Measuring Bidirectional Texture Reflectance with a Kaleidoscope

Jefferson Y. Han   Ken Perlin†

Media Research Laboratory
Department of Computer Science
New York University

Abstract

We describe a new technique for measuring the bidirectional texture function (BTF) of a surface that requires no mechanical movement, can measure surfaces *in situ* under arbitrary lighting conditions, and can be made small, portable and inexpensive. The enabling innovation is the use of a tapered kaleidoscope, which allows a camera to view the same surface sample simultaneously from many directions. Similarly, the surface can be simultaneously illuminated from many directions, using only a single structured light source. We describe the techniques of construction and measurement, and we show experimental results.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism - color, shading, shadowing, texture I.4.8 [Image Processing and Computer Vision]: Scene Analysis - color, shading

Keywords: BRDF, BTF, reflection models, image-based rendering, realistic image synthesis

1 Introduction

We introduce a new technique for measuring the appearance of a textured surface as it is viewed or illuminated from different directions. Previous such techniques have been somewhat cumbersome, relying on multiple successive measurements and requiring mechanical repositioning between each measurement. The new technique requires no mechanical movements and is very simply calibrated. Further, it has the benefit that it could be applied *in situ* directly to textured surfaces in their natural settings, under any lighting conditions.

This paper is structured as follows. First we describe previous techniques, and then we introduce our new method. After this we devote the bulk of the paper to practical and experimental issues, and to showing of experimental results, since the main focus of our paper is the introduction of an enabling technique, not theory. Our hope in disseminating this work is that this new method of measuring texture reflectance will make it easier for other researchers to produce new and better data-driven surface models.

†Email: {jhan, perlin}@mrl.nyu.edu

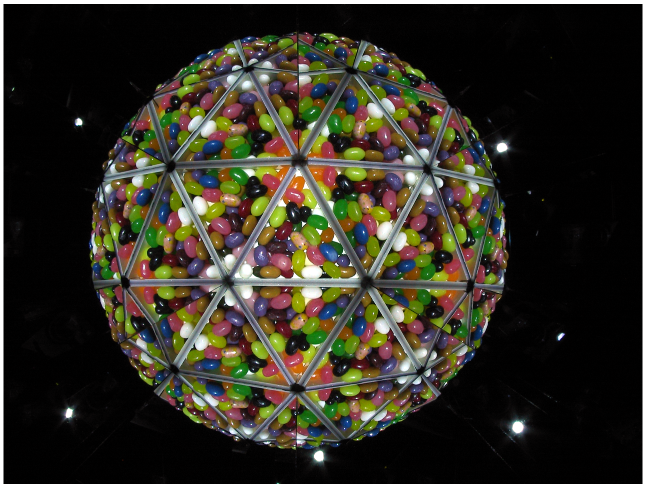

Figure 1: A view through a kaleidoscope

2 Background

Much recent work in realistic image synthesis has focused on the use of actual data measurements of real-world surfaces and materials, both in the search for better data-driven reflectance models, and for direct use in image-based rendering techniques.

The reflectance properties of a surface can be characterized by its Bidirectional Reflectance Distribution Function (BRDF) [Nicodemus et al. 1977], the four dimensional function that describes how much light from any incident direction (θ_i, ϕ_i) is transferred to any exitant direction (θ_e, ϕ_e):

$$BRDF(\theta_i, \phi_i, \theta_e, \phi_e)$$

The field is quite mature in techniques for measuring BRDFs, and for representing them accurately and compactly. Real world surfaces, however, are not perfectly homogeneous- they exhibit local variations in microgeometry and in reflectance, which are not adequately represented by a single BRDF.

Dana et al. [1999] define the Bidirectional Texture Function (BTF) as the six dimensional function which extends the BRDF by allowing reflectance to vary spatially along the surface, parameterized by *(u,v)*:

$$BTF(u, v, \theta_i, \phi_i, \theta_e, \phi_e)$$

This representation is able to effectively capture the various subtleties of complexly textured surfaces, particularly those exhibiting such phenomena as self-occlusion and self-shadowing.

There have been recent advances in working with BTFs for realistic image synthesis. Because the BTF is a large unwieldy 6D function, it is difficult to obtain a dense sampling, and therefore current databases are relatively sparse. Yet recent successful research has shown that even a sparse sampling of the BTF can be adequate for rendering applications [Liu et al. 2001; Tong et al. 2002; Vasilescu and Terzopoulos 2003].

Increased quality of BTF sample data would also be of benefit to computer vision research. For example, algorithms that reconstruct geometry or motion from multiple views require correspondences to be found between these views. BTF data would allow robust testing of the identification of corresponding surface points, even as the appearance of each surface point varies with view angle. This data would also benefit shape-from-texture, texture segmentation, and texture recognition techniques.

3 Prior Work

Use of real-world reflectance is currently characterized by the difficulty of gathering the BRDF and the BTF, particularly due to the high dimensionality of this data.

