Linear Light Source Reflectometry

Andrew Gardner   Chris Tchou   Tim Hawkins   Paul Debevec

University of Southern California Institute for Creative Technologies Graphics Laboratory[1]

ABSTRACT

This paper presents a technique for estimating the spatially-varying reflectance properties of a surface based on its appearance during a single pass of a linear light source. By using a linear light rather than a point light source as the illuminant, we are able to reliably observe and estimate the diffuse color, specular color, and specular roughness of each point of the surface. The reflectometry apparatus we use is simple and inexpensive to build, requiring a single direction of motion for the light source and a fixed camera viewpoint. Our model fitting technique first renders a reflectance table of how diffuse and specular reflectance lobes would appear under moving linear light source illumination. Then, for each pixel we compare its series of intensity values to the tabulated reflectance lobes to determine which reflectance model parameters most closely produce the observed reflectance values. Using two passes of the linear light source at different angles, we can also estimate per-pixel surface normals as well as the reflectance parameters. Additionally our system records a per-pixel height map for the object and estimates its per-pixel translucency. We produce real-time renderings of the captured objects using a custom hardware shading algorithm. We apply the technique to a test object exhibiting a variety of materials as well as to an illuminated manuscript with gold lettering. To demonstrate the technique's accuracy, we compare renderings of the captured models to real photographs of the original objects.

1   Introduction

Visual richness in computer-generated images can come from a wide variety of sources: complex and interesting geometry; detailed and realistic texture maps, aesthetically designed and realistically simulated lighting, and expressive animation and dynamics. In recent years it has become possible to obtain many of these rich qualities by sampling the properties of the real world: scanning 3D objects to obtain geometry, using digital images for texture maps, capturing real-world illumination, and using motion capture techniques to produce animation. What remains difficult is to acquire realistic reflectance properties of real-world objects: their spatially-varying colors, specularities, roughnesses, and translucencies. As a result, these properties are usually crafted laboriously by hand using image-editing software.

The problem of measuring object reflectance is well-understood for the case of a homogeneous material sample: a device called a
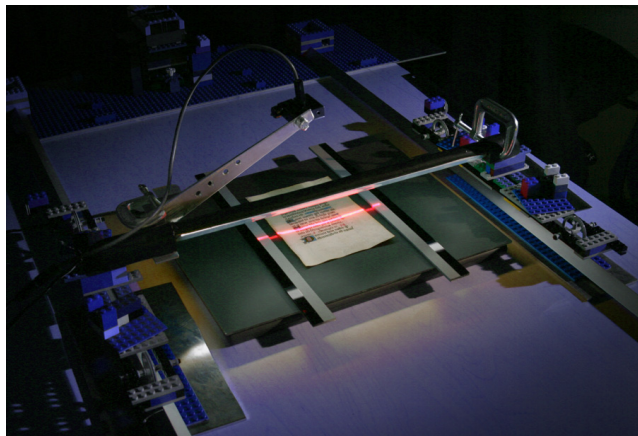
Figure 1: **The Linear Light Source Apparatus** uses a translation gantry to move a white neon tube across the captured object at a height of approximately 6cm. As the light moves across the object, a digital camera acquires a series of images which provide the necessary measurements to estimate per-pixel diffuse and specular reflectance parameters across the object. In this image the device is scanning a page from a medieval illuminated manuscript.

gonioreflectometer (e.g. [Ward 1992]) moves a light source across the full two-dimensional hemisphere of incident illumination directions (θ_i, ϕ_i) and for each incident direction measures the sample's radiance toward the full two-dimensional array of radiant directions (θ_o, ϕ_o). In this way the device captures the sample's Bidirectional Reflectance Distribution Function (BRDF), which can be used to render the material seen from any angle under any form of incident illumination. However, the capture process can be lengthy and requires specialized equipment.

To address these issues, significant work has developed mathematical *reflectance models* which can be fit to sparser samples of the BRDF. These mathematical models typically parameterize the BRDF in terms of its diffuse color for Lambertian reflectance, its specular color for glossy reflectance, and parameters relating to the relative sharpness or roughness of the specularity. However, for shiny materials the BRDF contains significant high-frequency components [Ramamoorthi and Hanrahan 2001], and parameters obtained by fitting to limited number of samples can grossly mischaracterize the specular behavior of the material.

In this paper, we attempt to reliably estimate both diffuse and specular properties by making use of a linear light source (LLS) as the illuminant in a novel reflectometry apparatus. By translating the LLS over an object, the LLS will provide illumination on the material from every possible direction of incident illumination. We show that although the different illumination directions are not observed independently, the resulting dataset provides ample information to reliably estimate the material's diffuse and specular reflectance properties. In this work, we have found that using a linear light source for reflectometry offers the following advantages:

- The light source needs to make just one pass over the object for a single viewing direction in order to have it exhibit both its diffuse and specular reflectance characteristics

- The contrast ratio between the specular and diffuse peaks is reduced, allowing for standard digital imaging systems to acquire the reflectance data.

- The mechanics for moving a light source in just one dimension are simple and inexpensive to build.

In the rest of this paper we describe the construction of our LLS apparatus, and how we process the image data to derive spatially varying reflectance parameters for an object. Our device handles objects that are mostly planar, and we demonstrate the capture technique on a flat test object as well as a medieval illuminated manuscript. We finally present a hardware-accelerated real-time rendering algorithm for the captured objects under novel illumination, and show that the technique realistically predicts the reflectance behavior of the objects under arbitrary point source illumination.

2 Background and Related Work

The work in this paper relates to the general problem of creating digitized representations of real-world objects, and draws from a wide variety of previous work in this area. The techniques can be roughly categorized as belonging to techniques for 3D scanning, reflectance modeling, reflectance measurement, and rendering.

Laser stripe systems such as those manufactured by Cyberware, Inc. have been used to acquire geometric models of mechanical parts, sculptures, faces, and bodies. Techniques using structured light patterns and/or stereo correspondence techniques have been used for geometry acquisition as well. A laser stripe system was recently used to scan Michaelangelo's David [Levoy et al. 2000] and a structured light technique was used to scan his Florentine Pieta [Rushmeier et al. 1998]. [Halstead et al. 1996] used a different approach to geometry measurement by estimating the shape of the eye's cornea from specular reflectance analysis. In our work, we use a laser stripe technique for recording the non-planar geometry of our scanned objects and a specular reflectance analysis to recover smaller-scale surface normal variations.

In reflectance modeling, Nicodemus et al [Nicodemus et al. 1977] described how a surface transforms incident illumination into radiant illumination as its four-dimensional Bidirectional Reflectance Distribution Function. Examples of BRDF models designed to be consistent with real-world reflectance properties are Torrance-Sparrow [Torrance and Sparrow 1967], Cook-Torrance [Cook and Torrance 1981], Poulin et al, [Poulin and Fournier 1990], Ward [Larson 1992], and Oren-Nayar [Oren and Nayar 1994]. More recent work has proposed models which also consider the subsurface scattering component of light reflection [Hanrahan and Krueger 1993; Jensen et al. 2001] for translucent surfaces. In our work, we make use of the Ward reflectance model [Larson 1992] with an additional simple translucency component.

