# The Spherical Target Gonioreflectometer: *A Proposal*

Robert R. Lewis
bobl@tricity.wsu.edu

School of Electrical Engineering and Computer Science
Washington State University at Tri-Cities
Richland, Washington 99325
USA

31 March, 1998

## Abstract

We discuss the design of a proposed instrument called a *spherical target gonioreflectometer* to measure the reflective properties of materials. By computing the image geometry of an illuminated sphere made of or coated with reflective material, we show that we can generate a sufficient amount of data to fill the parameter space needed to create bidirectional reflectance distribution functions (BRDFs) for real substances.

## 1 Introduction

In computer graphics, the fundamental equation governing reflectance (cf. Cohen and Wallace [2]) is usually taken to be

$$L_r(\omega_r) = L_e(\omega_r) + \int_{\Omega^R_{\mathbf{N}}} f_r(\omega_i', \omega_r)\, L_i(\omega_i')\, (\mathbf{N} \cdot \omega_i')\, d\omega_i' \quad (1)$$

where $L_r$ is the total radiance given off of a surface with normal $\mathbf{N}$, $L_e$ is the surface emissivity, $L_i$ is the incident radiance, $\Omega^R_{\mathbf{N}}$ is the reflection hemisphere (contains the surface normal $\mathbf{N}$, the reflected direction $\omega_r$, and the incident direction $\omega_i'$), $f_r$ is the bidirectional reflectance distribution function (hereafter, BRDF), $d\omega_i'$ is $\sin\theta_i' d\theta_i' d\phi_i'$, where $\theta_i'$ is the incident polar angle and $\phi_i'$ is the incident azimuthal angle. Figure 1 shows this geometry.

Gonioreflectometers are devices used to measure the reflective properties of materials. More specificially, they measure BRDFs. The configuration of a classic gonioreflectometer is quite similar to Figure 1: a sample of the material to be measured is placed where $dA$ is located and gantries hold a luminaire and photometer along directions $\omega_i'$ and $\omega_r$, respectively. The altitude and azimuth of each arm, which may be robotically controlled, allow a total of four degrees of freedom, the minimum necessary to measure an anisotropic BRDF. An automated version of such a system is under construction by Hanrahan and Levoy at the Stanford Computer Graphics Laboratory in connection with their work on light field measurement [10].

The gonioreflectometer in use at Cornell, described by Greenberg, et al. [6], places the sample on a pivot with two degrees of freedom. This allows the luminaire and photometer (a monochromator in this case) to lie in the plane. Revolving the luminaire around the sample adds the additional degree of freedom needed to measure isotropic BRDFs. With such features as liquid-cooled 12-bit analog-to-digital converters, this device offers great precision, but at a relatively high price.

Larson (ne' Ward) [12] (also [13]) presented an innovative approach which used a half-silvered plastic hemisphere with an arm-controlled luminaire and a CCD camera with a fisheye lens to collect reflectance data "in parallel": every CCD element received reflectance data from a slightly different geometry. The two-dimensional pixel array thus provided two degrees of freedom and the arm an additional one.

All of these systems are expensive and involve custom-built hardware. The system we will propose in this paper is relatively inexpensive and can, we believe, be built with off-the-shelf hardware. By providing a low-cost alternative, it is our hope to enable the capture and sharing of large amounts of reflectance data for incorporation into photorealistic rendering.

Figure 1: Surface Geometry for Reflection.

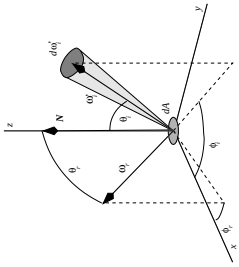


Figure 2: Sphere Illumination Geometry

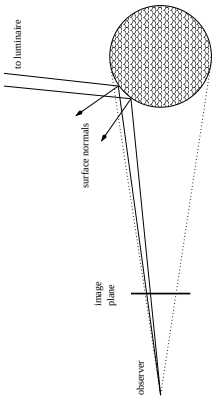


## 2 Directional Luminaires and the BRDF

Consider an object illuminated by a directional luminaire[1] coming from direction $\omega_i$. The incident radiance may then be represented with a Dirac $\delta$-function:

$$L_i = E_N \delta(\omega_i - \omega_i'). \tag{2}$$

where the constant $E_N$ is the "normal irradiance", the amount of power per unit area transferred to a surface with normal incidence[2].