3.1 BRDF Measurement

The straightforward approach to measuring the 4D BRDF is to mechanically position a light source and photometer around the hemisphere about the sample though the use of robotic armatures, as in [Murray-Coleman and Smith 1990]. Any such mechanical arrangement must have four degrees of freedom; data collection is tediously performed by sequentially stepping through each position.

Subsequent methods greatly improve the efficiency of data acquisition by reducing the number of mechanically scanned dimensions through the use of a 2D imaging element such as a CCD camera. Ward's LBL imaging gonioreflectometer [Ward 1992] uses a hemi-ellipsoidal mirror. A CCD camera equipped with a wide-angle-lens, and the surface sample are positioned at the mirror's two respective foci to effectively map pixel position to exitant angular position. This method requires mechanical repositioning of the light source. Also notable about Ward's device is that the mirror is semi-transparent, thereby permitting measurements when view and illumination angles are coincident. Others [Davis and Rawlings 1997; Mattison et al. 1998; Carter and Pleskot 1999] have thoroughly explored the various other possible arrangements of curved mirrors and beam splitters.

An alternative way to utilize an imaging element is to measure the BRDF on a *curved* sample. [Lu et al. 1998] arranges a sample patch onto a known cylinder. [Marschner et al. 1999] relaxes the sample geometry restriction by utilizing a range scanner, and improves acquisition flexibility by allowing for free positioning of the capture camera.

More recent work attempts to recover the BRDF from sampling environments that are even less structured. Boivin and Gagalowicz [2001] demonstrate recovering multiple BRDFs from a single photograph, with known geometry and light source positions. [Ramamoorth and Hanrahan 2001] describe a signal processing framework that generalizes the recovery of the BRDF under unknown lighting conditions.

3.2 BTF Measurement

The seminal work by Dana et al. on the BTF [1999] presents a 3DOF robotic system that incrementally tilts/rotates a patch of the sample in front of a light source. This method produces 205 total samples of the BTF, with a relatively even distribution of illumination directions, but, due to mechanical limitations, with a limited distribution of viewing angles. It also requires a sample patch of the surface to be affixed to the device, which makes *in situ* measurements impossible, particularly for skin.

Other research involving BTFs utilizes various other custom gantry rigs, such as that of [Furukawa et al. 2002], which uses 2 motorized concentric arcs carrying 6 cameras and 6 lights.

Later work by Dana [2001] introduces a BTF measurement device that utilizes a concave paraboloid mirror section, similar to that used in previous BRDF capture devices, but in concert with an aperture and a translation stage for the sample. Theoretically, this technique should be able to produce very high resolution sampling of the BTF in every dimension, with large flexibility in sample distribution, but at a slow capture rate. It also inherits the problems associated with the need to affix surface samples.

Note that this technique is representative of a general class of solutions to the BTF capture problem, which utilize a 4D BRDF measurement device, mechanically scanning the sample across the device to obtain the additional two dimensions.

3.3 Related Measurement

Other techniques measure that subset of the BTF for which the viewpoint is fixed, and only illumination is varied.

Debevec et al.'s "Light Stage" [2000], constructed to capture the complex reflectance of the human face, mechanically scans a directional light source at relatively high speeds through two degrees of freedom, capturing 64x32 illumination samples. Successive versions of the stage have replaced this single light source, first with a linear array of xenon strobes on a motorized arc, and then with a static 2D array of 156 LED clusters, allowing for the capture of subjects *in motion* under arbitrary illumination conditions [Debevec et al. 2002].

[Malzbender et al. 2001] describes a device for *in situ* surface reflectance measurement, wherein 50 inward-pointing light sources are distributed on a small, portable hemispherical frame, allowing for rapid automated acquisition. Polynomial curves are fitted to the lighting-dependent color at each pixel; these curves are used to generate images with novel lighting conditions that interpolate the light positions that were sampled.

4 Our Approach

Our measurement device is based on the principle of the kaleidoscope [Brewster 1819]. Generally used as a child's toy, a kaleidoscope is a hollow tube of polygonal cross-section, whose inner walls are lined with front-surface mirrors. Peering into a kaleidoscope creates an infinite "hall of mirrors" illusion; any surface sample placed at the far end will appear to "multiply" into many replicated images of itself.

A kaleidoscope can be tapered, so that its far end is smaller than its near end. When this is done, the surface sample at the far end will look like a faceted virtual sphere. This is because each successive reflection reorients the reflected image of the surface a little further away from the perpendicular, until eventually the reflected images disappear over the horizon of the sphere [Figure 1].

