

Estimating Surface Reflectance Properties of a Complex Scene under Captured Natural Illumination

Paul Debevec Chris Tchou Andrew Gardner Tim Hawkins
 Charis Poullis Jessi Stumpfel Andrew Jones Nathaniel Yun
 Per Einarsson Therese Lundgren Marcos Fajardo Philippe Martinez[†]

University of Southern California Institute for Creative Technologies Graphics Laboratory
 Ecole Normale Supérieure[†]



Fig. 1: (a) One of eight input photographs (b) Estimated reflectance properties (c) Synthetic rendering with novel lighting

ABSTRACT

We present a process for estimating spatially-varying surface reflectance of a complex scene observed under natural illumination conditions. The process uses a laser-scanned model of the scene's geometry, a set of digital images viewing the scene's surfaces under a variety of natural illumination conditions, and a set of corresponding measurements of the scene's incident illumination in each photograph. The process then employs an iterative inverse global illumination technique to compute surface colors for the scene which, when rendered under the recorded illumination conditions, best reproduce the scene's appearance in the photographs. In our process we measure BRDFs of representative surfaces in the scene to better model the non-Lambertian surface reflectance. Our process uses a novel lighting measurement apparatus to record the full dynamic range of both sunlit and cloudy natural illumination conditions. We employ Monte-Carlo global illumination, multiresolution geometry, and a texture atlas system to perform inverse global illumination on the scene. The result is a lighting-independent model of the scene that can be re-illuminated under any form of lighting. We demonstrate the process on a real-world archaeological site, showing that the technique can produce novel illumination renderings consistent with real photographs as well as reflectance properties that are consistent with ground-truth reflectance measurements.

1 Introduction

Digitizing objects and environments from the real world has become an important part of creating realistic computer graphics. Capturing geometric models has become a common process through the use of structured lighting, laser triangulation, and laser time-of-flight measurements. Recent projects such as [Levoy et al. 2000; Rushmeier et al. 1998; Ikeuchi 2001] have shown that accurate and detailed geometric models can be acquired of real-world objects using these techniques.

To produce renderings of an object under changing lighting as well as viewpoint, it is necessary to digitize not only the object's geometry but also its reflectance properties: how each point of the object reflects light. Digitizing reflectance properties has proven to be a complex problem, since these properties can vary across the

surface of an object, and since the reflectance properties of even a single surface point can be complicated to express and measure. Some of the best results that have been obtained [Rushmeier et al. 1998; Marschner 1998; Lensch et al. 2003] capture digital photographs of objects from a variety of viewing and illumination directions, and from these measurements estimate reflectance model parameters for each surface point.

Digitizing the reflectance properties of outdoor scenes can be more complicated than for objects since it is more difficult to control the illumination and viewpoints of the surfaces. Surfaces are most easily photographed from ground level rather than from a full range of angles. During the daytime the illumination conditions in an environment change continuously. Finally, outdoor scenes generally exhibit significant mutual illumination between their surfaces, which must be accounted for in the reflectance estimation process. Two recent pieces of work have made important inroads into this problem. [Yu and Malik 1998] estimated spatially varying reflectance properties of an outdoor building based on fitting observations of the incident illumination to a sky model, and [Yu et al. 1999] estimated reflectance properties of a room interior based on known light source positions and using a finite element radiosity technique to take surface interreflections into account.

In this paper, we describe a process that synthesizes previous results for digitizing geometry and reflectance and extends them to the context of digitizing a complex real-world scene observed under arbitrary natural illumination. The data we acquire includes a geometric model of the scene obtained through laser scanning, a set of photographs of the scene under various natural illumination conditions, a corresponding set of measurements of the incident illumination for each photograph, and finally, a small set of BRDF measurements of representative surfaces within the scene. To estimate the scene's reflectance properties, we use a global illumination algorithm to render the model from each of the photographed viewpoints as illuminated by the corresponding incident illumination measurements. We compare these renderings to the photographs, and then iteratively update the surface reflectance properties to best correspond to the scene's appearance in the photographs. Full BRDFs for the scene's surfaces are inferred from the measured BRDF samples. The result is a set of estimated reflectance prop-

erties for each point in the scene that most closely generates the scene's appearance under all input illumination conditions.

While the process we describe leverages existing techniques, our work includes several novel contributions. These include our incident illumination measurement process, which can measure the full dynamic range of both sunlit and cloudy natural illumination conditions, a hand-held BRDF measurement process suitable for use in the field, and an iterative multiresolution inverse global illumination process capable of estimating surface reflectance properties from multiple images for scenes with complex geometry seen under complex incident illumination.

The scene we digitize is the Parthenon in Athens, Greece, done in collaboration with the ongoing Acropolis Restoration project. Scaffolding and equipment around the structure prevented the application of the process to the middle section of the temple, but we were able to derive models and reflectance parameters for both the East and West facades. We validated the accuracy of our results by comparing our reflectance measurements to ground truth measurements of specific surfaces around the site, and we generate renderings of the model under novel lighting that are consistent with real photographs of the site. At the end of the paper we discuss avenues for future work to increase the generality of these techniques.

2 Background and Related Work

The process we present leverages previous results in 3D scanning, reflectance modeling, lighting recovery, and reflectometry of objects and environments. Techniques for building 3D models from multiple range scans generally involve first aligning the scans to each other [Besl and McKay 1992; Chen and Medioni 1992], and then combining the scans into a consistent geometric model by either "zippering" the overlapping meshes [Turk and Levoy 1994] or using volumetric merging [Curlless and Levoy 1996] to create a new geometric mesh that optimizes its proximity to all of the available scans.

In its simplest form, a point's reflectance properties can be expressed in terms of its Lambertian surface color - usually an RGB triplet expressing the point's red, green, and blue reflectance properties. More complex reflectance models can include parametric models of specular and retroreflective components; some commonly used models are [Larson 1992; Oren and Nayar 1994; Lafortune et al. 1997]. More generally, a point's reflectance can be characterized in terms of its Bi-directional Reflectance Distribution Function (BRDF) [Nicodemus et al. 1977], which is a 4D function that characterizes for each incident illumination direction the complete distribution of reflected illumination. [Marschner et al. 1999] proposed an efficient method for measuring a material's BRDFs if a convex homogeneous sample is available. Recent work has proposed models which also consider scattering of illumination within translucent materials [Jensen et al. 2001].

To estimate a scene's reflectance properties, we use an incident illumination measurement process. [Marschner and Greenberg 1997] recovered low-resolution incident illumination conditions by observing an object with known geometry and reflectance properties. [Sato et al. 1999] estimated incident illumination conditions by observing the shadows cast from objects with known geometry. [Debevec 1998] acquired high resolution lighting environments by taking high dynamic range images [Debevec and Malik 1997] of a mirrored sphere, but did not recover natural illumination environments where the sun was directly visible. We combine ideas from Debevec:1998:RSO, and [Marschner and Greenberg 1997] to record high-resolution incident illumination conditions in cloudy, partly cloudy, and sunlit environments.