Reflectance measurement, or *reflectometry*, has produced techniques for estimating the parameters of a reflectance model of real-world surfaces based on photographs of the surface under various lighting conditions and viewing directions. For one example, the Image-Based BRDF measurement technique in [Marschner et al. 1999] fitted cosine lobe [Lafortune et al. 1997] reflectance parameters based on a small set of images of a convex material sample. [Nayar et al. 1994] used pictures of an object illuminated by different extended light sources to recover object shape, lambertian reflectance, and mirror-like specular reflectance. [Dana et al. 1997] recorded spatially-varying BRDFs for various texture samples using images of the sample rather than single radiance measurements. Another technique from [Marschner 1998] recorded a spatially-varying diffuse albedo map for a laser-scanned object by considering the visibility of each object surface to a moving light

source, and [Sato et al. 1997] in addition estimated uniform specular reflectance parameters for an object by observing the intensity of a specular lobe from various lighting directions. [Yu et al. 1999] estimated spatially varying diffuse parameters and piecewise constant specular parameters for images of a scene in a mutual illumination context. [Debevec et al. 2000] used light sources positioned both horizontally and vertically around a person's face to obtain spatially varying parameters for both diffuse and specular reflection as well as surface normals; [Malzbender et al. 2001] used an array of lights on the upper hemisphere to capture diffuse properties and surface normals of artifacts for real-time rendering. [Lensch et al. 2001] estimated spatially varying diffuse and specular properties making use of a principal component analysis of the possible reflectance properties of the object. [Zongker et al. 1999; Chuang et al. 2000] used patterns of light from behind an object on a video monitor to characterize the object's reflectance functions for reflected and refracted light.

[McAllister 2002] addressed the problem of capturing spatially-varying reflectance information for flat surfaces by directly recording and fitting the four-dimensional BRDF for each point on the object's surface. Using a larger captured data set, pan and tilt robotics, and a least-squares BRDF-fitting method, the technique produced accurate BRDF measurements including anisotropy, but at the cost of relatively slow capture. In our work we use simpler equipment and less data, additionally acquire surface normals, displacements, and translucency, but do not at this stage acquire anisotropy.

These previous techniques address important aspects of the reflectometry problem for spatially-varying surfaces, but none is at once general, efficient, and simple. In our work, we have found that using a linear light source as the illuminant allows us to efficiently capture both diffuse and specular properties using a simple translational gantry and fixed camera pose relative to the subject.

3 The Linear Light Source Apparatus

In this section we describe the design and construction of our Linear Light Source reflectometry apparatus, including the linear light source, the translation gantry, the laser stripe, and the camera system.

The Linear Light Source (LLS) chosen for our device is a white neon tube purchased from the automobile decoration company Street Glow, Inc. for $20. The light runs on 12V DC and was chosen over a fluorescent tube due to its low flicker and whiter spectrum, and over an incandescent light for its low heat output. The emissive light tube source was suitably small at 1cm in diameter and 50 cm long, and sits within a protective 2cm diameter clear plastic tube.

The spectrum of the white neon light tube is shown in Fig. 2. The white "neon" light produces light by exciting an argon-mercury gas mixture, which produces spectral emission lines in the ultraviolet, blue, green, and orange parts of the spectrum. The inside of the light tube is coated with white phosphors which are excited by the higher-frequency emission lines to produce energy across the visible spectrum. The final emitted light also includes contributions from the original Ar-Hg gas; for a future version we would like to filter out these spectral peaks since they bias the recorded colors toward the object's reflectance at these wavelengths. To block any potential harm to a material sample from ultraviolet light, we envelop the entire tube with a thin sheet of ultraviolet-blocking filter from Edmund Optics, producing the final spectrum seen in red in Fig. 2.

The Translation Gantry (Fig.1) moves the LLS horizontally across over the the surface of the subject at a fixed distance h. We built our system using LEGO MindStorms Robotics Invention System 2.0. This package provided a programmable robotics system with sufficient accuracy for less than $200. The light source
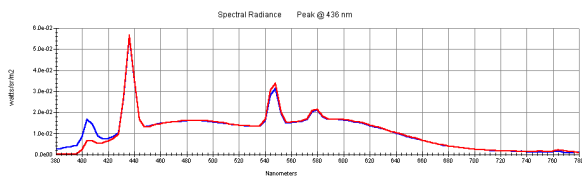
Figure 2: **Linear Light Source Spectrum** The spectral emission of the linear light source is shown before (blue) and after (red) adding the ultraviolet filter.

is mounted at the ends to the middle of two Lego trolleys whose grooved wheels ride along two $\frac{3}{4}''$ aluminum L-beams. The trolleys are attached at one end to the Lego motor system using fishing line running through pulleys. Attached to the other end of the trolleys are lines which attach to two weights which pull the trolleys forward. The LLS is attached to the trolleys via a single flat LEGO piece glued to each of its ends. To capture varying surface normals, the LLS can be mounted to the trolleys in three configurations; perpendicular to the trolley direction, and at ± 15 degrees.

The Laser Scanning System is used to record surface geometry variations, and consists of a miniature structured light diode laser with a 60 degree line lens mounted to project a laser stripe parallel to and 38mm behind the LLS. Surface variations will cause deformations in the stripe as seen from the camera, which we use to recover variations in the geometric structure of the subject. The laser and lens were purchased from Edmund Optics for $400.

The Translucency Measurement System records how a subject transmits light as well as reflects it, and was made by placing a $11\frac{3}{4} \times 9$ inch Cabin light box from Mamiya America Corporation beneath the LLS to act as the resting surface for the subject. The light box uses two diffused cathode tubes to provide an even diffuse white light across its surface. To reduce the amount of translucent light observed during the capture of the reflective properties, we place a diffuse dark gel over the surface of the light box so that its surface has an albedo of less than 0.07. This has the effect of absorbing most of the light the LLS would transmit through the object while still allowing sufficient transmitted light to illuminate the object from below for capturing its translucent properties.

The Camera System we use is a Canon EOS-D60 digital camera placed at approximately a 55 degree angle of incidence from perpendicular to the subject. The camera captures high resolution JPEG photographs of 3072×2048 pixels every five seconds. A program was created using the Canon camera control API and LEGO Mindstorms API to photograph and translate the gantry in a synchronized manner. In a typical session we capture up to three passes of the LLS with one picture taken every millimeter over a distance of 400mm. We used a long focal length Canon 200mm EF lens at f/16 aperture to record the subject with sufficient depth of field and at a nearly orthographic view. We calibrated the response curve of the D60 camera using a radiometric self-calibration similar to that in [Debevec and Malik 1997].

4 Capturing and Registering the Data

4.1 Data Capture

To capture a dataset using the LLS device, we begin by placing the object to be captured on the light box between the two rails as in Fig. 1. On either side of the object we place two calibration strips which consist of diffuse and sharp specular materials.