The effect is analogous to having an entire array of cameras all pointing toward the surface sample from different directions, which is precisely what is needed to measure the BTF. A single camera pointed at a

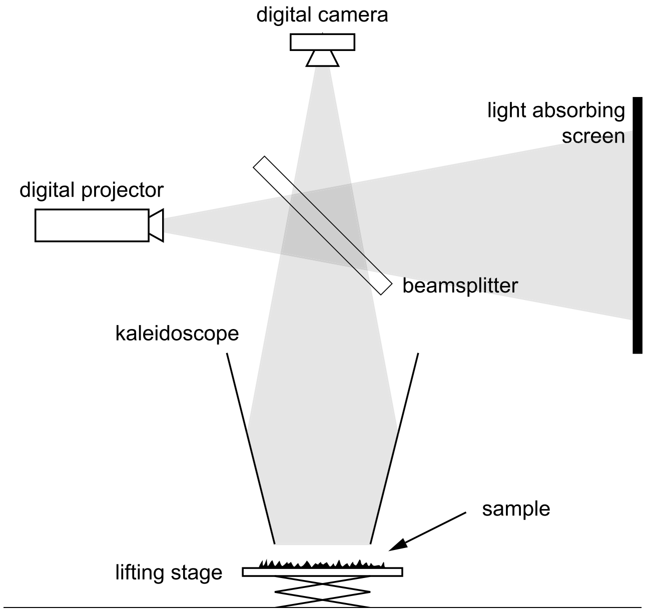

Figure 2: Schematic of optical components

surface sample which is on the far end of a tapered kaleidoscope will be able to see that same surface sample simultaneously from many different angles. These differently angled views of the surface sample appear to the camera as different facets of the virtual sphere.

4.2 Illumination

A nice benefit of this approach is that we can also use it as an illumination technique, using a single projector to illuminate the same surface sample from many different directions. When we point a projector down into the tapered kaleidoscope, different pixels of the projected image will arrive at the sample after having reflected off the kaleidoscope walls in different ways, and therefore will approach the sample from various directions. In effect, different regions of the projected image behave like separate light sources. By keeping only selected pixels of the projected image bright, we can choose a particular direction from which to illuminate the sample.

The optical paths of the camera and projector need to be merged together, so that both can be pointed down into the kaleidoscope. We do this through the use of a 45° beam splitter. Light from the projector reflects off this beam splitter down into the kaleidoscope. Light emerging back out of the kaleidoscope is transmitted through the beam splitter and is then captured by the camera. This arrangement allows the projected image to be coaxial with the image seen by the camera. Figure 2 shows an optical schematic of the device.

4.3 Procedure

Measurement of the surface BTF proceeds by taking a sequence of successive sub-measurements, one after the other. During each sub-measurement, exactly one region of the illumination image is bright, and all others are dark. Because each region of the illumination image corresponds to a unique sequence of reflections of light off of the kaleidoscope walls, that region will illuminate the surface sample from a unique sub-range of incoming light directions. A complete measurement consists of successive illumination of the sample surface by each of the illumination regions in turn.

4.4 Advantages

Our approach has a number of advantages in comparison to previous methods for measuring the BTF.

The new method requires no moving parts, allowing for full measurement to be performed very quickly. Since no physical movement is required between sub-measurements, all sub-measurements are guaranteed to be perfectly registered to one another. This property allows for a quite significant improvement in accuracy over previous methods.

The device can be used to measure surfaces *in situ*, under any lighting conditions, without relocating the sample from its native setting. For some site-specific surfaces, such as living human skin, methods in current use for measuring BTF are simply not viable, since they all require isolating a sample into a light-controlled environment. Also, approaches that require the sample to be physically repositioned between measurements cannot be used to measure loose samples such as rice, dirt or pebbles.

The new method requires only a single CCD camera or equivalent image capture device. This property allows the device to be fabricated at a low cost in comparison with previous methods that require multiple CCD cameras or equivalent image capture devices.

The new method richly samples the BTF. Even our first prototype captured 484 illumination/view angle pairs, which exceeds the 205 pairs captured by the technique of Dana et al. [1999].

The technique is also versatile enough to allow the device to be portable and hand-held; we discuss this further in the future work section.

All of these qualities make for a valuable new measurement tool, for use in situations for which current techniques are too bulky or unwieldy, or are simply impossible. For example, during a motion picture production, a member of the visual effects crew could use a device employing our method to measure the BTF of the skin of various parts of an actor's face, or the fabric of a costume or couch, or any prop or desk, wall, or floor surface of the set. With this information in hand, the appearance of these items can then be duplicated digitally with highly convincing realism and fidelity. Once the entire BTF has been captured, the filmmaker is free to make arbitrary decisions about lighting and camera placement, which the virtual objects can be synthesized to match.

5 Design Parameters

The kaleidoscope approach to BTF measurement is an extremely flexible one, with many design parameters to consider, depending on the objective.

5.1 Choice of the number of sides

In general, the kaleidoscope can be made as a regular polygon of n sides, for n>=3. We implemented a ray-tracer to better understand the effects of various values of n (see Figure 3).

It is apparent that not every virtual facet is complete; many are *fragmented*, appearing to have real and virtual mirror seams slicing through them. For simplicity we decided to consider only the *unfragmented* facets as usable data. As a result, the effect of n on fragmentation is a major factor in kaleidoscope design, since the proportion of these facets varies with n.