Considerable recent work has presented techniques to measure spatially-varying reflectance properties of objects. [Marschner 1998] used photographs of a 3D scanned object taken under point-

light source illumination to estimate its spatially varying diffuse albedo. This work used a texture atlas system to store the surface colors of arbitrarily complex geometry, which we also perform in our work. The work assumed that the object was Lambertian, and only considered local reflections of the illumination. [Sato et al. 1997] used a similar sort of dataset to compute a spatially-varying diffuse component and a sparsely sampled specular component of an object. [Rushmeier et al. 1998] used a photometric stereo technique (e.g. [Ikeuchi and Horn 1979; Nayar et al. 1994]) to estimate spatially varying Lambertian color as well as improved surface normals for the geometry. [Rocchini et al. 2002] used this technique to compute diffuse texture maps for 3D scanned objects from multiple images. [Debevec et al. 2000] used a dense set of illumination directions to estimate spatially-varying diffuse and specular parameters and surface normals. [Lensch et al. 2003] presents an advanced technique for recovering spatially-varying BRDFs of real-world objects, performing principal component analysis of relatively sparse lighting and viewing directions to cluster the object's surfaces into patches of similar reflectance. In this way, many reflectance observations of the object as a whole are used to estimate spatially-varying BRDF models for surfaces seen from limited viewing and lighting directions. Our reflectance modeling technique is less general, but adapts ideas from this work to estimate spatially-varying non-Lambertian reflectance properties of outdoor scenes observed under natural illumination conditions, and we also account for mutual illumination.

Capturing the reflectance properties of surfaces in large-scale environments can be more complex, since it can be harder to control the lighting conditions on the surfaces and the viewpoints from which they are photographed. [Yu and Malik 1998] solved for the reflectance properties of a polygonal model of an outdoor scene modeled with photogrammetry. The technique used photographs taken under clear sky conditions, fitting a small number of radiance measurements to a parameterized sky model. The process estimated spatially varying diffuse and piecewise constant specular parameters, but did not consider retroreflective components. The process derived two *pseudo-BRDFs* for each surface, one according to its reflectance of light from the sun and one according to its reflectance of light from the sky and environment. This allowed more general spectral modeling but required every surface to be observed under direct sunlight in at least one photograph, which we do not require. Using room interiors, [Yu et al. 1999; Loscos et al. 1999; Boivin and Galgowicz 2002] estimate spatially varying diffuse and piecewise constant specular parameters using inverse global illumination. The techniques used knowledge of the position and intensity of the scene's light sources, using global illumination to account for the mutual illumination between the scene's surfaces. Our work combines and extends aspects of each of these techniques: we use pictures of our scene under natural illumination conditions, but we image the illumination directly in order to use photographs taken in sunny, partly sunny, or cloudy conditions. We infer non-Lambertian reflectance from sampled surface BRDFs. We do not consider full-spectral reflectance, but have found RGB imaging to be sufficiently accurate for the natural illumination and reflectance properties recorded in this work. We provide comparisons to ground truth reflectance for several surfaces within the scene. Finally, we use a more general Monte-Carlo global illumination algorithm to perform our inverse rendering, and we employ a multiresolution geometry technique to efficiently process a complex laser-scanned model.

3 Data Acquisition and Calibration

3.1 Camera Calibration

In this work we used a Canon EOS D30 and a Canon EOS 1Ds digital camera, which were calibrated geometrically and radiomet-

rically. For geometric calibration, we used the Camera Calibration Toolbox for Matlab [Bouguet 2002] which uses techniques from [Zhang 2000]. Since changing the focus of a lens usually changes its focal length, we calibrated our lenses at chosen fixed focal lengths. The main lens used for photographing the environment was a 24mm lens focussed at infinity. Since a small calibration object held near this lens would be out of focus, we built a larger calibration object 1.2m \times 2.1m from an aluminum honeycomb panel with a 5cm square checkerboard pattern applied (Fig. 2(a)). Though nearly all images were acquired at f/8 aperture, we verified that the camera intrinsic parameters varied insignificantly (less than 0.05%) with changes of f/stop from f/2.8 to f/22.

In this work we wished to obtain radiometrically linear pixel values that would be consistent for images taken with different cameras, lenses, shutter speeds, and f/stops. We verified that the "RAW" 12-bit data from the cameras was linear using three methods: we photographed a gray scale calibration chart, we used the radiometric self-calibration technique of [Debevec and Malik 1997], and we verified that pixel values were proportional to exposure times for a static scene. From this we found that the RAW pixel values exhibited linear response to within 0.1% for values up to 3000 out of 4095, after which saturation appeared to reduce pixel sensitivity. We ignored values outside of this linear range, and we used multiple exposures to increase the effective dynamic range of the camera when necessary.




Figure 2: (a) 1.2m \times 2.1m geometric calibration object (b) Lens falloff measurements for 24mm lens at f/8 (c) Lens falloff curve for (b)

Most lenses exhibit a radial intensity falloff, producing dimmer pixel values at the periphery of the image. We mounted each camera on a Kaidan nodal rotation head and photographed a diffuse disk light source at an array of positions for each lens at each f/stop used for data capture (Fig. 2(b)). From these intensities recorded at different image points, we fit a radially symmetric 6th-order even polynomial to model the falloff curve and produce a flat-field response function, normalized to unit response at the image center.

The digital cameras used each had minor variations in sensitivity and color response. We calibrated these variations by photographing a MacBeth color checker chart under natural illumination with each camera, lens, and f/stop combination, and solved for the best 3×3 color matrix to convert each image into the same color space. Finally we used a utility for converting RAW images to floating-point images using the EXIF metadata for camera model, lens, ISO, f/stop, and shutter speed to apply the appropriate radiometric scaling factors and matrices. These images were organized in a PostgreSQL database for convenient access.

3.2 BRDF Measurement and Modeling

In this work we measure BRDFs of a set of representative surface samples, which we use to form the most plausible BRDFs for the rest of the scene. Our relatively simple technique is motivated by the principal component analyses of reflectance properties used in [Lensch et al. 2003] and [Matusik et al. 2003], except that we choose our basis BRDFs manually. Choosing the principal BRDFs in this way meant that BRDF data collected under controlled illumination could be taken for a small area of the site, while the large-scale scene could be photographed under a limited set of natural

illumination conditions.

3.2.1 Data Collection and Registration

The site used in this work is composed entirely of marble, but its surfaces have been subject to different discoloration processes yielding significant reflectance variations. We located an accessible 30cm \times 30cm surface that exhibited a range of coloration properties representative of the site. Since measuring the reflectance properties of this surface required controlled illumination conditions, we performed these measurements during our limited nighttime access to the site and used a BRDF measurement technique that could be executed quickly.