We originate our coordinate system at the center of the top surface of the light box. We place the camera on a tripod looking down at the object at approximately a 55 degree angle of incidence from perpendicular, an angle that adequately separates the observed
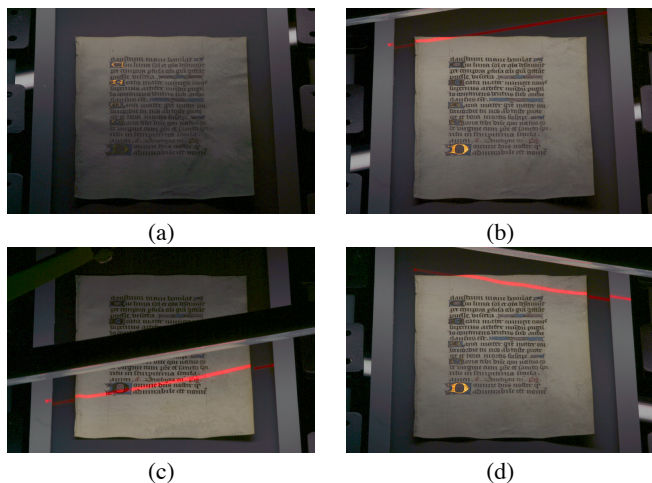

(a)

(b)

(c)

(d)

Figure 3: **A Linear Light Source Dataset** for an illuminated manuscript. In **(a)**, the LLS is above the image and forms a mirror angle with the top two gold initials; its reflection can also be seen in the shiny part of the calibration strips at left and right. In **(b)**, the LLS has moved so that the mirror angle is on the final gold "D" and the laser stripe has begun to scan across the manuscript. In **(c)**, the LLS itself is visible, and the diffuse peak it generates is at the bottom of the page. **(d)** shows an image from a second pass of the LLS at a different angle. The data from this second pass allows us to estimate the left-right component of the surface normal.

position of the diffuse and specular peaks without unduly foreshortening the object's appearance in the camera.

We attach the LLS to the gantry in one of several positions according to whether we wish to capture surface normal variations for the object. If the object is very close to being flat, we perform a single pass with the LLS mounted perpendicularly to the direction of the gantry motion. If the object has varying surface normals and displacements, we perform two passes of the LLS, one with the LLS mounted at $\beta = +15°$ (rotated about the y axis) and one with the LLS mounted at $\beta = -15°$.

We begin with the gantry high enough so that the mirror-angle reflection of the LLS can no longer be seen in the shiny parts of the calibration strips. The control software then alternates between translating the gantry and taking photographs of the subject. We stop the recording when the laser stripe of the LLS has moved past the final edge of the object, which takes approximately 30 minutes. Some individual images from an LLS dataset can be seen in Fig. 3.

If we are capturing surface normals, we perform a second pass with the light source at the opposite LLS angle. Once all passes are completed, we rewind the gantry once more and turn off the LLS and laser stripe. We turn on the light box to light the object from behind and take an additional photograph to record its translucent qualities. We then remove the object and take a photograph of the light box. Finally, we photograph a checkerboard calibration pattern centered at $(0,0,0)$ aligned with the x and z axes on the light box to complete the capture process.

4.2 Data Registration

With the data collected, we first use the image of the checkerboard to derive the camera's position v, orientation, and focal length using photogrammetric modeling techniques as found in [Debevec et al. 1996]. This allows us to construct the function V that projects image coordinates (u,v) to their corresponding 3D location P on the xz plane.

In our experiments each pass of the LLS produces approximately 400 images of the object. We refer to this set of images as $I_{u,v}(t)$

(a) Diffuse vellum

(b) Shiny black paper
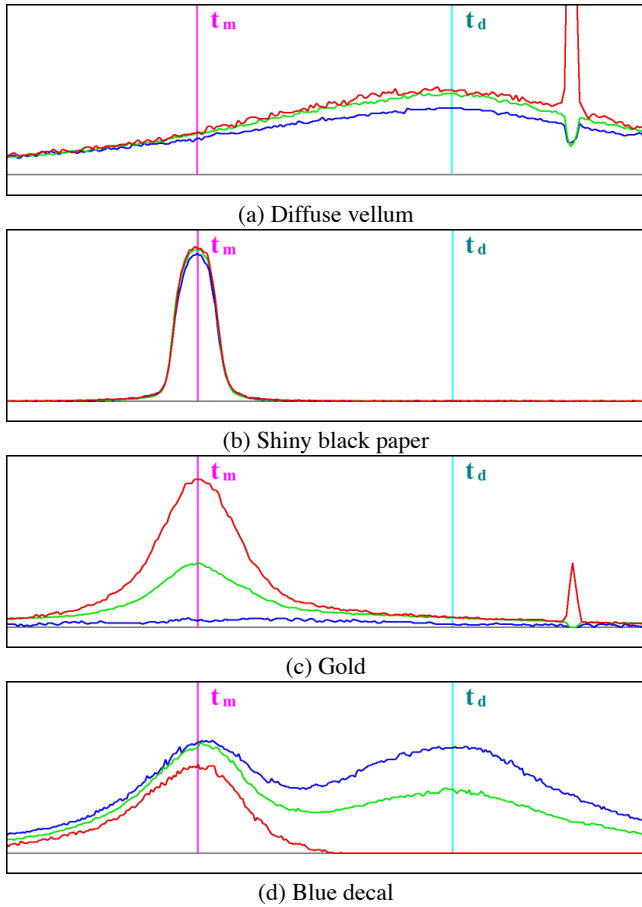
(c) Gold

(d) Blue decal

Figure 4: **Registered Reflectance Traces** $I_{u,v}(t)$ are shown for a variety of materials. As the light moves across the object, the specular lobe reflection appears first near the mirror angle t_m, followed by the diffuse reflection at t_d when the light passes directly over the surface point. In our fitting process we determine ρ_d based on the height of the diffuse lobe and ρ_s and α based on the height and breadth of the specular lobe. Traces (b) and (c) are sampled from the test object in Fig. 11; (a) and (c) are sampled from the illuminated manuscript in Fig. 13. In (a) and (c) the laser stripe used to determine the surface point's height is visible at right.

where t ranges from 1 to the number of images in the pass. For any pixel (u,v), $I_{u,v}(t)$ is a graph of the pixel's intensity as the LLS moves over the corresponding point on the object's surface; we call such a graph the pixel's *reflectance trace*. In a reflectance trace, we typically first see a peak corresponding to the specular reflection of the pixel, followed by a broader peak corresponding to the diffuse peak of the material, followed by a red spike corresponding to the time at which the laser stripe passes over the material. Since our camera is radiometrically calibrated, we assume the $I_{u,v}(t)$ to be proportional to scene radiance. Several reflectance traces for different types of object surfaces are shown in Fig. 4.

Our reflectance fitting procedure requires that we know the line equation of the linear light source $l(t)$ for each image $I(t)$. We determine this by first determining the time-varying position of the laser stripe $\hat{l}(t)$ which we do by marking the line of the laser stripe's intersection with the xz plane in each of two images in the image sequence, one at t_a when the laser is near the top of the frame and one at t_b where the laser is near the bottom of the frame. We then project each image-space line onto the xz plane in world space to compute $\hat{l}(t_a)$ and $\hat{l}(t_b)$. All other values of $\hat{l}(t)$ are based on a linear
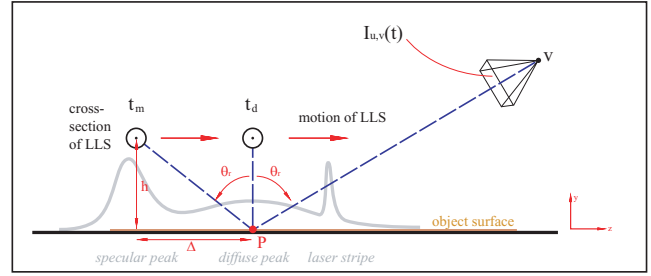

Figure 5: **Registering the Data** For each pixel (u,v) corresponding to surface point P, we use the 3D pose of the camera and the model of the light source's movement to determine the time t_d when the LLS is directly above the corresponding surface, and t_m, when the LLS is at the mirror angle to the surface with respect to the camera at v. The gray reflectance trace in the background shows the amount of light reflected by P toward the camera as the LLS moves across the object; its shape traces out the specular peak, the diffuse peak, and the laser stripe.

interpolation of its values at t_a and t_b. Since the light source is at height h and a fixed distance d in front of the laser, we compute $l(t) = \hat{l}(t) + (d \sin\beta, h, d\cos\beta)$.