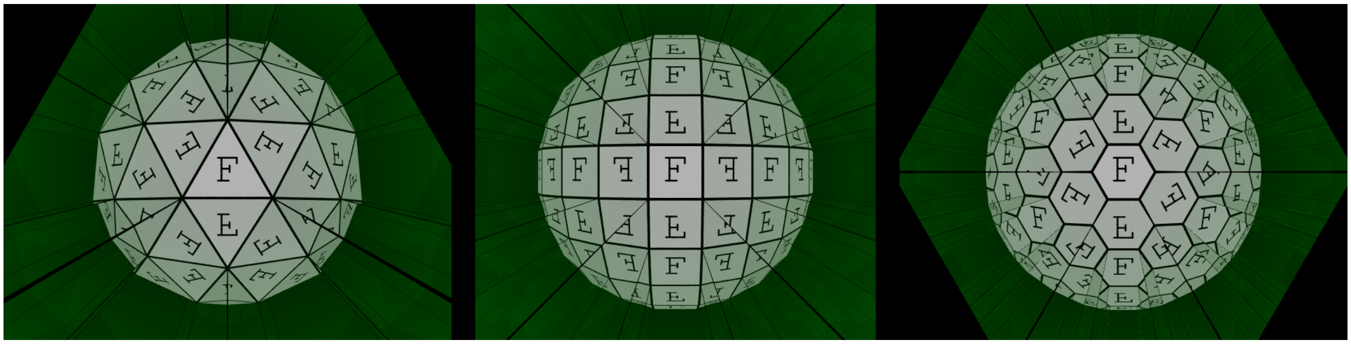

Figure 3: Kaleidoscope simulations for n = { 3, 4, 6 }

The value of *n* also directly determines the shape of the base as the regular *n*-gon. However, image processing is most easily performed on rectangular images, so for any *n*!=4, we only utilize the area of the largest inscribable square.

We ultimately focused on the triangular *n*=3 case, because of its simplicity in construction, and its highest proportion of whole unfragmented facets, though it does compromise on sample area and capture efficiency.

5.2 Choice of taper angle

Varying the angle of taper also significantly affects what is seen through the kaleidoscope. Angle of taper refers to the amount that the kaleidoscope narrows from one end to the other, and is defined as the tilt angle between the mirrored side and the kaleidoscope's optical axis.

A larger taper angle causes each successive reflection to tilt further away from the surface normal, which produces fewer facets that are visible before eventually disappearing over the horizon (elevation exceeds 90°). Conversely, a smaller angle of taper, forming a kaleidoscope with walls that are more parallel, produces a greater number of visible facets with finer angular steps. However, capturing a greater number of facets in a single view results in fewer pixels for each facet, and thus a reduction in spatial resolution.

We use kaleidoscopes with a relatively large angle of taper (and correspondingly fewer, larger facets) to capture relief surfaces with high self-shadowing, such as pebbles, cloth, and jellybeans. This optimizes for greater spatial resolution within the sample; the tradeoff is fewer different angular directions. We use tall slender kaleidoscopes with a smaller angle of taper (and correspondingly more numerous, smaller facets) to capture shiny surfaces with sharp specular peaks in reflectance.

We can calculate an optimal taper angle given a desired angular resolution, and desired final grazing angle.

Ultimately we chose a taper that tilts from vertical angle by 9°. This provides 4 orders of reflections to the horizon, a final grazing facet elevation angle of 76°, and 22 complete views of the surface sample, providing $22^2 = 484$ distinct view/illumination angle pairs. See Figure 4 for a tabulation of the actual angles of this design, along with a visualization of those spherical coordinates on the unit hemisphere.

5.3 Other physical dimensions

The remaining design parameter decisions consist of determining the scale of the kaleidoscope that will best: (i) accommodate a surface sample of a desired size, and (ii) work with a given camera field of view and projector field of view without the use of any additional lenses or optics.

Before constructing the device, we created a simple OpenGL-based visualization tool to balance the various interrelated parameters. This allowed us to vary, in simulation, taper angle, base patch size, kaleidoscope height, and field of view and distance of the camera and the projector.

At this stage we realized that for a given sample size and tilt angle (a smaller angle produces a larger virtual sphere), the height of the kaleidoscope (and therefore the bulk and expense of the front-surface mirrors) is determined by the field of view of the camera and projector: the kaleidoscope's height can be reduced if a wider field of view is used. The camera we used had a vertical field of view of 39°; the projector had a vertical field of view of 21°. The smaller of these (the projector) was the limiting factor, which ultimately determined the kaleidoscope height.