Figure 3: BRDF Samples are measured from a 30cm square region exhibiting a representative set of surface reflectance properties. The technique used a hand-held light source and camera and a calibration frame to acquire the BRDF data quickly.

The BRDF measurement setup (Fig. 3), includes a hand-held light source and camera, and uses a frame placed around the sample area that allows the lighting and viewing directions to be estimated from the images taken with the camera. The frame contains fiducial markers at each corner of the frame's aperture from which the camera's position can be estimated, and two glossy black plastic spheres used to determine the 3D position of the light source. Finally, the device includes a diffuse white reflectance standard parallel to the sample area for determining the intensity of the light source.

The light source chosen was a 1000W halogen source mounted in a small diffuser box, held approximately 3m from the surface. Our capture assumed that the surfaces exhibited isotropic reflection, requiring the light source to be moved only within a single plane of incidence. We placed the light source in four consecutive positions of 0° , 30° , 50° , 75° , and for each took hand-held photographs at a distance of approximately 2m from twenty directions distributed on the incident hemisphere, taking care to sample the specular and retroreflective directions with a greater number of observations. Dark clothing was worn to reduce stray illumination on the sample. The full capture process involving 83 photographs required forty minutes.

3.2.2 Data Analysis and Reflectance Model Fitting

To calculate the viewing and lighting directions, we first determined the position of the camera from the known 3D positions of the four fiducial markers using photogrammetry. With the camera positions known, we computed the positions of the two spheres by tracing rays from the camera centers through the sphere centers for several photographs, and calculated the intersection points of these rays. With the sphere positions known, we determined each light position by shooting rays toward the center of the light's reflection in the

spheres. Reflecting the rays off the spheres, we find the center of the light source position where the two rays most nearly intersect. Similar techniques to derive light source positions have been used in [Masselus et al. 2002; Lensch et al. 2003].

From the diffuse white reflectance standard, the incoming light source intensity for each image could be determined. By dividing the overall image intensity by the color of the reflectance standard, all images were normalized by the incoming light source intensity. We then chose three different areas within the sampling region best corresponding to the different reflectance properties of the large-scale scene. These properties included a light tan area that is the dominant color of the site’s surfaces, a brown color corresponding to encrusted biological material, and a black color representative of soot deposits. To track each of these sampling areas across the dataset, we applied a homography to each image to map them to a consistent orthographic viewpoint. For each sampling area, we then obtained a BRDF sample by selecting a 30x30 pixel region and computing the average RGB value. Had there been a greater variety of reflectance properties in the sample, a PCA analysis of the entire sample area as in [Lensch et al. 2003] could have been used.

To extrapolate the BRDF samples to a complete BRDF, we fit the BRDF to the Lafortune cosine lobe model (Eq. 1) in its isotropic form with three lobes for the Lambertian, specular, and retroreflective components. As suggested in [Lafortune et al. 1997], we then use a non-linear Levenberg-Marquardt optimization algorithm to determine the parameters of the model from our measured data. We first estimate the Lambertian component ρ_d , and then fit a retroreflective and a specular lobe separately before optimizing all the parameters in a single system. The resulting BRDFs (Fig. 5(b), back row) show mostly Lambertian reflectance with noticeable retroreflection and rough specular components at glancing angles. The brown area exhibited the greatest specular reflection, while the black area was the most retroreflective.

$$f(\vec{u}, \vec{v}) = \rho_d + \sum_i [C_{xy,i}(u_x v_x + u_y v_y) + C_{z,i} u_z v_z]^N \quad (1)$$

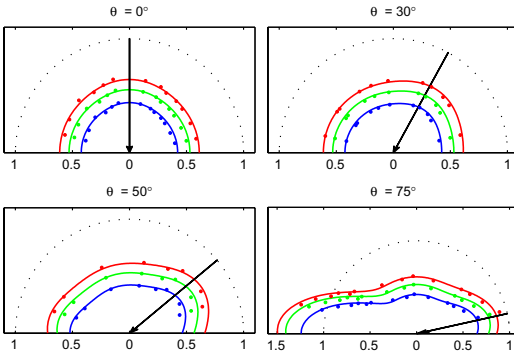


Figure 4: **BRDF Data and Fitted Reflectance Lobes** are shown for the RGB colors of the tan material sample for the four incident illumination directions. Only measurements within 15° of in-plane are plotted.

3.2.3 BRDF Inference

We wish to be able to make maximal use of the BRDF information obtained from our material samples in estimating the reflectance properties of the rest of the scene. The approach we take is informed by the BRDF basis construction technique from [Lensch et al. 2003], the data-driven reflectance model presented in [Matusik et al. 2003], and spatially-varying BRDF construction tech-

nique used in [Marschner et al. 2000]. Because the surfaces of the rest of the scene will often be seen in relatively few photographs under relatively diffuse illumination, the most reliable observation of a surface’s reflectance is its Lambertian color. Thus, we form our problem as one of inferring the most plausible BRDF for a surface point given its Lambertian color and the BRDF samples available.

We first perform a principal component analysis of the Lambertian colors of the BRDF samples available. For RGB images, the number of significant eigenvalues will be at most three, and for our samples the first eigenvalue dominates, corresponding to a color vector of (0.688, 0.573, 0.445). We project the Lambertian color of each of our sample BRDFs onto the 1D subspace S (Fig. 5(a) formed by this eigenvector. To construct a plausible BRDF f for a surface having a Lambertian color ρ_d , we project ρ_d onto S to obtain the projected color ρ'_d . We then determine the two BRDF samples whose Lambertian components project most closely to ρ'_d . We form a new BRDF f' by linearly interpolating the Lafortune parameters (C_{xy}, C_z, N) of the specular and retroreflective lobes of these two nearest BRDFs f_0 and f_1 based on distance. Finally, since the retroreflective color of a surface usually corresponds closely to its Lambertian color, we adjust the color of the retroreflective lobe to correspond to the actual Lambertian color ρ_d rather than the projected color ρ'_d . We do this by dividing the retroreflective parameters C_{xy} and C_z by $(\rho'_d)^{1/N}$ and then multiplying by $(\rho_d)^{1/N}$ for each color channel, which effectively scales the retroreflective lobe to best correspond to the Lambertian color ρ_d . Fig. 5(b) shows a rendering with several BRDFs inferred from new Lambertian colors with this process.




Figure 5: (a) Inferring a BRDF based on its Lambertian component ρ_d (b) Rendered spheres with measured and inferred BRDFs. Back row: the measured black, brown, and tan surfaces. Middle row: intermediate BRDFs along the subspace S . Front row: inferred BRDFs for materials with Lambertian colors not on S .

