

A Dual Light Stage

Tim Hawkins, Per Einarsson, and Paul Debevec

USC Institute for Creative Technologies

Abstract

We present a technique for capturing high-resolution 4D reflectance fields using the reciprocity property of light transport. In our technique we place the object inside a diffuse spherical shell and scan a laser across its surface. For each incident ray, the object scatters a pattern of light onto the inner surface of the sphere, and we photograph the resulting radiance from the sphere's interior using a camera with a fisheye lens. Because of reciprocity, the image of the inside of the sphere corresponds to the reflectance function of the surface point illuminated by the laser; that is, the color that point would appear to a camera along the laser ray when the object is lit from each direction on the surface of the sphere. The measured reflectance functions allow the object to be photorealistically rendered from the laser's viewpoint under arbitrary directional illumination conditions. Since each captured reflectance function is a high-resolution image, our data reproduces sharp specular reflections and self-shadowing more accurately than previous approaches. We demonstrate our technique by scanning objects with a wide range of reflectance properties and show accurate renderings of the objects under novel illumination conditions.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Capturing Real-World Data for Rendering

1. Introduction

Image-based relighting techniques simulate novel illumination on a subject based on images acquired in different basis lighting conditions. Most commonly, the basis images of the subject are taken under a discrete set of directional lighting conditions, and a linear combination of the basis images formed to produce a rendering of the subject under novel illumination. Distributing the lighting directions throughout the sphere of incident illumination allows arbitrary distant lighting environments to be simulated accurately.

In theory, this basic relighting process can reproduce the full range of reflectance phenomena an object can exhibit under distant illumination, including diffuse and specular reflection, self-shadowing, translucency, and caustics. In practice, however, the discretization of the incident lighting directions limits the technique's ability to accurately reproduce high-frequency reflectance characteristics: a shiny surface reflecting a diffuse lighting environment can appear to reflect many small light sources, and shadows cast by a moving virtual light source can appear to progress in a series of steps rather than with continuous motion.

We present a novel technique for capturing reflectance functions that exploits the reversibility of light transport, a property known as reciprocity. Our device, which we call a *dual* light stage, measures reflectance functions by reversing the traditional roles of camera and light source. Where

a camera pixel would measure radiance along an incoming ray of light R , we instead use a laser to send light out along the reversed ray $-R$. Conversely, where a light source would normally be placed to illuminate the object, we instead sense the light radiating from the object toward the same direction. We sense this reflected light by placing a diffuse spherical surface around the object, photographing the image radiated onto this sphere with a camera. While real cameras capture many pixels in parallel, and real lighting conditions must be applied one at a time, our dual light stage reverses these characteristics: for a virtual camera pixel corresponding to the current laser ray, our camera captures the response of that pixel to all illumination directions simultaneously. From the captured images, which represent reflectance functions, we can produce novel renderings of the object under arbitrary distant illumination conditions. As expected, these images appear to be acquired from the position of the laser beam, rather than from the position of the camera sensing the object's reflectance functions.

2. Background and Related Work

2.1. Image-Based Relighting

From the additive nature of light, a rendering of a scene under novel illumination can be created as a linear combination of renderings under basis lighting conditions [Hae92, NSD94]. [DHT*00] used a light stage device with




Figure 1: (a) A traditional light stage, where an object is progressively photographed as illuminated by a number of lighting directions. (b) A dual light stage, where a laser is scanned across the object and the scattered light forms images on the inside of a diffuse sphere. The reflectance function images are recorded by a camera with a fisheye lens that views the entire sphere.

a single spiraling light source to capture reflectance functions of faces and objects with either 2048 or 8192 directions on the sphere, and showed relightings of the objects using HDR image-based lighting environments acquired as in [Deb98]. [KBMK01] used a robotic arm to move a light source to different positions around an object to acquire reflectance functions of approximately 150 lighting directions. [MGW01] captured reflectance functions of the upper hemisphere of relatively diffuse objects, and created realistic real-time point-source relightings based on a parabolic fit to the reflectance function data. In our work, we capture similar datasets, but we reverse the direction of the illumination to capture reflectance functions at much higher resolution than this previous work at over 100,000 pixels, and our imaging process allows us to record the reflectance functions without aliasing as a continuous image. As a result, we are able to better reproduce high-frequency reflectance phenomena such as sharp specularities, self-shadowing, and caustics. In our work, we have used a relatively low spatial image resolution of approximately 200×200 pixels in favor of high reflectance resolution of the to keep the datasets below a reasonable size of 8GB.