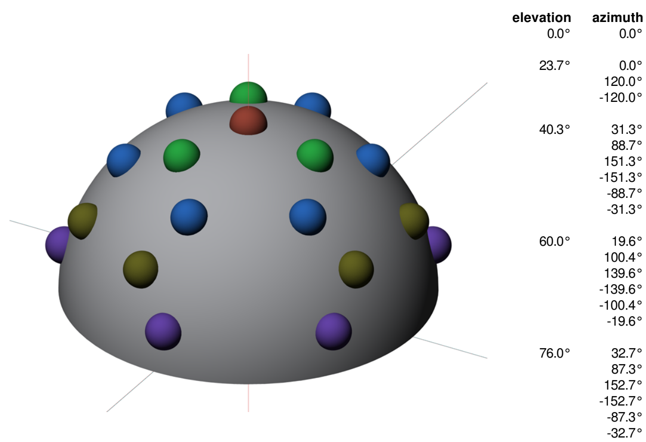

Figure 4: Distribution of viewpoint and illumination angles

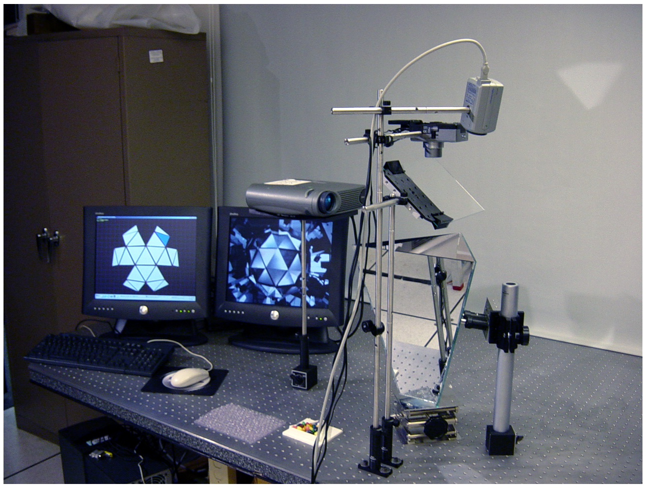

Figure 5: Experimental Setup

The kaleidoscope has a triangular base edge length of 4", providing a maximally inscribed active sample area of 2.3" square, and has a total height of 14.7". The three trapezoidal front-surface mirrors needed for this design were cut for us from standard stock by a professional stained glass cutter.

5.4 Miscellaneous Issues

For the beam splitter, we use an ordinary plate of glass, which has approximately 96% transmission at a 45° incident angle. Because the projector has a high luminance, this glass reflects more than sufficient illumination down into the kaleidoscope.

We found LCD projectors to be unsuitable for our purposes, because the reflectivity of the beam splitter varied with polarization. For this reason, our experiments were conducted with a DLP projector, which provides unpolarized illuminance. The camera was a Canon PowerShot G1, which has a capture resolution of 2048×1536. We maintained a small aperture so as to maximize depth of focus.

A large proportion of light was transmitted through the beam splitter, and ended up being projected onto the wall of our laboratory. Some of this reflected light made its way back to the beam splitter, and a small portion of that light was reflected up into the camera. We placed a matte black surface on the wall, which absorbed almost all of this unwanted light. The color calibration step compensated for what little was left.

To maintain precision in the experiments, it was important not to continually jar the kaleidoscope. For this reason, we installed the entire apparatus on an optical table. A sample to be measured was first slid underneath the kaleidoscope, upon a mechanical stage. The stage was then elevated until the sample was flush with the kaleidoscope opening. Our laboratory setup is shown in Figure 5.

6 Calibration

6.1 Radiometric Calibration

Deviations in brightness and color balance came from many sources, including mirror coatings, mirror absorption, and mismatch between the projector "white" color and the camera "white" color. In our measurements dichroic mirror coatings caused slight color shifts at different incident angles, which showed up as variations in hue between different facets of the virtual sphere.

There was also a dropoff per unit sample area at the outer facets, simply due to the fact that a tilted facet presents fewer pixels to the projector. An inherent shortcoming of our method is that both spatial resolution and brightness drop off at the most extreme angles.

To compensate for all of these deviations, as well as others we might not have accounted for, we calibrated the device *in situ* using a Kodak standard color chart. We needed to do this calibration only once, since the projector, camera, beam splitter and mirrors all remained unchanged. Were our experiments to take place over a long time frame, we expect that we would need to periodically recalibrate to account for gradual shifts in the projector lamp as it ages.

6.2 Sub-Image Registration and Calibration

We used image processing to identify and extract the many reflected images of the surface sample. This procedure needed to be performed only once, using the following *in situ* calibration:

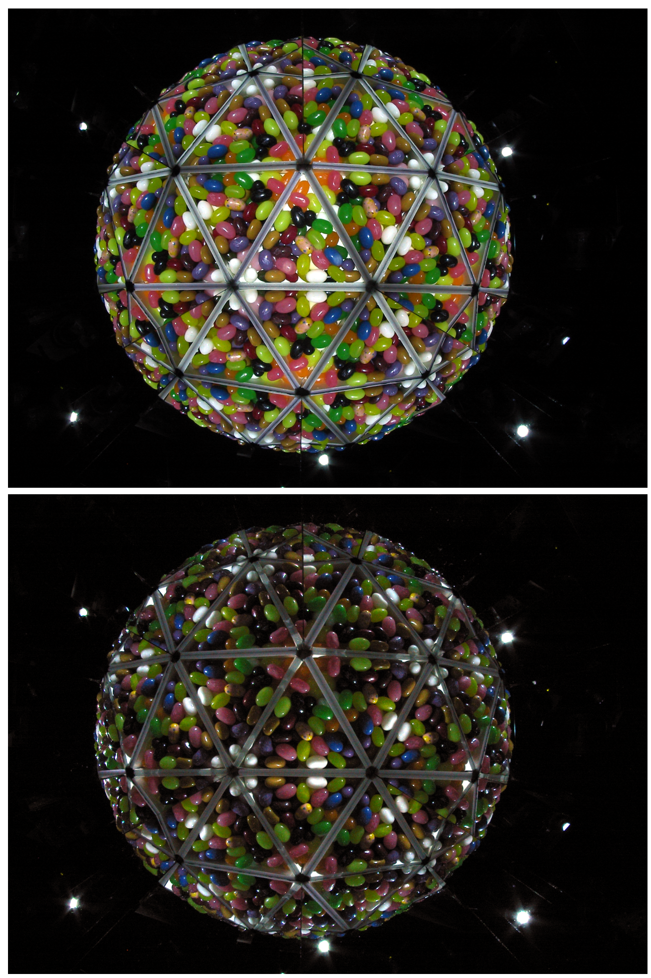

Figure 6: Two multi-view captures of "jellybeans", under different illumination directions

We placed a planar 3×3 checkerboard test pattern under the kaleidoscope and performed corner detection to identify the sub-pixel coordinates of each reflected checkerboard image. Those points were used to compute the best homography transform that maps each patch to the unit square.

Those transformations were in turn applied to each of the 22 illumination imaging shots. The resulting 22 square sub-images were each clipped out, and saved to disk. The result was a 22×22 array of images indexed by projector facet and camera facet.

Correction for the lens distortion of the camera needed to be done only once, using the technique of [Zhang 1999].

6.3 Illumination Alignment

We needed to determine which pixels in the projected image illuminated each kaleidoscopically reflected image of the surface sample. For purposes of the current work, we did this manually, implementing a triangle editor in software. Using the actual image from a video camera peering into the kaleidoscope as a guide, this editor allowed a user to quickly outline each of the 22 triangles.

Ideally this step should be done automatically as follows: The projector would project a known tracking pattern, which the camera would record. This data would then be used to recover, in a single step, the projection matrix of the projector itself, as well as all the projection matrices of all the reflected images of the surface sample. This calibration also would need to be performed only once.

7 Experimental Results

Figure 6 shows two multi-view image captures of a sample of jellybeans, taken with two different illumination angles, and Figure 9 shows the full, structured 484 image BTF after sub-image extraction has been performed.

Figure 7 shows a multi-view image of a coin captured with a kaleidoscope having a relatively small taper. This device has 79 unfragmented facets, and can capture $79^2 = 6241$ distinct view/illumination angle pairs. A small-taper kaleidoscope is particularly useful for measuring fine variations in reflectance due to small differences in angle.

7.1 BSSTF trials

For surfaces which have appreciable sub-surface scattering, it is useful to measure the BSSRDF (Bidirectional Scattering Surface Reflectance Distribution Function) of the surface by illuminating only a small spot of the surface sample, and then to measure the light which emerges from locations within the larger region that surrounds this spot [Jensen et al. 2001]. By incrementally *moving* this illuminated spot and taking associated measurements at each successive spot position, we can measure what can be termed the sample's BSSTF (Bidirectional Scattering Surface Texture Function):

$$BSSTF(u_i, v_i, u_e, v_e, \theta_i, \phi_i, \theta_e, \phi_e)$$

The BSSTF, also described as the *reflectance field* in [Debevec et al. 2001], is an eight dimensional function: two for the entry point of the light into the sample, two for the exit point of the light out of the sample, two for incoming spherical angle, and two for outgoing spherical angle. Because our technique requires no physical movement, it is now feasible to accumulate the many measurements