3.3 Natural Illumination Capture

Each time a photograph of the site was taken, we used a device to record the corresponding incident illumination within the environment. The lighting capture device was a digital camera aimed at three spheres: one mirrored, one shiny black, and one diffuse gray. We placed the device in a nearby accessible location far enough from the principal structure to obtain an unshadowed view of the sky, and close enough to ensure that the captured lighting would be sufficiently similar to that incident upon the structure. Measuring the incident illumination directly and quickly enabled us to make use of photographs taken under a wide range of weather including sunny, cloudy, and partially cloudy conditions, and also in changing conditions.

3.3.1 Apparatus Design

The lighting capture device is designed to measure the color and intensity of each direction in the upper hemisphere. A challenge in capturing such data for a natural illumination environment is that the sun’s intensity can exceed that of the sky by over five orders of magnitude, which is significantly beyond the range of most digital




Figure 6: (a) The incident illumination measurement device at its chosen location on the site (b) An incident illumination dataset.

image sensors. This dynamic range surpassing 17 stops also exceeds that which can conveniently be captured using high dynamic range capture techniques. Our solution was to take a limited dynamic range photograph and use the mirrored sphere to image the sky and clouds, the shiny black sphere to indicate the position of the sun (if visible), and the diffuse grey sphere to indirectly measure the intensity of the sun. We placed all three spheres on a board so that they could be photographed simultaneously (Fig. 6). We painted the majority of the board gray to allow a correct exposure of the device to be derived from the camera’s auto-exposure function, but surrounded the diffuse sphere by black paint to minimize the indirect light it received. We also included a sundial near the top of the board to validate the lighting directions estimated from the black sphere. Finally, we placed four fiducial markers on the board to estimate the camera’s relative position to the device.

We used a Canon D30 camera with a resolution of 2174×1446 pixels to capture images of the device. Since the site photography took place up to 300m from the incident illumination measurement station, we used a radio transmitter to trigger the device at the appropriate times. Though the technique we describe can work with a single image of the device, we set the camera’s internal auto-exposure bracketing function to take three exposures for each shutter release at -2 , $+0$, and $+2$ stops. This allowed somewhat higher dynamic range to better image brighter clouds near the sun, and to guard against any problems with the camera’s automatic light metering.

3.3.2 Sphere Reflectance Calibration

To achieve accurate results, we calibrated the reflectance properties of the spheres. The diffuse sphere was painted with flat gray primer paint, which we measured as having a reflectivity of (0.30, 0.31, 0.32) in the red, green, and blue color channels. We further verified it to be nearly spectrally flat using a spectroradiometer. We also exposed the paint to several days of sunlight to verify its color stability. In the above calculations, we divide all pixel values by the sphere’s reflectance, producing values that would result from a perfectly reflective white sphere.

We also measured the reflectivity of the mirrored sphere, which was made of polished steel. We measured this reflectance by using a robotic arm to rotate a rectangular light source in a circle around the sphere and taking a long-exposure photograph of the resulting reflection (Fig. 7(a)). We found that the sphere was 52% reflective at normal incidence, becoming more reflective toward grazing angles due to Fresnel reflection (Fig. 7(b)). From the measured reflectance data we used a nonlinear optimization to fit a Fresnel curve to the data, arriving at a complex index of refraction of $(2.40 + 2.98i, 2.40 + 3.02i, 2.40 + 3.02i)$ for the red, green, and blue channels of the sphere.

Light from a clear sky can be significantly polarized, particularly in directions perpendicular to the direction of the sun. In our work we assume that the surfaces in our scene are not highly specular, which makes it reasonable for us to disregard the polarization of the incident illumination in our reflectometry process. However, since




Figure 7: (a) Mirrored sphere photographed under an even ring of light, showing an increase in brightness at extreme grazing angles (the dark gap in the center is due to light source occluding the camera). (b) Fitted Fresnel reflectance curves.

Fresnel reflection is affected by the polarization of the incoming light, the clear sky may reflect either more or less brightly toward the grazing angles of the mirrored sphere than it should if it were photographed directly. To quantify this potential error, we photographed several clear skies reflected in the mirrored sphere and at the same time took hemispherical panoramas with a 24mm lens. Comparing the two, we found an RMS error of 5% in sky intensity between the sky photographed directly and the sky photographed as reflected in the mirrored sphere (Fig. 8). In most situations, however, unpolarized light from the sun, clouds, and neighboring surfaces dominates the incident illumination on surfaces, which minimizes the effect of this error. In Section 6, we suggest techniques for eliminating this error through improved optics.




Figure 8: (a) Sky photographed as reflected in a mirrored sphere (b) Stitched sky panorama from 16 24mm photographs, showing slightly different reflected illumination due to sky polarization.

3.3.3 Image Processing and Deriving Sun Intensity

To process these images, we assemble each set of three bracketed images into a single higher dynamic range image, and derive the relative camera position from the fiducial markers. The fiducial markers are indicated manually in the first image of each day and then tracked automatically through the rest of day, compensating for small motions due to wind. Then, the reflections in both the mirrored and shiny black spheres are transformed to 512×512 images of the upper hemisphere. This is done by forward-tracing rays from the camera to the spheres (whose positions are known) and reflecting the rays into the sky, noting for each sky point the corresponding location on the sphere. The image of the diffuse sphere is also mapped to the sky’s upper hemisphere, but based on the sphere’s normals rather the reflection vectors. In the process, we also adjust for the reflectance properties of the spheres as described in Section 3.3.2, creating the images that would have been produced by spheres with unit albedo. Examples of these unwarped images are shown in Fig. 9.

If the sun is below the horizon or occluded by clouds, no pixels in the mirrored sphere image will be saturated and it can be used directly as the image of the incident illumination. We can validate the accuracy of this incident illumination map by rendering a synthetic diffuse image D' with this lighting and checking that it is consistent




Figure 9: Sphere images unwarped to the upper hemisphere for the (a) Mirrored sphere (b) Shiny black sphere (c) Diffuse sphere D . Saturated pixels are shown in black.

with the appearance of the actual diffuse sphere image D . As described in [Miller and Hoffman 1984], this lighting operation can be performed using a diffuse convolution filter on the incident lighting environment. For our data, the root mean square illumination error for our diffuse sphere images agreed to within 2% percent for a variety of environments.

When the sun is visible, it usually saturates a small region of pixels in the mirrored sphere image. Since the sun’s bright intensity is not properly recorded in this region, performing a diffuse convolution of the mirrored sphere image will produce a darker image than actual appearance of the diffuse sphere (Compare D' to D in Fig. 10). In this case, we reconstruct the illumination from the sun as follows. We first measure the direction of the sun as the center of the brightest spot reflected in the shiny black sphere (with its darker reflection, the black sphere exhibits the most sharply defined image of the sun.) We then render an image of a diffuse sphere D^* lit from this direction of illumination, using a unit-radiance infinite light source 0.53 degrees in diameter to match the subtended angle of the real sun. Such a rendering can be seen in the center of Figure 10.