[MPDW03] captures 6-dimensional reflectance fields that represent an object's response to spatially-varying illumination using a movable video projector, structured light patterns, and a fixed camera viewpoint. However, the resolution in each of the four lighting dimensions was necessarily relatively limited. In our work, we capture only 4D reflectance fields, but at significantly higher resolution than in previous work.

2.2. Environment Matting

[ZWCS99, CZH*00] has addressed the problem of capturing high-resolution reflectance function behavior for light emanating from the background behind an object. They do this by analyzing the reflectance of structured light patterns projected behind the object., and fit a parametric model of the reflectance function in this region to achieve compelling composited results for diffuse, specular, translucent, and refractive materials. Our work using reversed lighting directions in effect yields non-parametric environment mattes from the laser light scattering directly onto the area of the sphere behind the object; as a result, our renderings appear to have the background composited behind them. However, our backgrounds are effectively much lower resolution than those achieved in [ZWCS99, CZH*00] due to the limited fisheye image resolution.

2.3. Hybrid Techniques

[MPN*02] combines the image-based visual hull techniques of [MBR*00], the reflectance field acquisition technique of [DHT*00], and the environment matting technique of [ZWCS99] to capture high resolution parametric reflectance functions and 3D geometry hull using LCD monitors and a sparse set of light sources. [MPZ*02] further extended these techniques to apply to specularly reflective and refractive objects. Our work has not focussed on three-dimensional acquisition, but shows a unified approach to environment matte and reflectance function acquisition by imaging the complete sphere of scattered light from an object with relatively detailed resolution.

2.4. Spatially-Varying Reflectance Measurement

[LKG*01] uses multiple illumination conditions and viewpoints and a BRDF clustering technique to recover spatially varying diffuse and specular reflectance behavior of objects with known geometry. In [GLL*04] the subsurface scattering properties of translucent objects are captured by illuminating the object with a dense set of laser points. For each surface position, its impulse response is registered by measuring the outgoing light intensity at every other surface points. Our work resembles [GLL*04] in that we also scan a laser across the object surface and we record its light scattering characteristics. Unlike [GLL*04] we do not relate our reflectance function images to object geometry and do not render from new viewpoints or spatially-varying illumination. However, for single viewpoints and distant lighting, our process simulate a wider range of reflectance properties including diffuse and specular reflection.

2.5. Reciprocity

Reciprocity is a fundamental property of light transport that is widely used in computer graphics. It is most commonly used to refer to the invariance of the BRDF with respect to exchanging the outgoing and incoming angle: the ratio of the scattered light will be identical if the two directions are interchanged.

$$f_r(\omega_i \rightarrow \omega_o) = f_r(\omega_o \rightarrow \omega_i) \quad (1)$$

However, reciprocity can be applied more generally to the global light transport in any static scene. This is often referred to as the Helmholtz reciprocity principle [vH25], although, as observed in [Vea98] and elsewhere, he stated only a restricted form of the principle which was extended by others [Ray00]. This principle states that *any* path of a light beam is always reversible, and that the relative power loss is the same for propagation in both directions. Rendering by ray tracing implicitly assumes this more general reciprocity condition.

[Vea96, Vea98] provide analyses of the limitations and proper interpretation of reciprocity. [Pot04] also provides a thorough overview.

Helmholtz reciprocity has been exploited for imaging in a number of ways, including the flying spot scanners used in early experimental television and in telecine systems, and more recently in scanning optical microscopy, color and range scanning [BRG92], and in the design of compact scanning endoscopes [SSBR01]. Reciprocity has also been used in recent years to aid stereo correspondence for non-Lambertian surface reconstruction [MKZB01]. [ZBK02] exploits reciprocity for surface reconstruction in an approach that jointly estimates accurate surface normals and stereo correspondences.

In a similar spirit to our work, [SCG*] demonstrates the

fact that cameras and video projectors are dual through reciprocity. They show that spatially varying illumination effects can be produced from multiple lighting directions using an array of cameras to produce views from the viewpoint of a video projector. Our work differs in that we use a laser rather than a video projector allowing us to efficiently capture impulse response reflectance characteristics. While we do not capture the effects of spatially varying illumination, we are able to record directional illumination as a continuous image, which removes the possibility of aliasing in the reflectance functions.