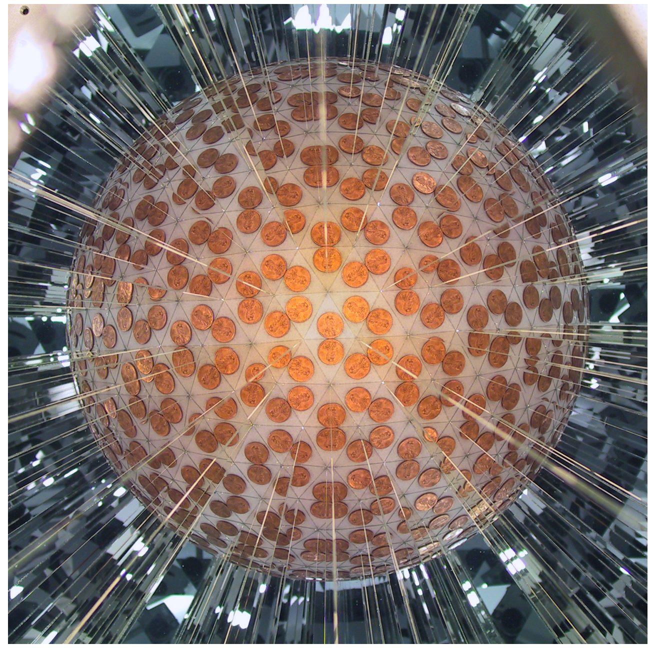

Figure 7: A multiview image of a penny, using a longer taper kaleidoscope

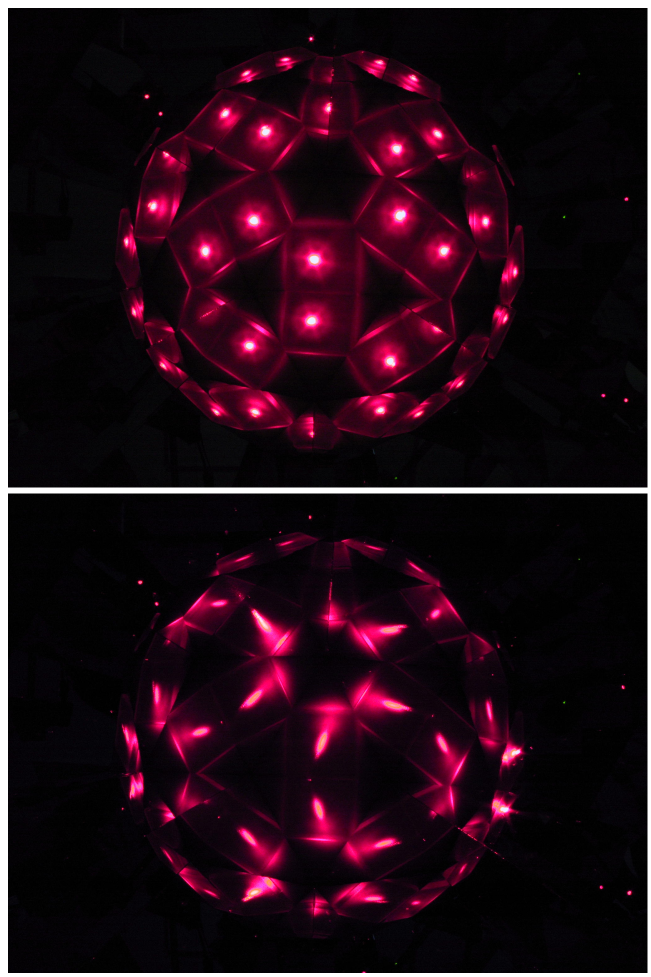

Figure 8: Two measurements of a BSSTF

Figure 9: The full 22×22 image BTF measurement of "jellybeans"

needed to build this eight dimensional function in a timely manner without any loss of precision from mechanical movement.

We have begun measuring BSSTFs through our primary kaleidoscope. Figure 8 shows two sub-steps of a measurement of a translucent block, in these early tests illuminated by a laser. Analogously, the projector would scan an image of a fine spot across the surface sample area.

8 Future Work

In near term work, we would like improve the image extraction process to utilize the data in the many visible but fragmented facets that we currently ignore. Though the resulting dataset would no longer be rectangular, it would provide an even richer BTF sampling.

We also plan to add high dynamic range (HDR) capture capability, by taking multiple image captures of varying exposure lengths, as in [Devebec and Malik 1997].

The setup we required for the current work was rather large, to accommodate the need for flexibility of research and experimentation. Also, we used an off-the-shelf camera and projector, both of which had relatively long focal length lenses. This required the kaleidoscope to be fairly tall. A commercial embodiment of the device would be able to use custom wider-angle lenses and a correspondingly smaller taper angle for the kaleidoscope walls. The device would then be far more compact, since the ratio of tube height to sample width would consequently be smaller.

We believe that the most generally useful embodiment of the technique would be a small, hand held, battery operated device, which would be used *in situ* to measure surface reflectance in somewhat the way a light meter is currently used to measure illuminance. The device would be held against any surface sample to be captured. The only essential component change will be a replacement of the projector by a set of small individually collimated white light LEDs. Because the device would lie flush against the sample, unwanted ambient light could be excluded from the measurement through the use of a light curtain. This would allow the measurement to be made under uncontrolled lighting conditions. In this embodiment, we believe the technique will have the greatest ability to positively impact the motion picture industry, by helping to reduce costs and increase flexibility for digital set construction and digital actor replacement.

At the other end of the scale, we also plan to implement a room-scale version of the technique. In this arrangement, each wall of a high-ceiling room would be a trapezoidal mirror. A two dimensional array of downward-pointing cameras and projectors would be mounted on the ceiling. This technique would provide a relatively economic way to simultaneously capture a live performance from a large number of camera angles under controllable lighting conditions.

Acknowledgements

The authors would like to thank Ned Greene, Henrik Jensen, and Ravi Ramamoorthi for their suggestions, Joel Kollin for optical engineering assistance, and everyone else on the 12th floor for their support.

References

BOIVIN, S. AND GAGALOWICZ, A. 2001. Image-Based Rendering of Diffuse, Specular and Glossy Surfaces from a Single Image. In Proceedings of ACM SIGGRAPH 2001, ACM Press / ACM SIGGRAPH, New York. E. Fiume, Ed., Computer Graphics Proceedings, Annual Conference Series, ACM, 107-116.

BREWSTER, D. 1819. A Treatise on the Kaleidoscope, A. Constable.

CARTER, R. R., AND PLESKOT, L. K. 1999. Imaging scatterometer. US Patent 5912741, June 1999.