Figure 10: Solving for sun intensity α based on the appearance of the diffuse sphere D and the convolved mirrored sphere D' .

We can then write that the appearance of the real diffuse sphere D should equal the sphere lit by the light captured in the mirrored sphere D' plus an unknown factor α times the sphere illuminated by the unit sun D^* , i.e.

$$D' + \alpha D^* = D$$

Since there are many pixels in the sphere images, this system is overdetermined, and we compute the red, green, and blue components of α using least squares as $\alpha D^* \approx D - D'$. Since D^* was rendered using a unit radiance sun, α indicates the radiance of the sun disk for each channel. For efficiency, we keep the solar illumination modeled as the directional disk light source, rather than updating the mirrored sphere image M to include this illumination. As a result, when we create renderings with the measured illumination, the solar component is more efficiently simulated as a direct light source.

We note that this process does not reconstruct correct values for the remaining saturated pixels near the sun; the missing illumination from these regions is effectively added to the sun’s intensity. Also, if the sun is partially obscured by a cloud, the center of the saturated region might not correspond precisely to the center of the

sun. However, for our data the saturated region has been sufficiently small that this error has not been significant. Fig. 11 shows a lighting capture dataset and a comparison rendering of a model of the capture apparatus, showing consistent captured illumination.




Figure 11: (a) Real photograph of the lighting capture device (b) Synthetic rendering of a 3D model of the lighting capture device to validate the lighting measurements.

3.4 3D Scanning

To obtain 3D geometry for the scene, we used a time-of-flight panoramic range scanner manufactured by Quantapoint, Inc, which uses a 950nm infrared laser measurement component [Hancock et al. 1998]. In high resolution mode, the scanner acquires scans of 18,000 by 3,000 3D points in 8 minutes, with a maximum scanning range of 40m and a field of view of 360 degrees horizontal by 74.5 degrees vertical. Some scans from within the structure were scanned in low-resolution, acquiring one-quarter the number of points. The data returned is an array of (x,y,z) points as well as a 16-bit monochrome image of the infrared intensity returned to the sensor for each measurement. Depending on the strength of the return, the depth accuracy varied between 0.5cm and 3cm. Over five days, 120 scans were acquired in around the site, of which 53 were used to produce the model in this paper.




Figure 12: Range measurements, shaded according to depth (top), and infrared intensity return (bottom) for one of 53 panoramic laser scans used to create the model. A fiducial marker appears at right.

3.4.1 Scan Processing

Our scan processing followed the traditional process of alignment, merging, and decimation. Scans from outside the structure were initially aligned during the scanning process through the use of checkerboard fiducial markers placed within the scene. After the site survey, the scans were further aligned using an iterative closest point (ICP) algorithm [Besl and McKay 1992; Chen and Medioni 1992] implemented in the CNR-Pisa 3D scanning toolkit [Callieri et al. 2003]. To speed the alignment process, three or more subsections of each scan corresponding to particular scene areas were cropped out and used to determine the alignment for the entire scan.

For merging, the principal structure of the site was partitioned into an $8 \times 17 \times 5$ lattice of voxels 4.3 meters on a side. For convenience, the grid was chosen to align with the principal architectural features of the site. The scan data within each voxel was merged by a volumetric merging algorithm [Curless and Levoy 1996] also from the CNR-Pisa toolkit using a volumetric resolution of 1.2cm.

Finally, the geometry of a $200m \times 200m$ area of surrounding terrain was merged as a single mesh with a resolution of 40cm.

Several of the merged voxels contained holes due to occlusions or poor laser return from dark surfaces. Since such geometric inconsistencies would affect the reflectometry process, they were filled using semi-automatic tools with Geometry Systems, Inc. GSI Studio software (Fig. 13).




Figure 13: (a) Geometry for a voxel colored according to texture atlas regions (b) The corresponding texture atlas.

Our reflectometry technique determines surface reflectance properties which are stored in texture maps. We used a texture atlas generator [Graphite 2003] based on techniques in [Lévy et al. 2002] to generate a 512×512 texture map for each voxel. Then, a low-resolution version of each voxel was created using the Qslim software [Qslim 1999] based on techniques in [Garland and Heckbert 1998]. This algorithm was chosen since it preserves edge polygons, allowing low resolution and high resolution voxels to connect without seams, and since it preserves the texture mapping space, allowing the same texture map to be used for either the high or low resolution geometry.




Figure 14: Complete model assembled from the 3D scanning data, including low-resolution geometry for the surrounding terrain. High and medium resolution voxels used for the multiresolution reflectance recovery are indicated in white and blue.

The complete high-resolution model of the main structure used 89 million polygons in 442 non-empty voxels (Figure 14). The lowest-resolution model contained 1.8 million polygons, and the surrounding environment used 366K polygons.

3.5 Photograph Acquisition and Alignment

Images were taken of the scene from a variety of viewpoints and lighting conditions using the Canon 1Ds camera. We used a semi-automatic process to align the photographs to the 3D scan data. We began by marking approximately 15 point correspondences between each photo and the infrared intensity return image of one or more 3D scans, forming a set of 2D to 3D correspondences. From this we estimated the camera pose using Intel’s OpenCV library, achieving a mean alignment error of between 1 and 3 pixels at 4080×2716 pixel resolution. For photographs with higher

alignment error, we use an automatic technique to refine the alignment based on comparing the structure’s silhouette in the photograph to the model’s silhouette seen through the recovered camera as in [Lensch et al. 2001], using a combination of gradient-descent and simulated annealing.

4 Reflectometry

In this section we describe the central reflectometry algorithm used in this work. The basic goal is to determine surface reflectance properties for the scene such that renderings of the scene under captured illumination match photographs of the scene taken under that illumination. We adopt an inverse rendering framework as in [Boivin and Gagalowicz 2002; Debevec 1998] in which we iteratively update our reflectance parameters until our renderings best match the appearance of the photographs. We begin by describing the basic algorithm and continue by describing how we have adapted it for use with a large dataset.

4.1 General Algorithm

The basic algorithm we use proceeds as follows:

1. Assume initial reflectance properties for all surfaces
2. For each photograph:
 - (a) Render the surfaces of the scene using the photograph’s viewpoint and lighting
 - (b) Determine a reflectance update map by comparing radiance values in the photograph to radiance values in the rendering
 - (c) Compute weights for the reflectance update map
3. Update the reflectance estimates using the weightings from all photographs
4. Return to step 2 until convergence

For a pixel’s Lambertian component, the most natural update for a pixel’s Lambertian color is to multiply it by the ratio of its color in the photograph to its color in the corresponding rendering. This way, the surface will be adjusted to reflect the correct proportion of the light. However, the indirect illumination on the surface may change in the next iteration since other surfaces in the scene may also have new reflectance properties, requiring further iterations.