Assuming that the object is non-emissive, (1) becomes

$$L_r(\omega_r) = E_N \, f_r(\omega_i, \omega_r) \, (\mathbf{N} \cdot \omega_i) \tag{3}$$

so, assuming a non-participating medium so that $L_{\text{pixel}} = L_r(\omega_r)$, pixel values are geometrically related to BRDF values. If we choose a shape whose geometry is known and can relate individual pixels to that geometry, we can determine the BRDF. We refer to this phenomenon as "shading from shape".

[1] We present the case for a directional (i.e., collimated) luminaire, but a point luminaire would work as well, introducing the slight complication of inverse-square behavior.

[2] $E_N$ is the value of the irradiance $E$ when the light is coming from the normal direction $\mathbf{N}$.

## 3 Reflection on Spherical Targets

Let us consider a very simple object geometry: spheres. Dating back at least as far as Blinn [1], synthetically illuminated spheres have been used to illustrate illumination models. The material editor of at least one major modelling and animation system, Alias|Wavefront, draws such spheres for the user to illustrate the effects of user-specified diffusivity and specularity parameters.

There is an intuitive notion that an illuminated sphere is a good exemplar of reflective properties. Let us examine this

how complete our coverage is, let us consider the experimental setup shown in Figure 3. Measured from the center of the sphere, there is an angle $\Theta$ between the camera and luminaire directions.

We then consider a camera resolution of $32 \times 32$ pixels and assume that the projection of the sphere's diameter subtends the image area. Most cameras are capable of much greater resolution than this, so our coverage here is guaranteed to be a lower limit.

Figure 4 shows the result. The first three plots are for individual $\Theta$ values of $30°$, $60°$, and $90°$, respectively. The fourth plot is a composite of data from $\Theta = 0°$ to $\Theta = 140°$ in increments of $10°$.

Since not all pixel rays intersect the sphere, not all intersection points are illuminated, and there is vertical symmetry, the actual number of points generated for a single value of $\Theta$ varies from 406 for $\Theta = 0°$ to 4 for $\Theta = 140°$. Nevertheless, it is evident from Figure 4 that there are no major gaps in parameter space for a spherical target.

The only possible undersampling problems might be at the extrema of $\theta_r$ and $\phi_r$, but these will be improved considerably at more realistic image resolutions.

Figure 4: Parametric Coverage. These were generated assuming a 32 × 32 pixel array.