DANA, K. J., GINNEKEN, B. VAN, NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and Texture of Real World Surfaces. ACM Transactions on Graphics, 18, 1, 1-34.

DANA, K. J. 2001. BRDF/BTF Measurement Device. In Proceedings of Eighth IEEE International Conference on Computer Vision (ICCV), IEEE Computer Society, vol. 2, pp. 460-6, Vancouver, British Columbia, July 2001.

DAVIS, K. J., AND RAWLINGS, D. C. 1997. Directional reflectometer for measuring optical bidirectional reflectance. US Patent 5637873, June 1997.

DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the Reflectance Field of a Human Face. In Proceedings of ACM SIGGRAPH 2000, ACM Press / ACM SIGGRAPH, New York. Computer Graphics Proceedings, Annual Conference Series, ACM, 145-156.

DEBEVEC, P. E., MALIK, J. 1997. Recovering High Dynamic Range Radiance Maps from Photographs. In Proceedings of ACM SIGGRAPH 1997, ACM Press / ACM SIGGRAPH, New York. Computer Graphics Proceedings, Annual Conference Series, ACM, 369-378.

DEBEVEC, P., WENGER, A., TCHOU, C., GARDNER, A., WAESE, J., AND HAWKINS, T. 2002. A Lighting Reproduction Approach to Live-Action Compositing. ACM Transactions on Graphics, 21, 3, 547-556.

FURUKAWA, R., KAWASAKI, H., IKEUCHI, K., AND SAKAUCHI, M. 2002. Appearance based object modeling using texture database: Acquisition, compression and rendering. In Proceedings of the 13th Eurographics Workshop on Rendering Techniques, pp. 257-266, 2002.

JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A Practical Model for Subsurface Light Transport. In Proceedings of ACM SIGGRAPH 2001, ACM Press / ACM SIGGRAPH, New York. E. Fiume, Ed., Computer Graphics Proceedings, Annual Conference Series, ACM, 511-518.

LIU, X., YU, Y., AND SHUM, H. Y. 2001. Synthesizing Bidirectional Texture Functions for Real-World Surfaces. In Proceedings of ACM SIGGRAPH 2001, ACM Press / ACM SIGGRAPH, New York. E. Fiume, Ed., Computer Graphics Proceedings, Annual Conference Series, ACM, 97-106.

LU, R., KOENDERINK, J. J., AND KAPPERS, A. M. L. 1998. Optical properties (bidirectional reflectance distribution functions) of velvet. Applied Optics, 37, 25, 5974-5984.

MALZBENDER, T., GELB, D., AND WOLTERS, H. 2001. Polynomial Texture Maps. In Proceedings of ACM SIGGRAPH 2001, ACM Press / ACM SIGGRAPH, New York. E. Fiume, Ed., Computer Graphics Proceedings, Annual Conference Series, ACM, 519-528.

MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., TORRANCE, K. E., AND GREENBERG, D. P. 1999. Image-based BRDF Measurement Including Human Skin. In Proceedings of the 10th Eurographics Workshop on Rendering, pp. 131-144, June 1999.

MATTISON, P. R., DOMBROWSKI, M. S., LORENZ, J., DAVIS, K., MANN, H., JOHNSON, P., AND FOOS, B. 1998. Hand-held directional reflectometer: an angular imaging device to measure BRDF and HDR in real-time. In Proceedings of SPIE, The International Society for Optical Engineering, Scattering and Surface Roughness II, 3426:240-251, July 1998.

MURRAY-COLEMAN, J. F., AND SMITH, A. M. 1990. The Automated Measurement of BRDFs and their Application to Luminaire Modeling. Journal of the Illuminating Engineering Society, pp. 87-99, Winter 1990.

NICODEMUS, F. E., RICHMOND, J. C., AND HSIA, J. J. 1977. Geometric Considerations and Nomenclature for Reflectance, U.S. Dept. of Commerce, National Bureau of Standards, October 1977.

RAMAMOORTHI, R. AND HANRAHAN, P. 2001. A Signal-Processing Framework for Inverse Rendering. In Proceedings of ACM SIGGRAPH 2001, ACM Press / ACM SIGGRAPH, New York. E. Fiume, Ed., Computer Graphics Proceedings, Annual Conference Series, ACM, 117-128.

TONG, X., ZHANG, J., LIU, L., WANG, X., GUO, B., AND SHUM, H. Y. 2002. Synthesis of Bidirectional Texture Functions on Arbitrary Surfaces. ACM Transactions on Graphics, 21, 3, 665-672.

VASILESCU, M.A.O., AND TERZOPOULOS, D. 2003. TensorTextures. ACM SIGGRAPH 2003 Conference Abstracts and Applications, July 2003.

WARD, G. J. 1992. Measuring and Modeling Anisotropic Reflection. In Computer Graphics (Proceedings of ACM SIGGRAPH 92), 26, 2, ACM, 255-263

ZHANG, Z. 1999. Flexible Camera Calibration By Viewing a Plane From Unknown Orientations. International Conference on Computer Vision (ICCV '99), Corfu, Greece, pages 666-673, September 1999.