Since each photograph will suggest somewhat different reflectance updates, we weight the influence a photograph has on a surface’s reflectance by a confidence measure. For one weight, we use the cosine of the angle at which the photograph views the surface. Thus, photographs which view surfaces more directly will have a greater influence on the estimated reflectance properties. As in traditional image-based rendering (e.g. [Buehler et al. 2001]), we also downweight a photograph’s influence near occlusion boundaries. Finally, we also downweight an image’s influence near large irradiance gradients in the photographs since these typically indicate shadow boundaries, where small misalignments in lighting could significantly affect the reflectometry.

In this work, we use the inferred Lafortune BRDF models described in Sec. 3.2.3 to create the renderings, which we have found to also converge accurately using updates computed in this manner. This convergence occurs for our data since the BRDF colors of the Lambertian and retroreflective lobes both follow the Lambertian color, and since for all surfaces most of the photographs do not observe a specular reflection. If the surfaces were significantly more specular, performing the updates according to the Lambertian component alone would not necessarily converge to accurate reflectance estimates. We discuss potential techniques to address this problem in the future work section.




Figure 15: **Computing reflectance properties for a voxel** (a) Iteration 0: 3D model illuminated by captured illumination, with assumed reflectance properties (b) Photograph taken under the captured illumination projected onto the geometry (c) Iteration 1: New reflectance properties computed by comparing (a) to (b). (d) Iteration 2: New reflectance properties computed by comparing a rendering of (c) to (b).

4.2 Multiresolution Reflectance Solving

The high-resolution model for our scene is too large to fit in memory, so we use a multiresolution approach to computing the reflectance properties. Since our scene is partitioned into voxels, we can compute reflectance property updates one voxel at a time. However, we must still model the effect of shadowing and indirect illumination for the rest of the scene. Fortunately, lower-resolution geometry can work well for this purpose. In our work, we use full-resolution geometry (approx. 800K triangles) for the voxel being computed, medium-resolution geometry (approx. 160K triangles) for the immediately neighboring voxels, and low-resolution geometry (approx. 40K triangles) for the remaining voxels in the scene. (The surrounding terrain is kept at a low resolution of 370K triangles.) The multiresolution approach results in over a 90% data reduction in scene complexity during the reflectometry of any given voxel.

Our global illumination rendering system was originally designed to produce 2D images of a scene for a given camera viewpoint using path tracing [Kajiya 1986]. We modified the system to include a new function for computing surface radiance for any point in the scene radiating toward any viewing position. This allows the process of computing reflectance properties for a voxel to be done by iterating over the texture map space for that voxel. For efficiency, for each pixel in the voxel’s texture space, we cache the position and surface normal of the model corresponding to that texture coordinate, storing these results in two additional floating-point texture maps.

1. Assume initial reflectance properties for all surfaces
2. For each voxel V :
 - (a) Load V at high resolution, V ’s neighbors at medium resolution, and the rest of the model at low resolution
 - (b) For each pixel p in V ’s texture space
 - i. For each photograph I :
 - A. Determine if p ’s surface is visible to I ’s camera. If not, break. If so, determine the weight for this image based on the visibility angle, and note pixel q in I corresponding to p ’s projection into I .
 - B. Compute the radiance l of p ’s surface in the direction of I ’s camera under I ’s illumination
 - C. Determine an updated surface reflectance by comparing the radiance in the image at q to the rendered radiance l .
 - ii. Assign the new surface reflectance for p as the weighted average of the updated reflectances from each I

3. Return to step 2 until convergence



Figure 16: Estimated surface reflectance properties for an East facade column in texture atlas form.

Figure 15 shows this process of computing reflectance properties for a voxel. Fig. 15(a) shows the 3D model with the assumed initial reflectance properties illuminated by a captured illumination environment. Fig. 15(b) shows the voxel texture-mapped with radiance values from a photograph taken under the captured illumination in (a). Comparing the two, the algorithm determines updated surface reflectance estimates for the voxel, shown in Fig. 15(c). The second iteration compares an illuminated rendering of the model with the first iteration’s inferred BRDF properties to the photograph, producing new updated reflectance properties shown in 15(d). For this voxel, the second iteration produces a darker Lambertian color for the underside of the ledge, which results from the fact that the *black* BRDF sample measured in Section 3.2 has a higher proportion of retroreflection than the average reflectance. The second iteration is computed with a greater number of samples per ray, producing images with fewer noise artifacts. Reflectance estimates for three voxels of a column on the East Facade are shown in texture atlas form in Fig. 16. Reflectance properties for all voxels of the two Facades are shown in Figs. 1(b) and 19(d). For our model, the third iteration produces negligible change from the second, indicating convergence.

5 Results

We ran our reflectometry algorithm on the 3D scan dataset, computing high-resolution reflectance properties for the two westmost and eastmost rows of voxels. As input to the algorithm, we used eight photographs of the East Facade (e.g. Fig. 1 (a)) and three of the West Facade, in an assortment of sunny, partly cloudy, and cloudy lighting conditions. Poorly scanned scaffolding which had been removed from the geometry was replaced with approximate polygonal models in order to better simulate the illumination transport within the structure. The reflectance properties of the ground were assigned based on a sparse sampling of ground truth measurements made with a MacBeth chart. We recovered the reflectance properties in two iterations of the reflectometry algorithm. For each iteration of the reflectometry, the illumination was simulated with two indirect bounces using the inferred Lafortune BRDFs. Computing the reflectance for each voxel required an average of ten minutes.

Figs. 1(b) and 19(c) show the computed Lambertian reflectance colors for the East and West Facades, respectively. Recovered texture atlas images for three voxels of the East column second from left are shown in Fig. 16. The images show few shading effects, suggesting that the maps have removed the effect of the illumination in the photographs. The subtle shading observable toward the back sides of the columns is likely the result of incorrectly computed indirect illumination due to the remaining discrepancies in the scaffolding.

Figs. 19(a) and (b) show a comparison between a real photograph and a synthetic global illumination rendering of the East Facade under the lighting captured for the photograph, indicating a consistent appearance. The photograph represents a significant variation in the lighting from all images used in the reflectometry

dataset. Fig. 19(d) shows a rendering of the West Facade model under novel illumination and viewpoint. Fig. 19(e) shows the East Facade rendered under novel artificial illumination. Fig. 19(f) shows the East Facade rendered under sunset illumination captured from a different location than the original site.




