two α's are not too large, Equation 5 will always yield a physically valid reflectance model.

The elliptical nature of the model arises from these two perpendicular slope distributions, and is apparent in the exponent of Equation 5a. A modified elliptical Phong reflection model was developed by Ohira and described by Yokoi and Toriwaki [Yokoi88], but this model is neither physically motivated nor properly normalized. Yokoi and Toriwaki also referred to an elliptical model developed by Takagi et al that took a better start from Blinn's model [Blinn77], but since Blinn's model is not normalized, it is probable that their model is not normalized, either. By starting with a normalized function, it is much easier to fit the model parameters to physical measurements as well as other specifications (such as appearance).

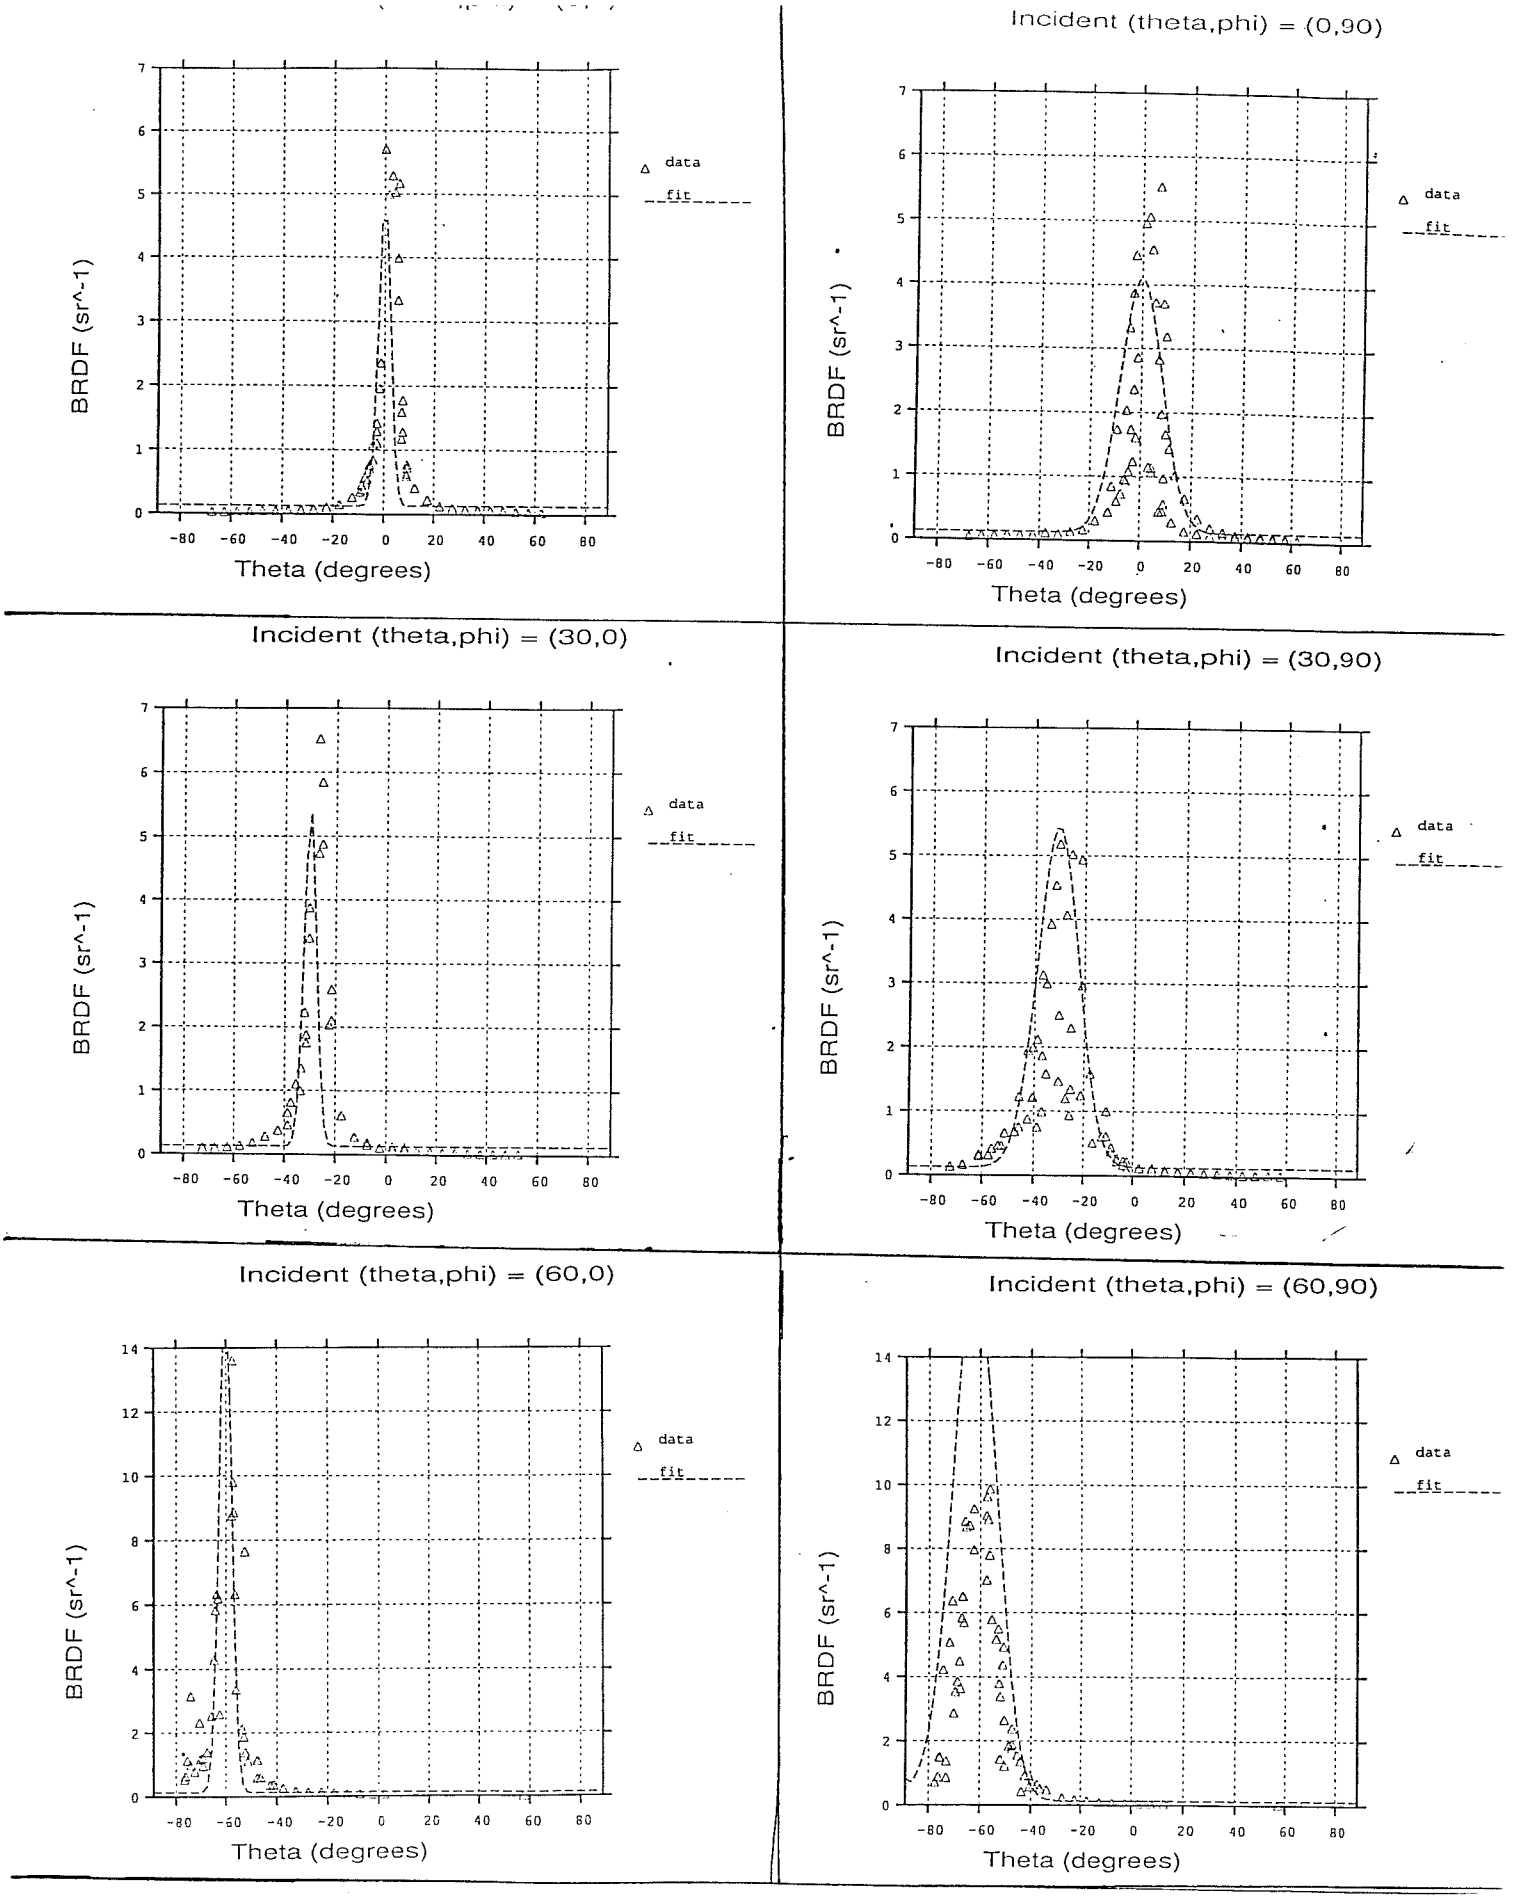

Figure 7. Measured data and elliptical Gaussian fit for unfinished aluminum. Unfinished aluminum exhibits anisotropy from rolling during its manufacture.