From the LLS position $l(t)$, we can compute for any point P on the object the time t_d at which the light source is directly above the point, and t_m, when the LLS is at the mirror angle to the surface with respect to the camera at v as in Fig. 5. Generally, t_m will lie near the specular peak and t_d will correspond to the diffuse peak. We finally let $\Delta = t_d - t_m$, which we use to delimit the region in which a reflectance trace will be analyzed to the area $t_m \pm \Delta$.

5 The Reflectance Model

The reflectance model which we use to fit the data observed with the LLS apparatus is the isotropic Gaussian lobe model by Ward [Larson 1992]. This model was chosen for its combination of simplicity and physical accuracy, though our technique can be used with other models as well. The Ward model specifies that a surface point will reflect light from the direction (θ_i, ϕ_i) to the direction (θ_r, ϕ_r) with the following distribution f_r:

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

In this equation, δ is the angle between the surface normal $\hat{n}$ and the half vector between the incident and reflected directions $\hat{h}$. The reflectance parameters of the surface are ρ_d, the diffuse component, ρ_s, the specular intensity component, and α, the specular roughness. These components are given with respect to each color channel, so the diffuse color of a surface point is given by the RGB triple $(\rho_{dr}, \rho_{dg}, \rho_{db})$. A point's specular color $(\rho_{dr}, \rho_{dg}, \rho_{db})$ is usually spectrally even for dielectric materials such as plastic and has a particular hue for the case of metals such as gold and bronze. Usually, the same α parameter is used for all color channels, so we determine our α parameter based on a weighted average of the red, green, and blue channels. In Section 6 we describe how we analyze the LLS data to estimate (ρ_d, ρ_s, α) for every point on the object surface. In Section 7 we describe how we can additionally recover a surface normal $\hat{n}$ and displacement d for each surface point, and in Section 8 we describe how to estimate each surface point's translucency ρ_{trans}. Using all three techniques we can recover the parameter vector $(\rho_d, \rho_s, \alpha, \rho_{trans}, \hat{n}, d)$ for every point on the object surface.

6 Model Fitting for a Flat Surface

Our model fitting process must address the discrepancy that while the reflectance model f_r specifies reflected light for each specific incident angle of illumination, our reflectance observations $I_{u,v}(t)$ are produced by light reflected from an entire area of incident illumination, namely the cylindrical linear light source. The approach we take is to create a virtual version of the linear light source and calculate the predicted appearance of its reflected light for a variety of diffuse and specular parameters. We then compare the statistical characteristics of the rendered reflectance traces to the statistical characteristics of a real pixel's reflectance trace to determine which reflectance model parameters best correspond to the observed data.

In this section we describe our model fitting process for the case of flat surfaces, such as the test object in Fig. 11. We first register the observed reflectance traces into a common coordinate space, and then compute the appropriate table of pre-rendered reflectance traces to be fit to the data. Finally, we compute the diffuse and specular parameters of each surface point by comparing statistics of the real data to the the rendered basis data. Specifically, we can describe our fitting process as one of finding values for ρ_d, ρ_s, and α such that for each pixel:

$$I(t) = \rho_d D(t) + \rho_s S_\alpha(t)$$

6.1 Creating the Reflectance Table

Reflectance model fitting for a linear light source illuminant is complicated by the fact that the reflectance function f_r is specified in terms of light arriving from a point light direction at (θ_i, ϕ_i). While Lambertian reflection from a linear light source has been solved in closed form [Nishita et al. 1985], reflection from generalized specular components is generally performed with integral approximations [Poulin and Amanatides 1991]. Because of this, and to keep our method generalizable to other reflectance models, we numerically integrate f_r across the light source surface for a range of surface roughness parameters α and linear light source positions relative to the pixel. We perform these integrations assuming that the viewing direction (θ_r, ϕ_r) is perpendicular to the linear light source (i.e. $\phi_r = 0$) and that θ_r is constant. Since our camera is nearly orthographic, we have found these assumptions to be sufficiently precise.

The results of our tabulations are stored in a series of functions $S_\alpha(t)$ which correspond to the height of a specular lobe with $\rho_s = 1$ for a range of light source positions. This table is visualized in Fig. 6. We choose our parameterization of t such that at $t = 0$ the light source is directly above the surface point and that at $t = 1$ the light source is at the mirror angle with respect to the surface patch and the viewing direction. In addition, we compute the corresponding function $D(t)$ which corresponds to the height of a diffuse lobe with $\rho_d = 1$ for the same range of light positions.

In practice we render the linear light as a cylindrical light equal to the 1cm diameter of our real linear light source so that sharp specular peaks will match accordingly. To render the cylindrical light, we uniformly sample the light transversely and radially. To speed up the specular fitting process, we also compute the mean μ and the standard deviation σ of the shape of each specular lobe in the table. These numbers characterize the position and breadth of the specular lobe as produced by the linear light source, and can be see in Fig. 6.

6.2 Estimating Diffuse Color ρ_s

In our approach we first fit the diffuse lobe of the LLS reflectance trace, which allows us to determine the ρ_d parameter for each color channel. Since for this version of the fitting process we assume that the material has a vertical surface normal, we know that the diffuse peak will happen at time t_d when the light is directly above
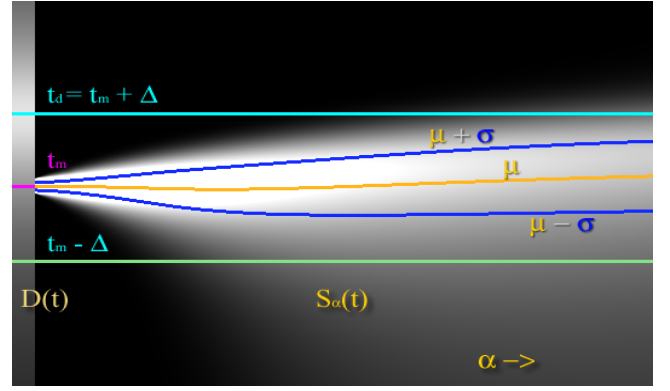

Figure 6: **The LLS Reflectance Table** $S_\alpha(t)$ stores the shape of the specular lobe for various roughness values α for a particular reflected angle $(\theta_r, 0)$ toward the camera. In this image of $S_\alpha(t)$, α ranges from 0 at left to 0.35 at right. For all traces $\rho_s = 1$. At left is the diffuse reflectance trace $D(t)$ with $\rho_d = 1$ and $\rho_s = 0$. The temporal position t_d where the linear light source is directly above the surface is indicated by the cyan line, coinciding with the diffuse peak, and t_m indicates where the mirror angle occurs. Statistics μ and σ are computed for each specular lobe in $S_\alpha(t)$.

the surface. We assume the specular peak to be a sufficient distance away from the diffuse peak so that its contribution to the reflectance trace is negligible at t_d, and we measure the height of diffuse peak l_d as the height of the reflectance trace at t_d. To calculate ρ_d, we also need to know the amount of light falling on pixel (u, v), which we determine by including a strip of diffuse material of known reflectance $\rho_{standard}$ along the side of the object. We measure the height of the diffuse peak $l_{standard}$ for pixels on this known sample and then compute: $\rho_d = l_d \cdot \rho_{standard} / l_{standard}$.