Figure 17: Left: Acquiring a ground truth reflectance measurement. Right: Reflectance comparisons for four locations on the East Facade

To provide a quantitative validation of the reflectance measurements, we directly measured the reflectance properties of several surfaces around the site using a MacBeth color checker chart. Since the measurements were made at normal incidence and in diffuse illumination, we compared the results to the Lambertian lobe directly, as the specular and retroreflective lobes are not pronounced under these conditions. The results tabulated in Fig. 17 show that the computed reflectance largely agreed with the measured reflectance samples, with a mean error of (2.0%, 3.2%, 4.2%) for the red, green, and blue channels.

6 Discussion and Future Work

Our experiences with the process suggest several avenues for future work. Most importantly, it would be of interest to increase the generality of the reflectance properties which can be estimated using the technique. Our scene did not feature surfaces with sharp specularities, but most scenes featuring contemporary architecture do. To handle this larger gamut of reflectance properties, one could imagine adapting the BRDF clustering and basis formation techniques in [Lensch et al. 2003] to photographs taken under natural illumination conditions. Our technique for interpolating and extrapolating our BRDF samples is relatively simplistic; using more samples and a more sophisticated analysis and interpolation as in [Matusik et al. 2003] would be desirable. A challenge in adapting these techniques to natural illumination is that observations of specular behavior are less reliable in natural illumination conditions. Estimating reflectance properties with increased spectral resolution would also be desirable.

In our process the photographs of the site are used only for estimating reflectance, and are not used to help determine the geometry of the scene. Since high-speed laser scan measurements can be noisy, it would be of interest to see if photometric stereo techniques as in [Rushmeier et al. 1998] could be used in conjunction with natural illumination to refine the surface normals of the geometry. [Yu and Malik 1998] for example used photometric stereo from different solar positions to estimate surface normals for a building's environment; it seems possible that such estimates could also be made given three images of general incident illumination with or without the sun.

Our experience calibrating the illumination measurement device showed that its images could be affected by sky polarization. We tested the alternative of using an upward-pointing fisheye lens to image the sky, but found significant polarization sensitivity toward

the horizon as well as undesirable lens flare from the sun. More successfully, we used a 91% reflective aluminum-coated hemispherical lens and found it to have less than 5% polarization sensitivity, making it suitable for lighting capture. For future work, it might be of interest to investigate whether sky polarization, explicitly captured, could be leveraged in determining a scene's specular parameters [Nayar et al. 1997].

Finally, it could be of interest to use this framework to investigate the more difficult problem of estimating a scene's reflectance properties under unknown natural illumination conditions. In this case, estimation of the illumination could become part of the optimization process, possibly by fitting to a principal component model of measured incident illumination conditions.

7 Conclusion

We have presented a process for estimating spatially-varying surface reflectance properties of an outdoor scene based on scanned 3D geometry, BRDF measurements of representative surface samples, and a set of photographs of the scene under measured natural illumination conditions. Applying the process to a real-world archaeological site, we found it able to recover reflectance properties close to ground truth measurements, and able to produce renderings of the scene under novel illumination consistent with real photographs. The encouraging results suggest further work be carried out to capture more general reflectance properties of real-world scenes using natural illumination.

Acknowledgements

We acknowledge the efforts of Brian Emerson, Marc Brownlow throughout this project, Andreas Wenger for helpful assistance, and John Shipley for his 3D scanning work. This project was produced by Diane Piepol, Lora Chen, and Maya Martinez. We wish to thank Tomas Lochman, Nikos Toganidis, Katerina Paraschis, Manolis Korres, Richard Lindheim, David Wertheimer, Neil Sullivan, and Angeliki Arvanitis, Cheryl Birch, James Blake, Bri Brownlow, Chris Butler, Elizabeth Cardman, Alan Chalmers, Yikuong Chen, Paolo Cignoni, Jon Cohen, Costis Dallas, Christa Deacy-Quinn, Paul T. Debevec, Naomi Dennis, Apostolos Dimopoulos, George Drettakis, Paul Egri, Costa-Gavras, Rob Groome, Christian Guillon, Youda He, Eric Hoffman, Leslie Ikemoto, Genichi Kawada, Cathy Kominos, Marc Levoy, Dell Lunceford, Mike Macedonia, Chrysostomos Nikias, Mark Ollila, Yannis Papoutsakis, John Parmentola, Fred Persi, Dimitrios Raptis, Simon Ratcliffe, Roberto Scopigno, Alexander Singer, Judy Singer, Laurie Swanson, Bill Swartout, Despoina Theodorou, Mark Timpson, Greg Ward, Karen Williams, and Min Yu for their help making this project possible. We also offer our great thanks to the Hellenic Ministry of Culture, the Work Site of the Acropolis, the Basel Skulpturhalle, the Musee du Louvre, the Herodion Hotel, Quantapoint, Alias, CNR-Pisa, and Geometry Systems, Inc. for their support of this project. This work was sponsored by TOPPAN Printing Co., Ltd., the University of Southern California office of the Provost, and Department of the Army contract DAAD 19-99-D-0046. Any opinions, findings and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the sponsors.

References

- BESL, P., AND MCKAY, N. 1992. A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Machine Intell.* 14, 239–256.
- BOIVIN, S., AND GAGALOWICZ, A. 2002. Inverse rendering from a single image. In *Proceedings of IS&T CGIV*.
- BOUGUET, J.-Y., 2002. Camera calibration toolbox for matlab. <http://www.vision.caltech.edu/bouguetj/calib.doc/>.




Figure 18: The East Facade model rendered under novel natural illumination conditions at several times of day.