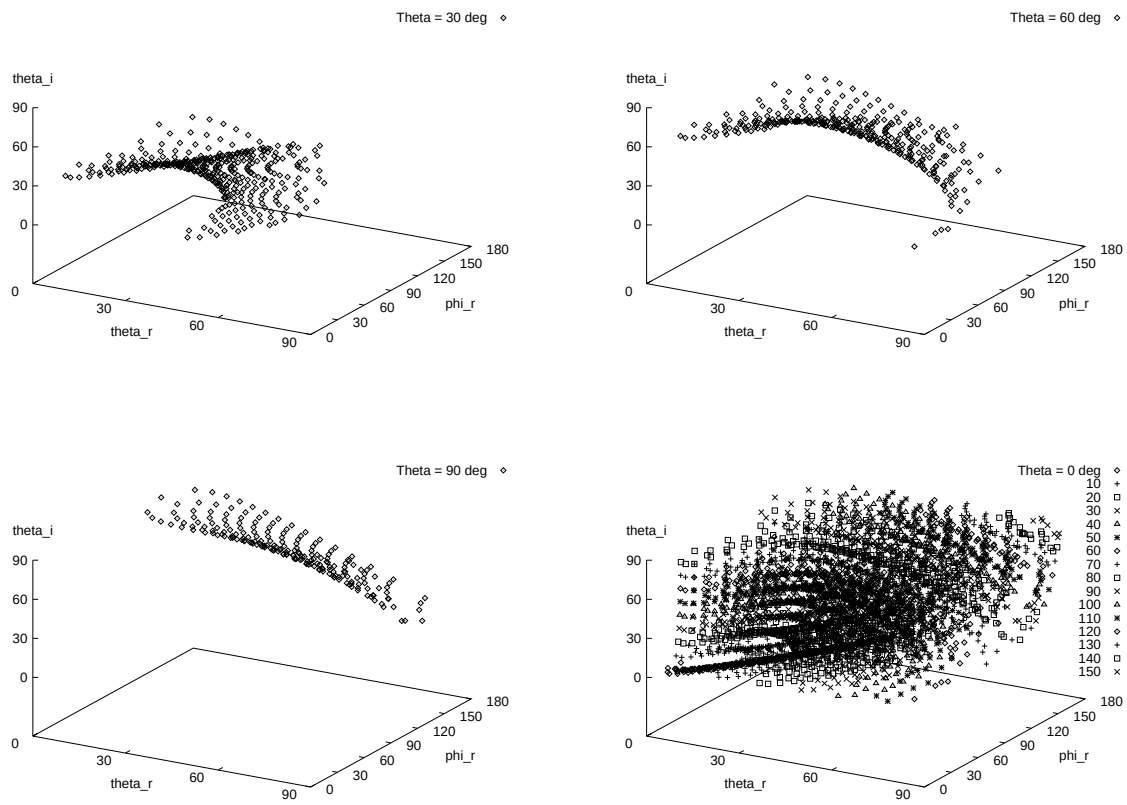


# 5 Implementation Considerations

We will refer to the (so far) hypothetical device built along the lines of Figure 3 as a *spherical target gonioreflectometer* or STG. There are three parts: a luminaire, a target, and a camera. We also need to calibrate the instrument.

## 5.1 Luminaire

As mentioned above, the two overall choices for luminaire are point and directional. One very inexpensive way to provide a point source would be a bare halogen desk lamp bulb. While this would require us to make the inverse square law modifications noted above, being able to vary the incident radiance by an easily-determined factor may have advantages in calibration.

The easiest way to provide a directional source would be with a collimator, as is used in spectroscopy. We would need to verify that the collimated light field was uniform and not subject to effects such as vignetting over the illuminated area.

In order to vary $\Theta$, either the luminaire or the camera would have to be moveable in a 180° arc surrounding the target, and the luminaire is likely to be the lighter of the two. Later phases of STG may automate this positioning.

## 5.2 Target

A major constraint of the STG is the need to guarantee the spherical geometry of the target. Targets should either be

isotropic substances that can be formed into spheres, such as metal, plastic, stone, or ceramic, or substances that can be uniformly applied to a sphere, such as paint, wax, or sandblasting.

In addition, at this stage of development, the STG requires targets that are spatially uniform on the scale of a projected pixel. It would work with a sandstone sphere, but not with a granite or marble one.

Both isotropicity and spatial uniformity can be guaranteed empirically by collecting data for several orientations of the target. The resulting BRDFs should be identical within experimental tolerances.

### 5.3 Camera

Raw STG data is a set of single frames, one for each $\Theta$ value. A conventional video stream is far more data than we need, but the individual frames are relatively low quality, noisy, and of poor contrast (specular BRDFs especially have high dynamic ranges).

Another approach would be to capture data with a still photographic camera and digitize the negatives with PhotoCD™ technology. This would, however, require us to calibrate out the film transfer function and introduce a delay for film processing.

There is an intriguing third alternative. In recent years, CCD technology has been adapted for use by amateur astronomers. This (comparatively) mass market has brought down the price considerably. For under US$2,000, it is now possible to purchase a CCD camera for astronomical use (see the review by Horne [7]) with a 336 by 224 pixel CCD array whose elements have 16 bits of resolution. This includes a filter wheel. Attached to a telescope or telephoto lens and connected to a PC to manage data capture, such a device should be quite appropriate for the STG. By the use of filters, short exposures, and stopping down, a camera intended for the observation of faint astronomical objects should work adequately in a laboratory setting.

### 5.4 Calibration

The STG needs to be calibrated geometrically and radiometrically. Geometric calibration relates camera pixel coordinates to coordinates in the idealized image plane of Figure 2. There is a substantial literature on this subject. Two apparently successful approaches cited by Debevec, et al. [4] are Tsai [11] and Faugeras and Toscani [5].

There are two degrees of radiometric calibration: absolute and relative. Absolute calibration relates pixel values to absolute radiances in physical units. Relative calibration relates pixel values to each other, deriving BRDFs within a multiplicative constant. The latter is usually adequate for computer graphics purposes – most people who do rendering are comfortable with a single scaling factor for each BRDF. As mentioned above, the prospective camera for the STG has 16 bits of precision, but should this prove inadequate, the techniques presented in Debevec and Malik [3] describe how to combine multiple exposures to extend camera dynamic range. Absolute calibration may even be possible if we are willing to obtain an appropriately-calibrated luminaire.