This performs the appropriate correction since $l_{standard} / \rho_{standard}$ is a measurement of the irradiance at the object surface.

6.3 Estimating Specular Intensity ρ_s and Roughness α

Once ρ_d is calculated, we proceed to determine the specular reflectance parameters of the sample point. Since the diffuse lobe is broad, we cannot assume that its contribution will be negligible in the region of the specular peak. To correct for this, we subtract the contribution of the modeled diffuse lobe from the reflectance trace before fitting the specular peak as in Fig. 7 to obtain $\hat{I}_{u,v}(t)$. We obtain the shape of the diffuse lobe from the first row of the reflectance table (Fig. 6) scaled to match the observed height of the lobe.

With the diffuse lobe subtracted, the specular lobe may be examined independently. Specular peaks generally lie near the mirror angle t_m, but because of off-specular peaks are usually further away from t_d. We characterize the specular peak in terms of its mean μ, standard deviation σ, and total energy I_{avg} as follows:

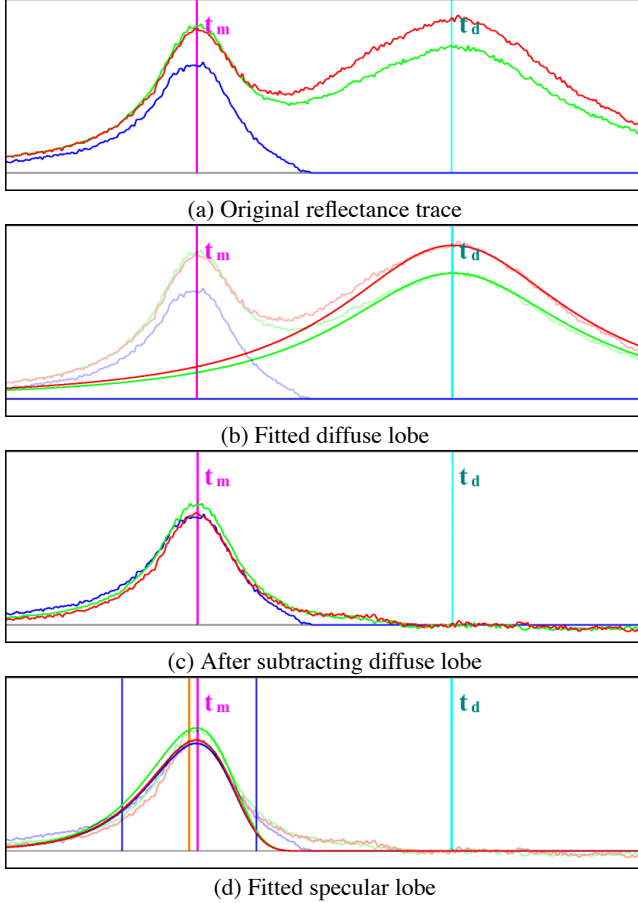$$\mu = \nu \frac{\sum t \cdot \hat{I}(t)}{\sum \hat{I}(t)} \tag{2}$$

$$\sigma^2 = \nu \frac{\sum (t - \mu)^2 \hat{I}(t)}{\sum \hat{I}(t)} \tag{3}$$

$$S = \nu \sum \hat{I}(t) \tag{4}$$

We place the limits of these summation to range from $t_m \pm \Delta$, which means that we consider the region of $\hat{I}(t)$ in the neighborhood of the t_m with a radius equal to the distance between t_d and t_m.

Because of slight perspective effects in the camera this region will be different widths for different pixels, so all values are normalized by the scaling factor $\nu = 1/(2\Delta)$.

With these characteristics calculated, we then look into our reflectance table to find a specular peak with the same σ as that observed in our specular peak. From this, we use the α value that produced that specular peak to be the specular roughness parameter of the pixel. To find the specular intensity ρ_s, we compute the ratio of the lobe's sum S to the sum of the specular lobe with roughness α in the table, and scale this value according to the same scaling factor used to scale the diffuse component as defined in Section 6.2.

(a) Original reflectance trace

(b) Fitted diffuse lobe

(c) After subtracting diffuse lobe

(d) Fitted specular lobe

Figure 7: **Fitting Diffuse and Specular Lobes** The reflectance fitting process determines the diffuse and specular parameters of a reflectance trace (a) from the yellow decal of our test object by choosing a ρ_d such that $\rho_d D(t)$ matches the height of the diffuse peak as in (b). Then, the fitted diffuse peak is subtracted from the reflectance trace to isolate the specular peak (c). The specular peak's standard deviation σ is computed and used to choose the appropriate specular roughness parameter α. Finally, ρ_s is determined as the ratio of the total energy of the specular peak compared to the energy of the specular sample $S_\alpha(t)$ in the reflectance table to produce the modeled specular peak (d).

7 Model Fitting for Non-Flat Surfaces

The reflectometry technique in the previous section assumes that the object surface is flat. If the object has surface normal variations, such as are found on embossed type, the surface will reflect light differently. Here we present modifications to the previous fitting procedure that allow the technique to recover a surface normal

map using two passes of the linear light source over the object at different diagonal angles. From these two passes it becomes possible to estimate variations in both components of the surface normal, which we describe in Section 7.2.

For non-flat surfaces, we also capture the larger-scale surface topography using the laser stripe system. We use the surface height map to tune the t_d and t_m values for each pixel, as well as to provide a displaced surface for rendering. Since the normal recovery technique of 7.2 assumes accurate values for t_d and t_s, we first describe in Section 7.1 the displacement estimation technique.

7.1 Estimating the Displacement d

As described in Section 3, our LLS capture device produces a red laser stripe which projects a vertical plane of light onto the object 38mm behind the light source. For each pixel (u,v), we detect the time at which the laser stripe crosses the corresponding point on the object by locating the pixel $I_{u,v}(t)$ which has the brightest red component. To obtain subpixel accuracy, we defocus the laser beam slightly and compute the peak of a quadratic function fit to the region near the laser peak. A space-time scanning technique [Curless and Levoy 1996] could also be employed.