3. Apparatus

Our scanning setup, seen in Figure 2, consists of a 140 cm diameter sphere, a high-speed color video camera with a fish-eye lens, a 3-watt white laser, and a 2-axis galvanometer. The sphere is constructed from two acrylic domes, which are painted with 33% reflective diffuse grey primer on the inside. The laser enters the sphere through a 13 cm diameter hole somewhat above the equator of the sphere. Similarly, we made a small observation hole at the top of the sphere for the camera to view into the sphere, and a small hole in the bottom of the sphere for a stand to hold the object we wish to scan.

Our laser is a mixed gas argon-krypton ion laser with strong spectral emission lines at 488, 512, and 635 nm, and with several weaker spectral lines. The laser output power is variable and at maximum power the laser emits 3 watts of visible radiation. (This Class 4 laser is not eye safe, but the fact that it immediately enters an enclosed sphere mitigates the safety issue.) The laser appears white to the human eye, and in conjunction with our color video camera allows us to capture reflectance functions in color. However, the small amount of energy in the yellow range of the spectrum yields poor color discrimination between yellow, orange, and red. White lasers with better spectral distributions are available, such as that used in [GLL*04].

Our spatial resolution is limited by the laser dot size, which is one millimeter. Higher resolution could be achieved by focusing the laser using additional optics.

Although we use a high-speed video camera and a high power laser in our setup, this is not a necessity. The high light output allows frames to be captured with short exposures ranging from 20 μ sec to 2 ms. In conjunction with the high framerate of the video camera, this speeds up the capture process significantly. For slower scanning, a more cost effective solution could use a less powerful laser and a low cost machine vision camera.

4. Data Capture

To capture a reflectance field of an object, the galvanometer scans the laser dot through about 200 horizontal scanlines




Figure 3: (a) A glass of wine illuminated by a laser spot, showing complex reflectance and scattering. (b)-(d) Three of the captured reflectance functions for the glass of wine. These are fisheye images of the inside of the diffuse sphere, taken from above. The object itself is visible in the center of each image. These central pixels are not valid reflectance function data, and we mask them for subsequent processing. The large red swath in (b) corresponds to a caustic focussed through the wine.



Figure 2: (a) A dual light stage comprises a sphere, a laser, a 2-axis galvanometer scanner, and a video camera. (b) The top half of the sphere has been removed to show the object inside. We can see the patterns of reflected light on the inside of the sphere.

over the surface of the object. Approximately 200 times during each scanline, the video camera captures an image of the reflected irradiance on the inside of the sphere, representing the reflectance function of a virtual pixel sampling the same

ray as the current laser ray. Examples of captured reflectance function images for two different objects are seen in Figure 3. The typical resolution of 200×200 is appropriate given the size of the scanned objects and the one millimeter diameter of the laser beam. The resolution of the reflectance functions themselves is determined by the resolution of the video camera, which in our case is set to 384×384 .

To cover the high dynamic range of the reflectance functions we repeat each scan three times using different exposure times. The three scans are processed into a single high dynamic range dataset. This approach requires that the galvanometers be very repeatable, which we found to be the case. Each of the three scans takes approximately three minutes, for a total scan time of less than ten minutes. However, because each scan represents 8 GB of data which must be transferred to a hard drive, the actual elapsed time for a high dynamic range scan is approximately one hour.

5. Geometric Mapping

The reflectance function images captured by the fisheye camera record the amount of light scattered by the object in all directions. Each pixel in the reflectance function corresponds to an outgoing direction from the center of the sphere. To find this mapping we first model the fisheye lens such that pixel $(u, v) \in [-1, 1] \times [-1, 1]$ maps to a vector $\omega = (\omega_x, \omega_y, \omega_z)$ as follows:

$$\begin{aligned} \theta &= \tan^{-1} \frac{-u}{v} \\ \phi &= 2 \sin^{-1} \sqrt{\frac{u^2 + v^2}{2}} \\ \omega &= (\sin\theta \cos\phi, \sin\theta \sin\phi, -\cos\theta) \end{aligned} \quad (2)$$

For each direction, we form a ray with origin at $(0, 1, 0)$ and direction $(\omega_x, \omega_y, \omega_z)$ and intersect it with a unit sphere to find the intersection (x, y, z) which corresponds to the reflected direction for pixel (u, v) relative to the center of the

sphere:

$$(x, y, z) = (-2\omega_z\omega_x, -2\omega_z\omega_y, 1 - 2\omega_z^2) \quad (3)$$

Finally, we rotate (x, y, z) to be in the laser's coordinate system. A lighting environment transformed to this coordinate system is shown in Figure 4.




Figure 4: A light probe image in angular map format (left), and our fisheye lens representation (right).