This simple four parameter model fits well the data we have gathered from anisotropic surfaces such as varnished wood and unfinished (rolled) or brushed metals. Because of its simplicity, it is easy to apply a least squares error minimization method to fit a set of parameters to measured data automatically. Automatic data fitting is essential to the economic modeling of surface reflectance for any significant database of materials. Figure 7 shows an example fit to the BRDF of an unfinished aluminum sample. Although the full hemisphere of reflected data was measured at 21 incident angles, it is difficult to visualize the 21 corresponding 3-dimensional point plots. We therefore present here only a slice of the data in the incident plane at 6 angles. The results section (6) lists the fitted parameters for this material as well as some other example surfaces.

5. Rendering Anisotropic Surfaces

The challenge to applying a new reflectance model to computer graphics is to approximate the luminance equation (1) in a manner that is unbiased and has low variance [Kajiya87]. Unfortunately, unbiased techniques (ie. pure Monte Carlo) tend to have high variance, while low variance approaches (ie. closed-form approximations) tend to be biased. To satisfy these conflicting requirements, we use a hybrid deterministic and stochastic ray tracing technique [Cook84]. A strictly deterministic calculation of the highlight contribution of sources, such as the widely used Whitted approximation [Whitted80], fails to pick up indirect semispecular contributions as demonstrated in Figure 8a. Conversely, using a stochastic sampling technique, such as that described by Cook [Cook86], causes the highlights from sources to show high variance in the form of excessive noise, even with 16 samples per pixel (Figure 8b). By combining the two techniques, using a deterministic solution for source contributions and a stochastic sampling for indirect contributions, we get a clean result without compromising accuracy. Figure 8c was calculated using the hybrid technique and the same number of samples as Figure 8b. Both figures took approximately 8 hours to compute on a Macintosh IIfx. Figure 8a took less than 15 minutes to compute since no sampling was required.

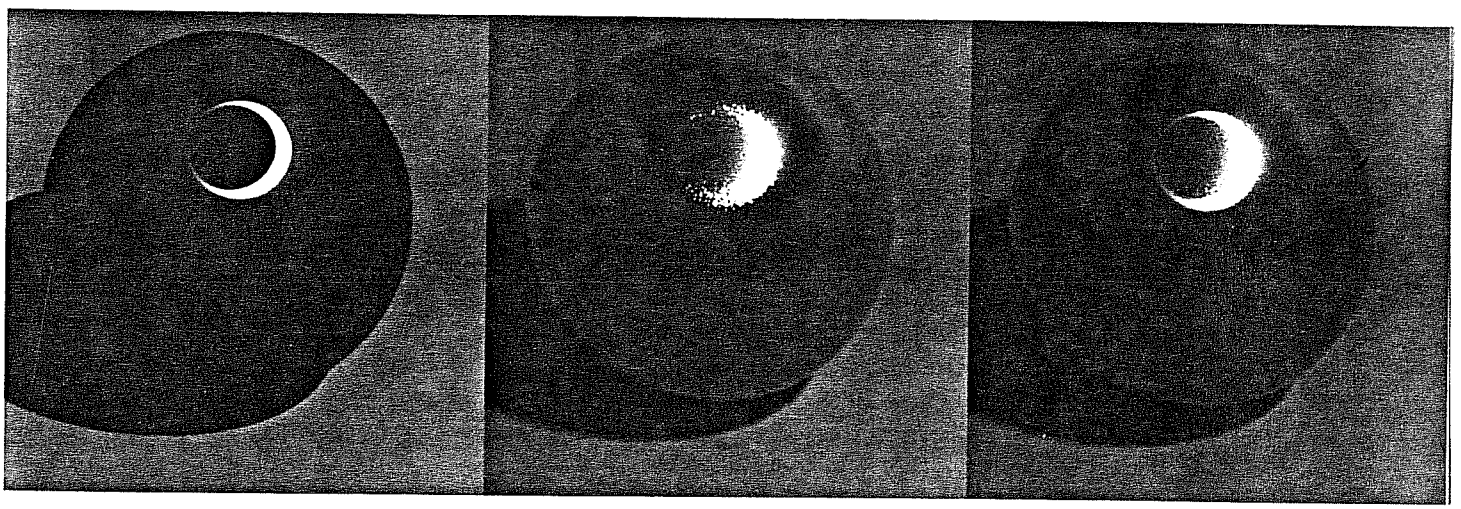

Figure 8a, 8b, 8c. Alternative rendering techniques for anisotropic reflection. 8a left shows deterministic technique with no sampling. 8b center shows strict Monte Carlo sampling approach. 8c right shows hybrid deterministic and Monte Carlo method.