## 6 Data Representation

Once we have captured the (non-uniformly spaced) BRDF data, in order to render with it we need to interpolate it for any incident and reflected geometry. There has been a considerable amount of work on this in recent years. Larson [12] fit Gaussian functions to his measured data. Lalonde and Fournier [8] fit that same data with wavelet basis functions. Another promising approach that we hope to apply to this data is Lee, et al.'s [9] multilevel B-spline algorithm.

## 7 Work in Progress

We are currently in the process of getting funding for a prototype STG as an initial step at establishing an image-based rendering laboratory at the Tri-Cities campus of Washington State University.

If the prototype is successful, it could conceivably be extended to include the following:

- *anisotropic materials*
  Capturing all four incident and reflected angles would require an additional degree of freedom, which could be provided easily by putting the target on a turntable. We need to investigate, though, if this would provide complete parametric coverage.
- *non-spherical targets*

The STG could be adapted to work with flat samples of materials which could be cut into polygons and assembled into polyhedra. (Of course, we would then have to call the instrument a "PTG".) This would make specification of the tangent geometry needed for anisotropic materials much easier.

- *textures*
  Having polygonal facets on the PTG would also allow us to incorpoate the spatial variation of a BRDF, otherwise known as surface texture, into our measurements. This would, of course, increase our data storage requirements.

- *BTDFs*
  If we could place a point light source within a spherical cavity *inside* a target sphere or polyhedron, it should be possible to make what, to the best of our knowledge, would be the first empirical measurements of a bidirectional transmittance distribution function (BTDF).

## References


[1] James F. Blinn. Models of light reflection for computer synthesized pictures. In James George, editor, *Computer Graphics (SIGGRAPH '77 Proceedings)*, volume 11, pages 192–198, San Jose, California, July 1977.

[2] Michael F. Cohen and John R. Wallace. *Radiosity and Realistic Image Synthesis*. Academic Press Professional, San Diego, CA, 1993.

[3] Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 369–378. ACM SIGGRAPH, Addison Wesley, August 1997. ISBN 0-89791-896-7.

[4] Paul E. Debevec, Camillo J. Taylor, and Jitendra Malik. Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 11–20. ACM SIGGRAPH, Addison Wesley, August 1996. held in New Orleans, Louisiana, 04-09 August 1996.

[5] Olivier Faugeras and Giorgio Toscani. The calibration problem for stereo. In *Proceedings IEEE CVPR 86*, pages 15–20, 1986.

[6] Donald P. Greenberg, Kenneth E. Torrance, Peter Shirley, James Arvo, James A. Ferwerda, Sumanta Pattanaik, Eric P. F. Lafortune, Bruce Walter, Sing-Choong Foo, and Ben Trumbore. A framework for realistic image synthesis. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 477–494. ACM SIGGRAPH, Addison Wesley, August 1997. ISBN 0-89791-896-7.

[7] Johnny Horne. Celestron's pixcel 255 ccd camera. *Sky & Telescope*, 94(4):53–57, October 1997.

[8] Paul Lalonde and Alain Fournier. Filtered local shading in the wavelet domain. In Julie Dorsey and Philipp Slusallek, editors, *Eurographics Rendering Workshop 1997*, pages 163–174, New York City, NY, June 1997. Eurographics, Springer Wein. ISBN 3-211-83001-4.

[9] Seungyong Lee, George Wohlberg, and Sung Yong Shin. Scattered data interpolation with multilevel b-splines. *IEEE Transactions on Visualization and Computer Graphics*, 3(3):228–244, July-September 1997.

[10] Marc Levoy and Pat Hanrahan. Light field rendering. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 31–42. ACM SIGGRAPH, Addison Wesley, August 1996. held in New Orleans, Louisiana, 04-09 August 1996.

[11] Roger Tsai. A versatile camera calibration technique for high accuracy 3d machine vision using off-the-shelf tv cameras and lenses. *IEEE Journal of Robotics and Computer Vision*, 3(4):323–344, August 1987.

[12] Gregory J. Ward. Measuring and modeling anisotropic reflection. In Edwin E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 265–272, July 1992.

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