6. Inverse Rendering

Since our sphere is concave, light will reflect multiple times inside the sphere before it goes towards the camera. In order to interpret the photographs of the reflected radiance onto the sphere's interior as reflectance functions, we need to subtract the effect of these interreflections in the sphere. We remove this indirect light from our data using an inverse rendering technique, related to but simpler than those of [Mar98] and [YDMH99].

We have found that the indirect light on the interior of a sphere is actually uniform over the entire sphere, and can be quickly computed and subtracted from each image to form an accurate reflectance function. The irradiance from indirect illumination E , at a surface point \mathbf{x} , is the integral of radiance L , over the whole surface area of the sphere:

$$E(\mathbf{x}) = \int_A L \cos\theta_i \frac{\cos\theta_o}{r^2} dA \quad (4)$$

As illustrated in Figure 5, θ_i is the angle between the incident light direction and the surface normal, θ_o is the angle between the exitant light direction and its corresponding surface normal, and r is the distance between the two points. From basic spherical geometry, we have $\theta_o = \theta_i$ and $r = 2\cos\theta$. It follows that the indirect irradiance is constant for every point inside the sphere:

$$E = \int_A \frac{L}{4} dA = \pi L_{avg} \quad (5)$$

where L_{avg} is the average radiance of the sphere. To compute the corrected reflectance function R for pixel (x, y) , we

subtract L_{avg} multiplied by the Lambertian reflectance of the sphere from each pixel:

$$R(x, y) = L(x, y) - \rho_d L_{avg} \quad (6)$$

An example of a reflectance function before and after this correction can be seen in Figure 6.

We note that the principle we are using to correct for indirect light also explains the success of integrating spheres in producing extremely even illumination fields. The principle implies that when a light source illuminates a point on the inside of a sphere, the radiosity of all other points inside the sphere is constant and independent of the position of the light source. Though no surface is perfectly Lambertian, we found the Lambertian assumption was sufficient to remove most of the indirect light from our reflectance functions. We also assume that the scanned object's contribution to indirect lighting is negligible since our objects are relatively small compared to the size of the sphere.




Figure 5: The geometry for inverse rendering.




Figure 6: A reflectance function before (a) and after (b) bounce light subtraction. Since the bounce light correction is just a uniform subtraction, the features of the reflectance function are unchanged, but it appears darker. Note that these reflectance functions are shown at a bright exposure to make the bounce light apparent.

7. Rendering

To render images from our reflectance function data set we first subtract the indirect light from the reflectance functions as described in Section 6. To light our scene using an environment map, we first remap the light probe into the same mapping space as the reflectance data. For each rendered pixel, we then compute the dot product between the pixel's reflectance function and the lighting environment being used. To account for the fact that a pixel near the outside of our fisheye mapping (corresponding to upward pointing directions) represents a smaller solid angle than pixels near the center (corresponding to downward pointing directions), we weight the contribution from pixel p by the cosine of the image-space distance from the center of the map.

8. Results

In Figure 7 we see renderings of a glass of wine, a glass sculpture of a snail, and a convex metal bowl, each illuminated by three different lighting environments. Animations of these objects under varying illumination may be seen on the accompanying video.

A blurred version of the environment is visible in the background of the renderings. This corresponds to pixels where the laser hits the sphere directly. The dynamic range of our datasets was sufficient to correctly image the resulting very bright spots, and the rendering algorithm naturally results in the background appearing in the correct place (strictly speaking there will be some parallax error if the environment is meant to be at a distance from the object different from the distance from the sphere to the object). The observed blurring is a consequence of the lower angular resolution of our video camera relative to the angular resolution of the laser scan.

9. Discussion

We briefly examine the advantages and disadvantages of the dual measurement approach. It would not be particularly difficult to rearrange the basic elements of our scanner to perform a more traditional reflectance function measurement. This could be accomplished by placing wide-angle optics in front of the laser to allow it to bounce off the sphere, providing directional illumination, while placing an appropriate lens on the video camera to capture images of the object.

Often high resolution in illumination space is not needed, and high resolution in image space is desired. In this case the loss of the parallelism of the camera for image space capture seems problematic. However, this need not be the case. Many industrial and machine vision cameras support very high frame rates (10,000 fps or higher) at lower resolution through binning. Since the laser can be scanned very rapidly over the object (1000 scanlines, corresponding to 1000 x

1000 image space resolution, can be scanned using a galvanometer in about 100 milliseconds, or in less than a millisecond using a rotating polygon mirror), the bottleneck in both the traditional and dual case becomes the serial transfer of data to storage, resulting in similar overall performance.