The hybrid approach reduces to the following equation:

$$L(\theta_r,\phi_r) = I\,\frac{\rho_d}{\pi} + L_s\,\rho_s + \sum_{i=1}^{N} L_i\,\omega_i\,\cos\theta_i f(\theta_i,\phi_i;\theta_r,\phi_r) \tag{6}$$

where:

I is the indirect irradiance at this point (a constant ambient level or the result of a diffuse interreflection or radiosity calculation)

L_s is the radiance value in the Monte Carlo sample direction given in Equation 7 below

L_i is the radiance of light source i

ω_i is the solid angle (in steradians) of light source i

N is the number of light sources

f is the elliptical Gaussian function defined in Equation 5 above

In applying this technique, it is very important not to bias the sample by overcounting the specular component. Bias is easily avoided by associating a flag with the stochastically sampled specular ray. If the ray hits a light source whose contribution is being included in a closed form calculation, then the ray is not counted. Few rays are wasted in this way, since light sources occupy a small amount of the visual space in most scenes.

5.1. Stochastic Sampling of Elliptical Gaussian

Because of its simplicity, the elliptical Gaussian model adapts easily to stochastic sampling techniques. Using standard Monte Carlo integral conversion methods [Rubenstein81], we can write the following formulas for obtaining uniformly weighted sample directions for each L_s ray in Equation 6:

$$\phi(u_1) = \tan^{-1}\left[\frac{\alpha_y}{\alpha_x}\tan(2\pi u_1)\right] \tag{7a}$$

$$\delta(u_2) = \left[\frac{-\log(u_2)}{\cos^2\phi/\alpha_x^2 + \sin^2\phi/\alpha_y^2}\right]^{1/2} \tag{7b}$$

where:

ϕ, δ are the angles appearing in Equation 5a

u_1, u_2 are uniform random variables in the range (0,1]

notes:

The tangent and arctangent in the Equation 7a should be computed carefully so as to keep the angle in its starting quadrant.

Uniformly weighted sample rays sent according to the above distribution will correctly reproduce the specified highlight. This is much more efficient than either distributing the samples evenly and then weighting the result, or using other techniques, such as rejection sampling, to arrive at the correct scattering. Readers familiar with Monte Carlo sampling techniques will immediately appreciate the advantage of having a formula for the sample point locations -- something that is impossible with more complicated reflectance models such as He-Torrance.

6. Results

The following table gives an abbreviated list of surfaces and their elliptical Gaussian fits. Color was not measured for any of the surfaces, but was matched by eye for the varnished plywood. Neither the white ceramic tile nor the glossy grey paper exhibited much anisotropy in their reflectance function, so they could just as well be approximated using the isotropic Gaussian model (4).

Material	ρ_d	ρ_s	α_x	α_y
unfinished aluminum	.4	.15	.03	.1
lightly brushed aluminum	.67	.19	.088	.13
heavily brushed aluminum	.73	.26	.045	.73
varnished plywood	(.7,.2,.05)	.025	.04	.11
white ceramic tile	.71	.031	.064	.071
glossy grey paper	.34	.053	.074	.081

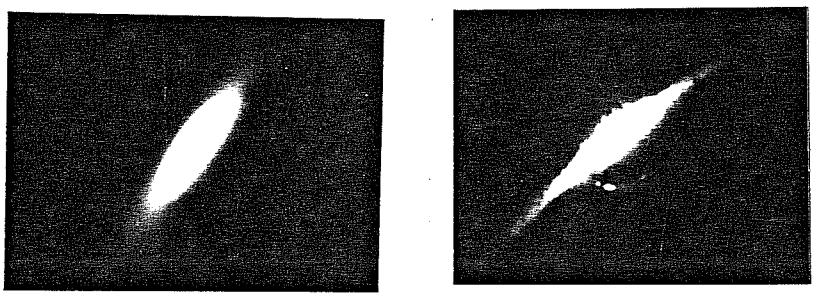

Figure 9a, 9b. Comparison between simulated and photographed highlight from the unfinished aluminum sample featured in Section 3.2. The simulation uses a fit to the elliptical Gaussian distribution described in Section 4.2.

Figure 9a shows a picture of the unfinished aluminum sample mentioned earlier with a small source arranged to produce a characteristic highlight. Figure 9b shows an elliptical Gaussian simulation of the material created using parameters matched to data measured by the imaging gonioreflectometer.

Figure 10a shows a photograph of a child's varnished wood chair with a small desk lamp immediately behind and above it. This arrangement results in a large anisotropic highlight in the seat of the chair. Figure 10b shows the closest simulation possible using a deterministic isotropic reflection model. Figure 10c shows a hybrid simulation with the elliptical Gaussian model. Notice how the hybrid rendering technique reproduces not only the highlight from the light source, but also the semispecular reflection from the back wall in the seat of the chair.

Figure 11 shows a table with anisotropic reflections in the wood varnish and the two candle holders. The lid of the silver box shown is also anisotropic, and demonstrates the use of local control to affect the reflectance properties of an anisotropic surface. A wave function determines the orientation of the brushed direction in the box lid, producing characteristic highlights. There are four low level light sources in the scene, the two candles on the table, an overhead light source above and to the right, and the moon shining in through a window. The moon is the source of the large white highlight on the left. Some sampling noise is evident from low probability specular interreflections, but this is the price we must pay for an unbiased solution. This image took approximately 24 hours to compute on a Sun 4/260 at a resolution of 1024x696 pixels and 9 samples per pixel.