To obtain the point P on the object surface corresponding to the pixel image coordinate (u,v), we intersect the camera ray corresponding to the pixel with the plane of the laser to obtain point P as in Fig. 8. The appropriate t_d and t_m values for point P can then be computed as before in Section 4.2.
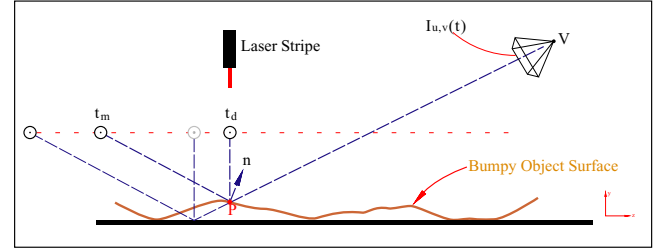
Figure 8: **Laser Stripe Scanning** Variations in surface height change the temporal location of the mirror angle, so to maintain registration a laser stripe is attached to the LLS to measure the surface shape of the object. Intersecting the camera ray VP with the plane of the laser at the time of the laser stripe crosses pixel (u,v) yields point Q on the surface. The displacement map is used to adjust the mirror angle position t_m for the raised surfaces, as well as to provide a geometric model of the object.

7.2 Estimating the Surface Normal $\hat{n}$

To estimate the surface normal, we scan the object twice with the LLS in its two different diagonal positions. Our estimation technique relies on the fact that if a pixel's surface normal $\hat{n}$ points toward the approaching light, its specular peak will occur earlier in $\hat{I}(t)$, and if it points away from the approaching LLS, its specular peak will occur later. The two light source directions thus allow us to estimate both components of the surface normal.

To determine how specular peaks respond to changes in surface normal, we render four different reflectance tables for the $\beta = 15°$ degree diagonal light source using the technique in Section 6.1. We render the first two reflectance tables with the surface normal set to $\pm\delta$ degrees rotated about the z axis (yielding a perturbation in the x component) and the second two tables $\pm\delta$ degrees about the x axis (yielding a perturbation in the z component). Reflectance tables for $\delta = 5°$ are shown in Fig. 9.

For each reflectance table, we locate the S_α lobe having the same standard deviation σ as the pixel's observed specular lobe in $\hat{I}(t)$.

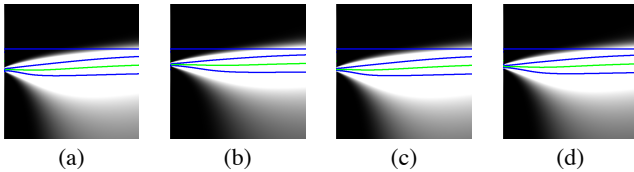<div style="text-align:center">(a)   (b)   (c)   (d)</div>

Figure 9: **Reflectance Tables for Varying Surface Normal** To analyze surface normal variations using the specular lobe, we render out $S_\alpha(t)$ reflectance tables for the $\beta = 15°$ LLS angle for surface normals pointing slightly (a) up, (b) down, (c) left, and (d) right from the light source. Perturbing the normal changes the time of the specular peak's mean μ, allowing us to estimate both components of a pixel's surface normal $\hat{n}$.

We then find the corresponding means $\mu_{x-}, \mu_{x+}, \mu_{z-}, \mu_{z+}$ of each of these lobes. We use these four means to produce a linear function $f_{+15}(\Delta x, \Delta z)$ mapping normal perturbation to mean. By symmetry, we can construct a similar function $f_{-15}(\Delta x, \Delta z)$ for $\beta = -15°$, by simply swapping the values of μ_{x-} and μ_{x+}. For the pixel being considered, we solve for the normal perturbation components Δx and Δz by equating $f_{+15}(\Delta x, \Delta z)$ and $f_{-15}(\Delta x, \Delta z)$ to the pixel's actual measured means for the two LLS runs, μ_{+15} and μ_{-15}. This yields the following 2×2 linear system:

$$\mu_{+15} = \frac{\mu_{x-} + \mu_{x+} + \mu_{z-} + \mu_{z+}}{4} + \frac{\Delta x}{2\delta}(\mu_{x+} - \mu_{x-}) + \frac{\Delta z}{2\delta}(\mu_{z+} - \mu_{z-})$$

$$\mu_{-15} = \frac{\mu_{x-} + \mu_{x+} + \mu_{z-} + \mu_{z+}}{4} + \frac{\Delta x}{2\delta}(\mu_{x-} - \mu_{x+}) + \frac{\Delta z}{2\delta}(\mu_{z+} - \mu_{z-})$$

Since this technique for computing surface normals depends on analyzing specular reflection, it works well on specular surfaces and poorly on diffuse surfaces. For this reason, we produce our final surface normal map by combining the surface normals estimated by specular reflection (Fig. 13(e) and (f), insets) with surface normals obtained by differentiating our laser-scanned displacement map d (Fig. 13(e) and (f)). To blend between these normal estimates, we found good results could be obtained by using the specular-derived normals when the ratio $\rho_s/\alpha > 0.5$, the displacement-derived normals when $\rho_s/\alpha < 0.2$, and a linear blend of the two estimates for intermediate values of ρ_s/α.

7.3   Adjusting for Surface Normal $\hat{n}$

Once we have estimated a surface normal for each pixel, we can re-estimate the diffuse and specular parameters based on this improved knowledge of the surface characteristics. In particular, we can adjust our model of the diffuse lobe to be rendered with this new surface normal. To facilitate this we pre-compute a 20×20 table of diffuse lobes $D'_{x,z}(t)$ for surface normals varying in x and z, and use bilinear interpolation to create a $D'(t)$ for the particular surface normal. From this we re-compute ρ_d and then produce a new specular reflectance trace $\hat{I}'(t) = I(t) - D'(t)$. From this new reflectance trace, we fir a new $S_\alpha(t)$ to obtain new parameters ρ_s and α. We find that one iteration of this procedure significantly increases the consistency of the specular roughness parameters.

8   Estimating the Translucency ρ_{trans}

Thin objects such as manuscripts can be translucent, which means that they will transmit light through the front when lit from behind. We solve for this translucency component ρ_{trans} by taking one additional image of the object $I_{backlit}$ as illuminated from behind by the light box, as well as an image of the light box with the object removed $I_{lightbox}$. A set of two such images are shown in Figure 10.
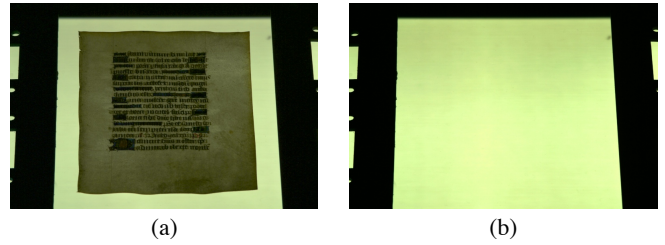

<div style="text-align:center">(a)   (b)</div>

Figure 10: **Capturing the Translucent Component** We measure the translucent component of the object by taking two images with the light box illuminated at the end of the LLS passes, one with and one without the object. Dividing the first image by the second yields the translucency component ρ_{trans} for each pixel and color in Fig. 13(f).

To determine the amount of transmitted light for each pixel of each color channel, we compute $\rho_{trans} = I_{backlit}/I_{lightbox}$ to obtain the percent of light transmitted at each pixel. In our reflectance model, we assume that any light transmitted through the material becomes diffused, so we do not measure ρ_{trans} for different incident and outgoing directions. This is a simplified model of translucency in comparison to subsurface scattering methods as in [Jensen et al. 2001], but we have found it to approximate the scattering characteristics of our test object sufficiently.