A key distinction between traditional and dual capture is that cameras can operate in parallel without interfering, while only a single lighting condition can be applied at one time. Thus where the traditional approach supports simultaneous capture of multiple views, the dual approach supports simultaneous capture of multiple lighting environments (see also [SCG*]). Any approach that captures images under varying illumination, such as environment matting and its extensions, laser-stripe scanning, the DISCO scans of [GLL*04], and the linear light source scans of [GTHD03] can be captured in parallel with each other and with our dual light stage data by designing appropriate sensors. Some of these datasets are already implicitly present in the dual light stage data, but specialized additional sensors may allow higher resolution capture with minimal additional data. The data captured by such sensors need not be enormous; only as many pixels are needed as lighting conditions typically captured. For example, linear light source reflectometry, laser-stripe scanning, and the high-quality environment matting technique of [CZH*00] can all be accomplished making use of linear CCD's equipped with a cylindrical lens, either directed toward the object (laser-stripe scanning) or toward the diffuse sphere (environment matting and linear light source reflectometry).

10. Future Work

Acquiring a traditional matte or an environment matte could increase the quality of the rendered backgrounds. Of course, a second camera zoomed on the background could provide a high resolution background matte, but this would greatly expand the size of the captured dataset. A matte could likely be computed using no additional data by analyzing the intensity in the expected location of the directly transmitted laser for each laser position. Acquiring an environment matte could be done with additional sensors as discussed in Section 9.

Our capture process proceeds without the use of feedback. Feedback could perhaps be used to identify areas of rapid change for more detailed scanning, similar to [PD03],

11. Conclusion

We have demonstrated a photometric scanning system that captures detailed reflectance functions of objects making use of the reciprocity property of light transport. The system allows the capture of effects that vary rapidly with illumination direction, such as those associated with highly specular objects, transparent and refractive objects, and objects with complex self-shadowing.

12. Acknowledgements

We gratefully acknowledge Bill Arkin, Shane Chen, Steve Gwynne, Peter James, Tom Pereira, and Laurie Swanson for assistance with building the dual light stage, and John Lai for video editing assistance. The project or effort described here has been sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM). Statements and opinions expressed do not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred.