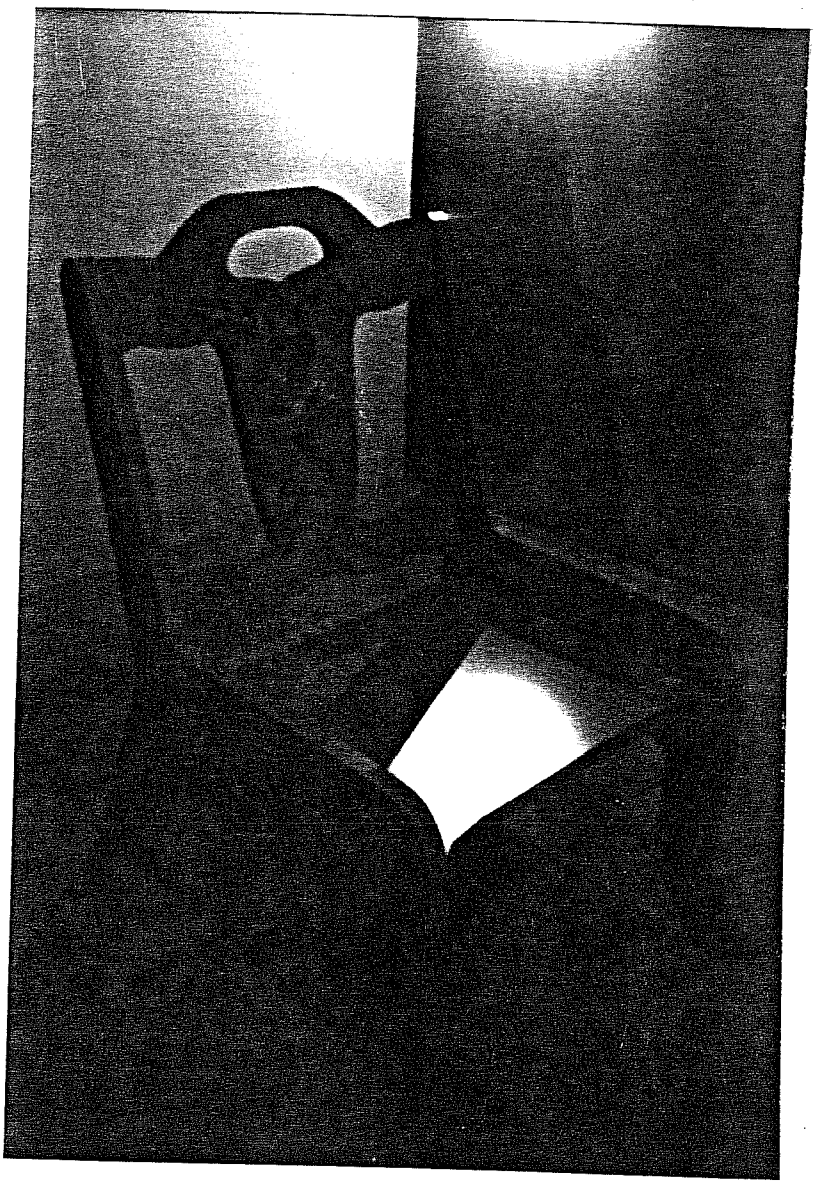

Figure 10a. A photograph of a child's varnished and painted chair.

Figure 10b. A simulation of the chair using a deterministic calculation with the isotropic Gaussian model given in Section 4.1.