9   Results

In this section we present the recovered reflectance parameters for two objects and renderings of the captured models under various lighting conditions. Our renderer was created by implementing the Ward model as a floating-point pixel shader in Direct X 9 on an ATI Radeon 9700 graphics card, allowing all surface maps and rendered values to be computed radiometrically linearly in high dynamic range. Our first object seen in Fig. 11 was chosen to test the algorithm's ability to capture varying diffuse and specular properties for the flat object case. It consists of an 8.5×11 inch shiny black card printed with black laser toner, which provides a dull specularity. Attached are four glossy color decals, and we placed an oily thumbprint in the center of the shiny black area. Since there were no significant variations in surface normal, we scanned the object with a single pass of the LLS to obtain its spatially-varying reflectance parameters ρ_d, ρ_s, α shown in 11.

For validation, we found that the recovered ρ_d parameters properly ranged between 0 and 1 and were consistent with measurements taken with a calibrated spectroradiometer. We also generated renderings of the surface under point-light source illumination in our interactive renderer as in 12(b) which are largely consistent with the actual appearance of the object. One noticeable difference is the greater breadth of the real specular lobe in the shiny black material, which we believe to be an artifact of the Ward reflectance model not precisely matching the toe behavior of this particular specular material, even if the overall height and width of the lobe are consistent. This result suggests that investigating fitting to more elaborate reflectance models could be useful. Also apparent is a bright rim of light on the upper right of the inner circle not present in the real image. We determined that this behavior resulted when fitting the behaviors of pixels which cover a heterogenous area of both sharp and dull specular material - the fitted peak produces a broader lobe than the sharp specularity and a higher reflectance in the off-specular region, yielding reflected light in the off-specular region not exhibited by the heterogenous material. This result suggests that investigating fitting mixed specular models in regions of high specular variance.

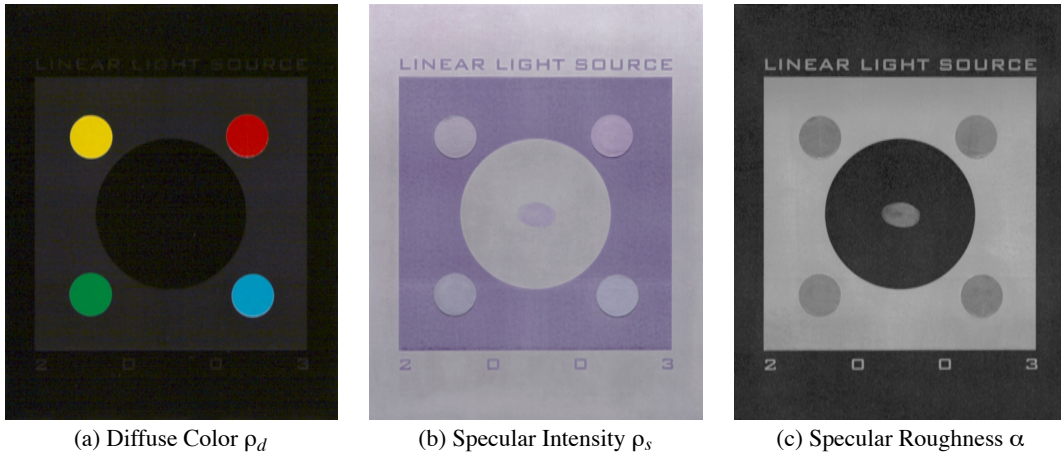For our second example we captured the reflectance properties of

(a) Diffuse Color ρ_d (b) Specular Intensity ρ_s (c) Specular Roughness α

Figure 11: **Recovered Reflectance Parameters for the Test Object**
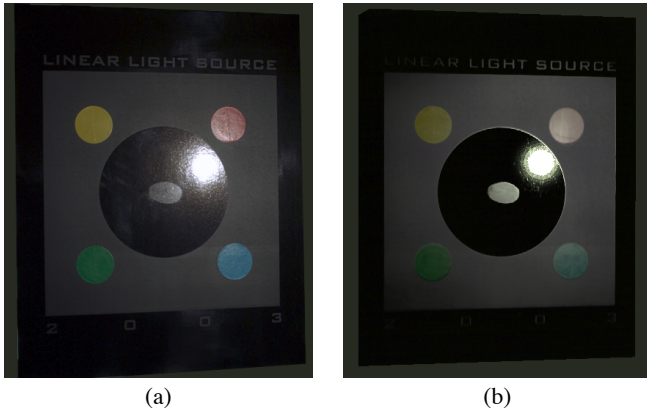

(a)   (b)

Figure 12: **Real and Synthetic Comparison** **(a)** Real photograph of test object. **(b)** Real-time rendering of the captured reflectance properties of the test object.
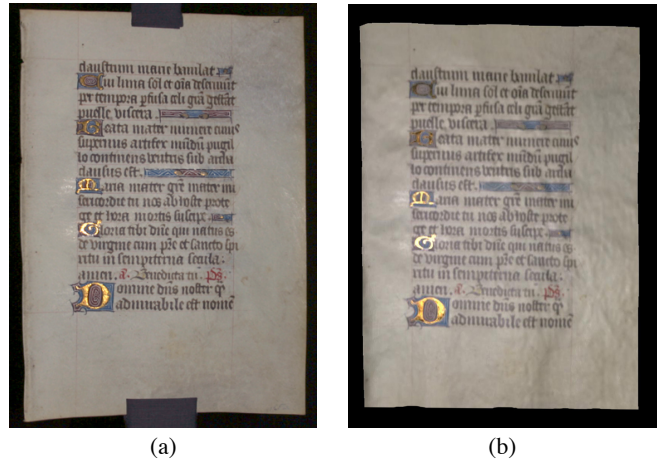

(a)   (b)

Figure 14: **Real and Synthetic Comparison** **(a)** Real photograph of the manuscript. **(b)** Real-time rendering of the modeled reflectance properties of the manuscript lit from a similar direction.

a 177×163mm page from a 15th-century illuminated manuscript seen in Fig. 14(a). The manuscript is written on vellum animal-skin parchment and features several embossed gold "illuminated" letters, which tested our method's ability to capture metallic specularity, surface normal variation, and translucency. Since the manuscript did not lie flat on the light box, it also exercised our technique's ability to capture and compensate for varying surface displacement. We captured two passes of the LLS at $\pm15°$ at 0.9 mm intervals. The manuscript's recovered ρ_d, ρ_s, α, d, $\hat{n}$, and ρ_{trans} parameters are shown in 13.

The algorithm found that the manuscript's vellum surface exhibited a dull specularity, and as expected found that the gold letters had a greater specular intensity and lower roughness than the vellum. The ρ_{trans} values indicated that the manuscript transmits approximately 17% of the red light, 14% of the green light, and 9% of the blue light. Renderings produced of the manuscript under various point light directions as in Fig. 14 were largely consistent with the actual appearance of the manuscript under similar lighting conditions. For some angles our modeled manuscript appeared to be slightly more specular than the original; a possible explanation is that the vellum exhibits a forward-scattering component which is mistaken as additional specularity by our lighting capture system. To test our surface normal recovery, we compared the appearance of the illuminated letter "D" in for real and synthetic images (Fig. 15). We found the normals to be largely consistent though the re-

flections had less texture and were somewhat duller; these results suggest that using a narrower LLS might improve the results for sharp specular surfaces. The video accompanying this paper shows real and synthetic animation of the manuscript under point light source illumination.