References

- [BRG92] BARIBEAU R., RIOUX M., GODIN G.: Color reflectance modeling using a polychromatic laser range sensor. *IEEE Trans. Pattern Anal. Mach. Intell.* 14, 2 (1992), 263–269.
- [CZH*00] CHUANG Y.-Y., ZONGKER D. E., HINDORFF J., CURLESS B., SALESIN D. H., SZELISKI R.: Environment matting extensions: Towards higher accuracy and real-time capture. In *Proceedings of SIGGRAPH 2000* (July 2000), pp. 121–130.
- [Deb98] DEBEVEC P.: Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings of SIGGRAPH 98* (July 1998), Computer Graphics Proceedings, Annual Conference Series, pp. 189–198.
- [DHT*00] DEBEVEC P., HAWKINS T., TCHOU C., DUIKER H.-P., SAROKIN W., SAGAR M.: Acquiring the reflectance field of a human face. *Proceedings of SIGGRAPH 2000* (July 2000), 145–156.
- [GLL*04] GOESELE M., LENSCH H. P. A., LANG J., FUCHS C., SEIDEL H.-P.: Disco: acquisition of translucent objects. *ACM Trans. Graph.* 23, 3 (2004), 835–844.
- [GTHD03] GARDNER A., TCHOU C., HAWKINS T., DEBEVEC P.: Linear light source reflectometry. *ACM Trans. Graph.* 22, 3 (2003), 749–758.
- [Hae92] HAEBERLI P.: Synthetic lighting for photography. Available at <http://www.sgi.com/grafica/synth/index.html>, January 1992.
- [KBMK01] KOUDELKA M. L., BELHUMEUR P. N., MAGDA S., KRIEGMAN D. J.: Image-based modeling and rendering of surfaces with arbitrary brdfs. In *2001 Conference on Computer Vision and Pattern Recognition (CVPR 2001)* (Dec. 2001), vol. 1, pp. 568–575.
- [LKG*01] LENSCH H. P. A., KAUTZ J., GOESELE M., HEIDRICH W., SEIDEL H.-P.: Image-based reconstruction of spatially varying materials. In *Rendering Techniques 2001: 12th Eurographics Workshop on Rendering* (June 2001), pp. 103–114.
- [Mar98] MARSCHNER S. R.: *Inverse Rendering for Computer Graphics*. 1998. Ph.D Thesis.
- [MBR*00] MATUSIK W., BUEHLER C., RASKAR R., GORTLER S. J., MCMILLAN L.: Image-based visual hulls. In *Proc. SIGGRAPH 2000* (July 2000), pp. 369–374.
- [MGW01] MALZBENDER T., GELB D., WOLTERS H.: Polynomial texture maps. *Proceedings of SIGGRAPH 2001* (August 2001), 519–528.
- [MKZB01] MAGDA S., KRIEGMAN D. J., ZICKLER T., BELHUMEUR P. N.: Beyond lambert: Reconstructing surfaces with arbitrary brdfs. In *ICCV* (2001), pp. 391–399.
- [MPDW03] MASSELUS V., PEERS P., DUTRÉ P., WILLEMS Y. D.: Relighting with 4d incident light fields. *ACM Transactions on Graphics* 22, 3 (July 2003), 613–620.
- [MPN*02] MATUSIK W., PFISTER H., NGAN A., BEARDSLEY P., ZIEGLER R., MCMILLAN L.: Image-based 3d photography using opacity hulls. In *Proceedings of ACM SIGGRAPH 2002* (July 2002), Computer Graphics Proceedings, Annual Conference Series, ACM Press / ACM SIGGRAPH, pp. 427–437.
- [MPZ*02] MATUSIK W., PFISTER H., ZIEGLER R., NGAN A., MCMILLAN L.: Acquisition and rendering of transparent and refractive objects. In *EGRW '02: Proceedings of the 13th Eurographics workshop on Rendering* (Aire-la-Ville, Switzerland, Switzerland, 2002), Eurographics Association, pp. 267–278.
- [NSD94] NIMEROFF J. S., SIMONCELLI E., DORSEY J.: Efficient re-rendering of naturally illuminated environments. In *Fifth Eurographics Workshop on Rendering* (June 1994), pp. 359–373.
- [PD03] PEERS P., DUTRE P.: Wavelet environment matting. In *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering* (Aire-la-Ville, Switzerland, Switzerland, 2003), Eurographics Association, pp. 157–166.
- [Pot04] POTTON R.: Reciprocity in optics. In *Reports on Progress in Physics* (2004), vol. 67, pp. 717–754.
- [Ray00] RAYLEIGH J.: On the law of reciprocity in diffuse reflection. *Philosophical Magazine*, 49 (1900), 324–325.
- [SCG*] SEN P., CHEN B., GARG G., MARSCHNER S., HOROWITZ M., LEVOY M., LENSCH H. P. A.: Dual photography. To appear in *ACM Transactions on Graphics (SIGGRAPH 2005)*.
- [SSBR01] SEIBEL E., SMITHWICK Q., BROWN C., REINHALL P.: Single fiber flexible endoscope: general design for small size, high resolution, and wide field of view. In *Proceedings of the SPIE, Biomonitoring and Endoscopy Technologies* (2001), vol. 4158, pp. 29–39.
- [Vea96] VEACH E.: Non-symmetric scattering in light transport algorithms. In *Proceedings of the eurographics workshop on Rendering techniques '96* (London, UK, 1996), Springer-Verlag, pp. 81–90.
- [Vea98] VEACH E.: *Robust monte carlo methods for light transport simulation*. PhD thesis, 1998. Adviser-Leonidas J. Guibas.
- [vH25] VON HELMHOLTZ H.: *Treatise on Physiological Optics*. Dover, New York, 1925.
- [YDMH99] YU Y., DEBEVEC P., MALIK J., HAWKINS T.: Inverse global illumination: recovering reflectance models of real scenes from photographs. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1999), ACM Press/Addison-Wesley Publishing Co., pp. 215–224.
- [ZBK02] ZICKLER T., BELHUMEUR P. N., KRIEGMAN D. J.: Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. In *7th European Conference on Computer Vision* (2002), pp. 869–894.
- [ZWCS99] ZONGKER D. E., WERNER D. M., CURLESS B., SALESIN D. H.: Environment matting and compositing. *Proceedings of SIGGRAPH 99* (August 1999), 205–214.



Figure 7: Renderings of three objects under the three different illumination environments shown at top. The environments are rotated differently for each object to best demonstrate the results. The renderings demonstrate accurate reproduction of sharp specular reflections, as well as other features of global light transport such as soft and hard shadows, caustics, and transparency and translucency. Additional lighting conditions including sharp directional lighting are included in the video.