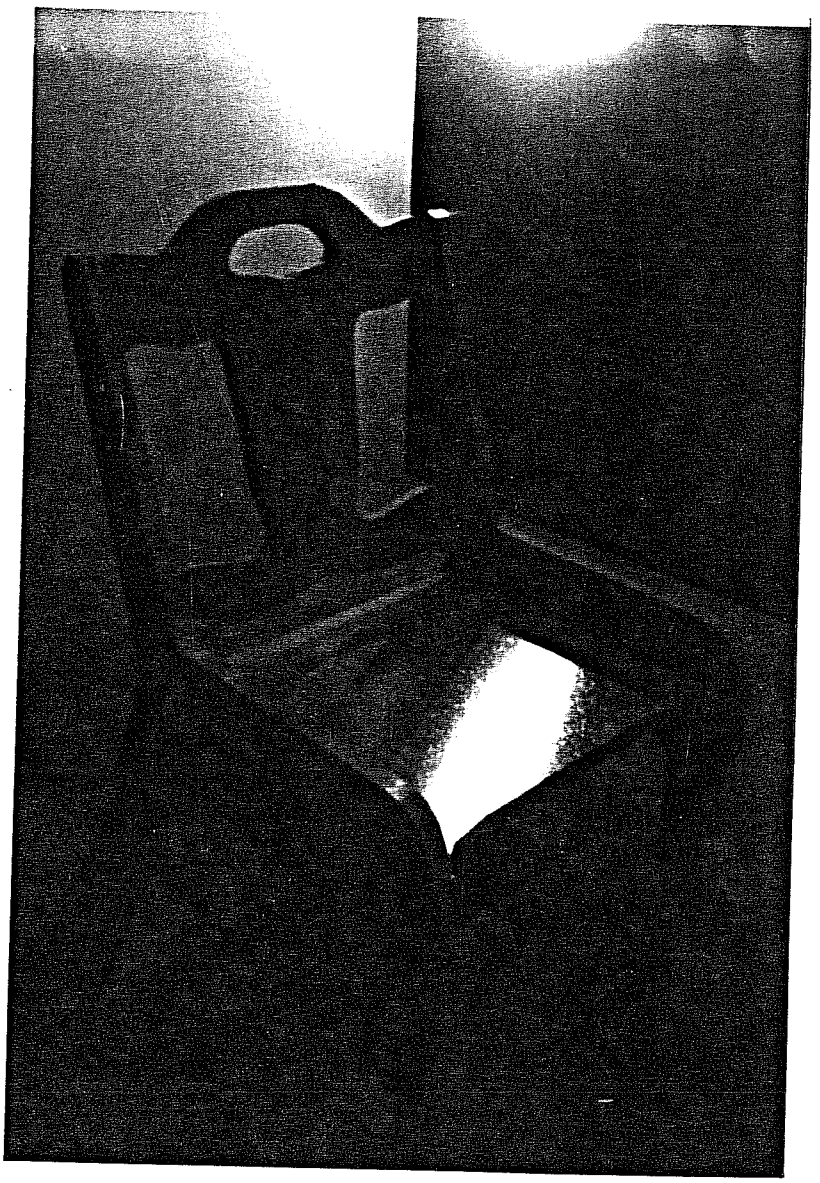

Figure 10c. A simulation of the chair using the elliptical Gaussian model from Section 4.2 and the hybrid rendering technique described in Section 5.