- BUEHLER, C., BOSSE, M., McMILLAN, L., GORTLER, S. J., AND COHEN, M. F. 2001. Unstructured lumigraph rendering. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, 425–432.
- CALLIERI, M., CIGNONI, P., GANOVELLI, F., MONTANI, C., PINGI, P., AND SCOPIGNO, R. 2003. Vclab's tools for 3d range data processing. In *VAST 2003 and EG Symposium on Graphics and Cultural Heritage*.
- CHEN, Y., AND MEDIONI, G. 1992. Object modeling from multiple range images. *Image and Vision Computing* 10, 3 (April), 145–155.
- CURLESS, B., AND LEVOY, M. 1996. A volumetric method for building complex models from range images. In *Proceedings of SIGGRAPH 96*, ACM SIGGRAPH / Addison Wesley, New Orleans, Louisiana, Computer Graphics Proceedings, Annual Conference Series, 303–312.
- DEBEVEC, P. E., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *Proceedings of SIGGRAPH 97*, Computer Graphics Proceedings, Annual Conference Series, 369–378.
- DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. *Proceedings of SIGGRAPH 2000* (July), 145–156.
- DEBEVEC, P. 1998. 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*, Computer Graphics Proceedings, Annual Conference Series, 189–198.
- GARLAND, M., AND HECKBERT, P. S. 1998. Simplifying surfaces with color and texture using quadric error metrics. In *IEEE Visualization '98*, 263–270.
- GRAPHITE, 2003. <http://www.loria.fr/levy/Graphite/index.html>.
- HANCOCK, J., LANGER, D., HEBERT, M., SULLIVAN, R., INGIMARSON, D., HOFFMAN, E., METTENLEITER, M., AND FROELICH, C. 1998. Active laser radar for high-performance measurements. In *Proc. IEEE International Conference on Robotics and Automation*.
- IKEUCHI, K., AND HORN, B. 1979. An application of the photometric stereo method. In *6th International Joint Conference on Artificial Intelligence*, 413–415.
- IKEUCHI, K. 2001. Modeling from reality. In *Proc. Third Intern. Conf on 3-D Digital Imaging and Modeling*, 117–124.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proceedings of SIGGRAPH 2001*, ACM Press / ACM SIGGRAPH, Computer Graphics Proceedings, Annual Conference Series, 511–518. ISBN 1-58113-292-1.
- KAJIYA, J. 1986. The rendering equation. In *SIGGRAPH 86*, 143–150.
- LAFORTUNE, E. P. F., FOO, S.-C., TORRANCE, K. E., AND GREENBERG, D. P. 1997. Non-linear approximation of reflectance functions. *Proceedings of SIGGRAPH 97*, 117–126.
- LARSON, G. J. W. 1992. Measuring and modeling anisotropic reflection. In *Computer Graphics (Proceedings of SIGGRAPH 92)*, vol. 26, 265–272.
- LENSCH, H. P. A., HEIDRICH, W., AND SEIDEL, H.-P. 2001. A silhouette-based algorithm for texture registration and stitching. *Graphical Models* 63, 4 (Apr.), 245–262.
- LENSCH, H. P. A., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2003. Image-based reconstruction of spatial appearance and geometric detail. *ACM Transactions on Graphics* 22, 2 (Apr.), 234–257.
- LEVOY, M., PULLI, K., CURLESS, B., RUSINKIEWICZ, S., KOLLER, D., PEREIRA, L., GINTON, M., ANDERSON, S., DAVIS, J., GINSBERG, J., SHADE, J., AND FULK, D. 2000. The digital michelangelo project: 3d scanning of large statues. *Proceedings of SIGGRAPH 2000* (July), 131–144.
- LÉVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. 2002. Least squares conformal maps for automatic texture atlas generation. *ACM Transactions on Graphics* 21, 3 (July), 362–371.
- LOSCOS, C., FRASSON, M.-C., DRETTAKIS, G., WALTER, B., GRANIER, X., AND POULIN, P. 1999. Interactive virtual lighting and remodeling of real scenes. In *Eurographics Rendering Workshop 1999*.
- MARSCHNER, S. R., AND GREENBERG, D. P. 1997. Inverse lighting for photography. In *Proceedings of the IS&T/SID Fifth Color Imaging Conference*.
- 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. *Eurographics Rendering Workshop 1999* (June).
- MARSCHNER, S., GUENTER, B., AND RAGHUPATHY, S. 2000. Modeling and rendering for realistic facial animation. In *Rendering Techniques 2000: 11th Eurographics Workshop on Rendering*, 231–242.
- MARSCHNER, S. 1998. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University.
- MASSELUS, V., DUTRÉ, P., AND ANRYS, F. 2002. The free-form light stage. In *Rendering Techniques 2002: 13th Eurographics Workshop on Rendering*, 247–256.
- MATUSIK, W., PFISTER, H., BRAND, M., AND McMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics* 22, 3 (July), 759–769.
- MILLER, G. S., AND HOFFMAN, C. R. 1984. Illumination and reflection maps: Simulated objects in simulated and real environments. In *SIGGRAPH 84 Course Notes for Advanced Computer Graphics Animation*.
- NAYAR, S. K., IKEUCHI, K., AND KANADE, T. 1994. Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Transactions on Robotics and Automation* 6, 4 (August), 418–431.
- NAYAR, S., FANG, X., AND BOULT, T. 1997. Separation of reflection components using color and polarization. *IJCV* 21, 3 (February), 163–186.
- NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. 1977. Geometric considerations and nomenclature for reflectance. *National Bureau of Standards Monograph* 160 (October).
- OREN, M., AND NAYAR, S. K. 1994. Generalization of Lambert's reflectance model. *Proceedings of SIGGRAPH 94* (July), 239–246.
- QSLIM, 1999. <http://graphics.cs.uiuc.edu/garland/software/qslim.html>.
- ROCCHINI, C., CIGNONI, P., MONTANI, C., AND SCOPIGNO, R. 2002. Acquiring, stitching and blending diffuse appearance attributes on 3d models. *The Visual Computer* 18, 3, 186–204.
- RUSHMEIER, H., BERNARDINI, F., MITTLEMAN, J., AND TAUBIN, G. 1998. Acquiring input for rendering at appropriate levels of detail: Digitizing a pietà. *Eurographics Rendering Workshop 1998* (June), 81–92.
- SATO, Y., WHEELER, M. D., AND IKEUCHI, K. 1997. Object shape and reflectance modeling from observation. In *SIGGRAPH 97*, 379–387.
- SATO, I., SATO, Y., AND IKEUCHI, K. 1999. Illumination distribution from shadows. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR'99)*, 306–312.
- TURK, G., AND LEVOY, M. 1994. Zippered polygon meshes from range images. In *Proceedings of SIGGRAPH 94*, ACM SIGGRAPH / ACM Press, Orlando, Florida, Computer Graphics Proceedings, Annual Conference Series, 311–318. ISBN 0-89791-667-0.
- YU, Y., AND MALIK, J. 1998. Recovering photometric properties of architectural scenes from photographs. In *Proceedings of SIGGRAPH 98*, Computer Graphics Proceedings, Annual Conference Series, 207–218.
- YU, Y., DEBEVEC, P., MALIK, J., AND HAWKINS, T. 1999. Inverse global illumination: Recovering reflectance models of real scenes from photographs. *Proceedings of SIGGRAPH 99* (August), 215–224.
- ZHANG, Z. 2000. A flexible new technique for camera calibration. *PAMI* 22, 11, 1330–1334.



Figure 19: **(a)** A real photograph of the East Facade, with recorded illumination **(b)** Rendering of the model under the illumination recorded for (a) using inferred Lafortune reflectance properties **(c)** A rendering of the West Facade from a novel viewpoint under novel illumination. **(d)** Front view of computed surface reflectance for the West Facade (the East is shown in Fig. 1(b)). A strip of unscanned geometry above the pediment ledge has been filled in and set to the average surface reflectance. **(e)** Synthetic rendering of the West Facade under a novel artificial lighting design. **(f)** Synthetic rendering of the East Facade under natural illumination recorded for another location. In these images, only the front two rows of outer columns are rendered using the recovered reflectance properties; all other surfaces are rendered using the average surface reflectance.