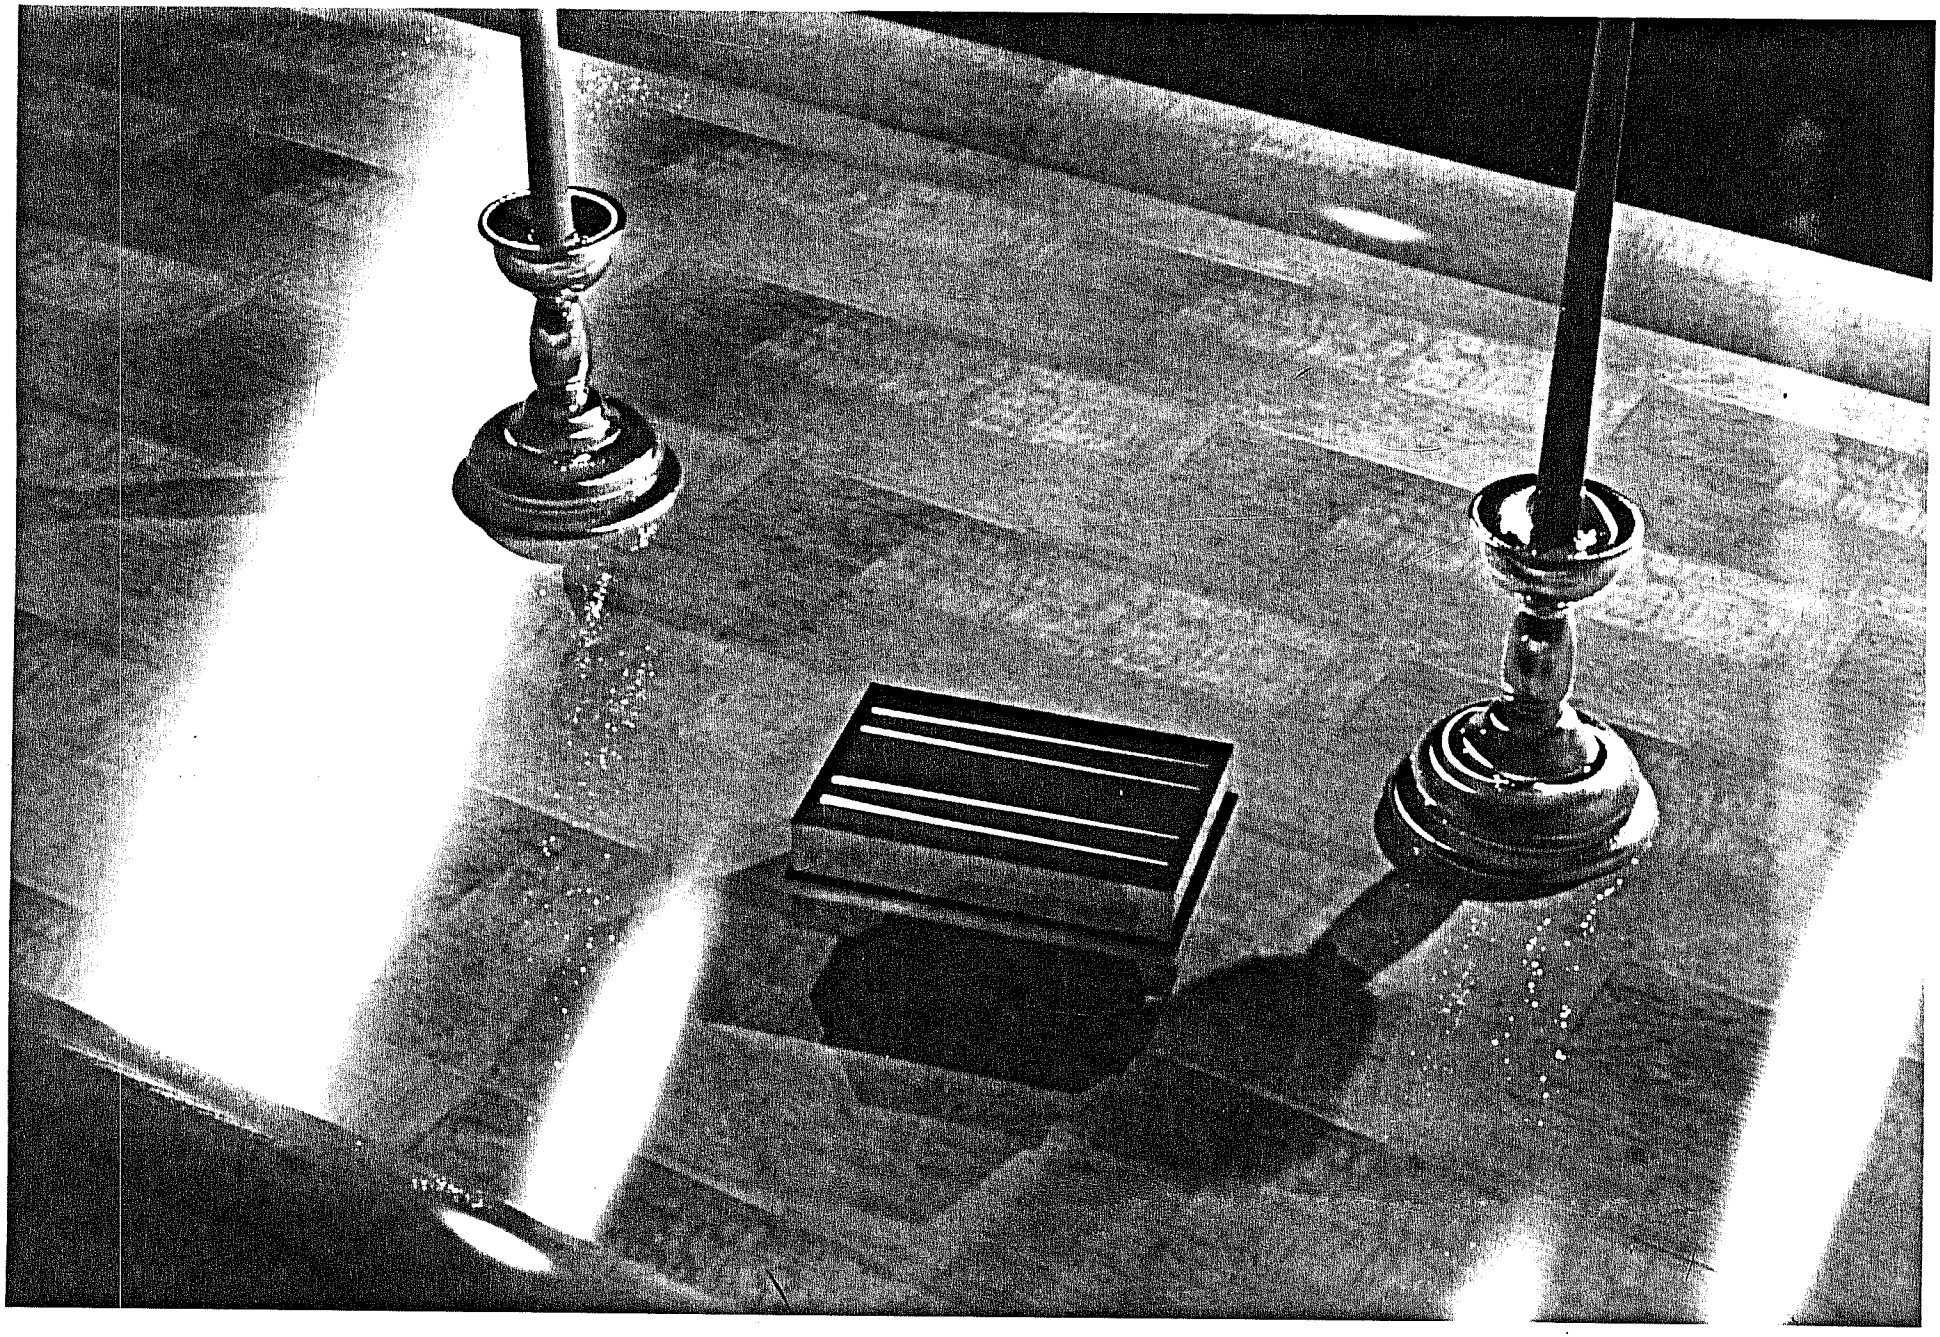

Figure 11. A table scene with anisòtropic reflection in metallic and varnished wood surfaces.

7. Conclusion

We have presented a new imaging gonioreflectometer for measuring anisotropic (as well as isotropic) BRDFs that is both practical and economical. It is possible with this device to quantify the reflectance functions of many architectural materials, thus opening up a whole new world of accuracy and realism in lighting simulation and rendering. Our prototype device has some limitations in its ability to measure grazing angles and highly specular materials, but these limitations are not inherent in the design and will be worked out in future implementations.

We have also presented a formulation of the Gaussian reflection model for anisotropic materials that nicely fits data we have gathered with the new device. It is difficult to justify going to a sophisticated model when a simple model fits the data and provides other advantages as well. It is easy and fast to fit data automatically, and easy to implement efficient software and hardware rendering solutions. Another advantage of our particular empirical model is that its four fitting parameters can be set in an intuitive fashion without regard to any reflectance measurements, and the results will still be properly normalized and physically meaningful.

8. Acknowledgements

This work was supported in part by a grant from Apple Computer Corporation, and by the Assistant Secretary for Conservation and Renewable Energy, Office of Building Energy Research and Development, Buildings Equipment Division of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

Ken Turkowski of Apple's Advanced Technology Division was jointly responsible for supervision of this project. Francis Rubinstein, Rudy Verderber and Sam Berman each gave invaluable support, and Robert Clear had a direct hand in overseeing the calibration and validation of the camera and reflectometer equipment.

Anat Grynberg participated in the initial design and construction of the imaging gonioreflectometer, and Lisa Stewart made numerous equipment modifications during the calibration stage and took most of the BRDF measurements used in this paper.

Carol Stieger did the tole painting on the real and simulated versions of the child's chair.

9. References

[Beckmann63]
Beckmann, Petr, Andre Spizzichino, *The Scattering of Electromagnetic Waves from Rough Surfaces*, Pergamon Press, NY, 1963.

[Blinn77]
Blinn, James F., "Models of Light Reflection for Computer Synthesized Pictures," *Computer Graphics*, Vol. 11, No. 2, July 1977.

[Cabral87]
Cabral, Brian, Nelson Max, Rebecca Springmeyer, "Bidirectional Reflection from Surface Bump Maps," *Computer Graphics*, Vol. 21, No. 4, July 1987.

[Cook82]
Cook, Robert L., Kenneth E. Torrance, "A Reflectance Model for Computer Graphics," *Computer Graphics*, Vol. 15, No. 3, August 1981.

[Cook84]
Cook, Robert L., Thomas Porter, Loren Carpenter, "Distributed Ray Tracing," *Computer Graphics*, Vol. 18, No. 3, July 1984.

[Cook86]
Cook, Robert L., "Stochastic Sampling in Computer Graphics," *ACM Transactions on Graphics*,

Vol. 5, No. 1, January 1986.

[He91]
He, X., K.E. Torrance, F.X. Sillion, D.P. Greenburg, "A Comprehensive Physical Model for Light Reflection," *Computer Graphics*, Vol. 25, No. 4, July 1991.

[Kajiya85]
Kajiya, James T., "Anisotropic Reflection Models," *Computer Graphics*, Vol. 19, No. 3, July 1985.

[Kajiya87]
Kajiya, James T., "The Rendering Equation," *Computer Graphics*, Vol. 20, No. 4, August 1986.

[Murray-Coleman90]
Murray-Coleman, J.F., A.M. Smith, "The Automated Measurement of BRDFs and their Application to Luminaire Modeling," *Journal of the Illuminating Engineering Society*, Winter 1990.

[Nicodemus77]
Nicodemus, F.E., J.C. Richmond, J.J. Hsia, *Geometrical Considerations and Nomenclature for Reflectance*, U.S. Department of Commerce, National Bureau of Standards, October 1977.

[Phong75]
Phong, B., "Illumination for Computer Generated Pictures," *Communications of the ACM*, Vol. 18, No. 6, June 1975.

[Poulin90]
Poulin, Pierre, Alain Fournier, "A Model for Anisotropic Reflection," *Computer Graphics*, Vol. 24, No. 4, August 1990.

[Rubenstein81]
Rubenstein, R.Y., *Simulation and the Monte Carlo Method*, J. Wiley, New York, 1981.

[Sandford85]
Sandford, Brian P., David C. Robertson, "Infrared Reflectance Properties of Aircraft Paints," *Proceedings IRIS Targets, Backgrounds, and Discrimination*, 1985.

[Sillion91]
Sillion, Francois, James Arvo, Donald Greenburg, "A Global Illumination Solution for General Reflectance Distributions," *Computer Graphics*, Vol. 25, No. 4, July 1991.

[Torrance67]
Torrance, K.E., E.M. Sparrow, "Theory for Off-Specular Reflection from Roughened Surfaces," *Journal of the Optical Society of America*, Vol. 57, No. 9, September 1967.

[Yokoi88]
Yokoi, Shigeki, Jun-ichiro Toriwaki, "Realistic Expression of Solids with Feeling of Materials," *JARECT*, Vol. 18, 1988.