To show a lighting configuration that would exhibit both the manuscript's reflective and translucent qualities, we implemented a multipass renderer to illuminate our manuscript with an omnidirectional lighting environment acquired inside a cathedral. Results of rendering both the manuscript and the test object in this lighting environment are shown in Fig. 16.

10 Discussion and Future Work

The capabilities and limitations of our linear light source reflectometry system suggest several avenues of work for future versions of the system.

In this work we recover trichromatic measurements of the object's diffuse, specular, and transmissive reflectance properties, owing to the fact that we use a standard RGB digital camera. Such an approximation to a description of the full spectral reflectance can cause errors in rendering the object under novel lighting conditions [Borges 1991]. This is a concern since a common application of
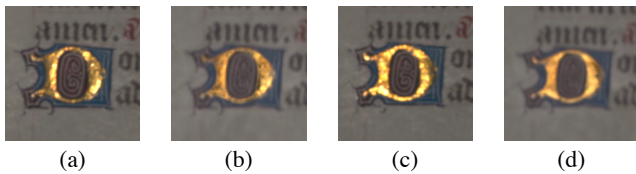
(a) Diffuse color ρ_d (b) Specular intensity ρ_s (c) Specular roughness α
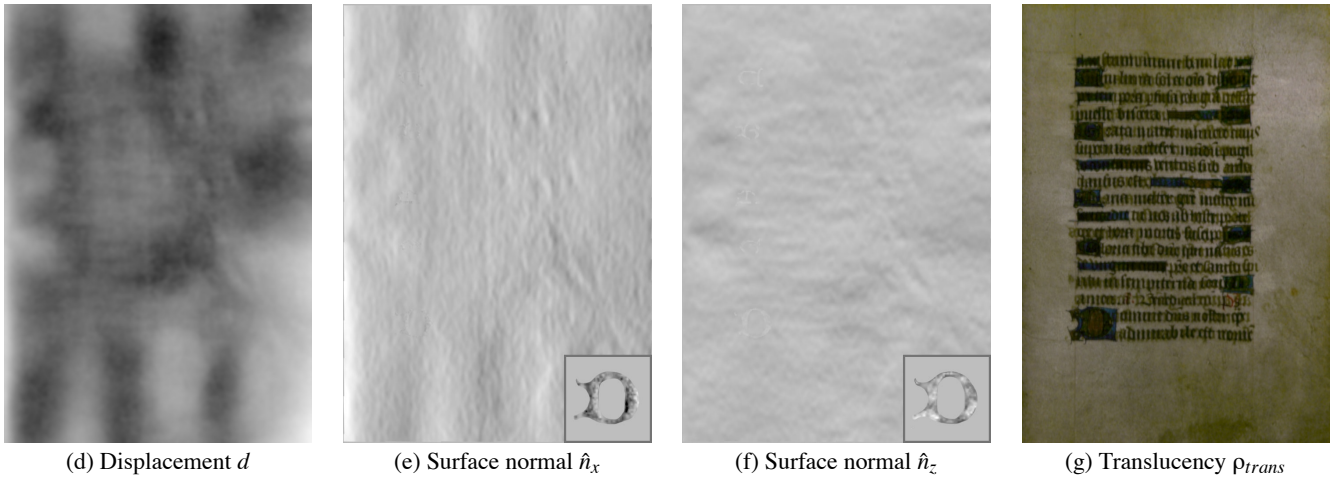
(d) Displacement d (e) Surface normal $\hat{n}_x$ (f) Surface normal $\hat{n}_z$ (g) Translucency ρ_{trans}

Figure 13: **Recovered Reflectance Parameters for an Illuminated Manuscript**

(a) (b) (c) (d)

Figure 15: **Real and Estimated Surface Normals** (a) and (c) show real photographs of the illuminated "D" under two directions of a point light source; (b) and (d) show renderings of the recovered reflectance properties under approximately the same illumination. The renderings show good correspondence between the real and recovered surface normals, but indicate that some smoothing of high-frequency components occurs in the recovery process.
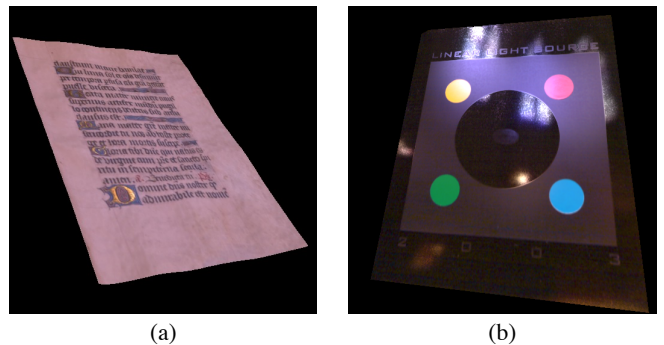
(a) (b)

Figure 16: **Renderings with Environmental Illumination** **(a)** Manuscript lit by the Grace Cathedral environment. **(b)** Test object lit by the Grace Cathedral environment.

capturing a real-world object would be to visualize it under both natural and artificial illumination conditions. In addition, having spectral reflectance data can be used in the analysis of the materials used to create a particular artifact. In the future one could capture multispectral object reflectance, for example using a standard camera outfitted with a liquid crystal tunable filter [Gat 1998].

One application of the LLS device would be to recover reflectance models of real-world object surfaces such as floor tiles and cobblestones. Similar to current texture map libraries, captur-

ing a variety of real-world materials could produce a database of realistic models easily inserted into computer-generated scenes.

A additional avenue of work is to extend the technique to capturing anisotropic reflectance characteristics. Many types of brushed metals and fabrics exhibit specular lobes whose roughness parameter α varies depending on the object orientation. Using a third pass of the LLS at an additional angle holds promise for estimating an object's per-pixel anisotropic reflectance parameters.

Finally, the design of our current LLS apparatus limits its application to relatively flat objects such as printed materials and fabrics. It would be desirable to extend these techniques to be able to capture the spatially varying reflectance characteristics of completely 3D objects such as human faces and painted sculptures. By improving the laser scanning system and assembling multiple LLS scans of an object from different angles, it could be possible to create 3D models of real-world objects which faithfully reproduce the object's reflectance properties.

11 Conclusion

In this paper we have presented a system for capturing the spatially varying diffuse and specular reflectance properties real-world object surfaces using a linear light source, and have applied our technique to a test object and an illuminated manuscript. The apparatus can be built at low cost, requires modest scanning time and computation, and produces realistic and largely faithful models of the real object's reflectance characteristics. Augmented by some of the techniques proposed as future work, this form of object capture could become a useful tool for computer graphics practitioners wanting to capture the realism of real-world reflectance properties, as well as to museum conservators and researchers wishing to digitally preserve, analyze, and communicate cultural artifacts.

References

BORGES, C. F. 1991. Trichromatic approximation for computer graphics illumination models. In *Computer Graphics (Proceedings of SIGGRAPH 91)*, vol. 25, 101–104.

CHUANG, Y.-Y., ZONGKER, D. E., HINDORFF, J., CURLESS, B., SALESIN, D. H., AND SZELISKI, R. 2000. Environment matting extensions: Towards higher accuracy and real-time capture. In *Proceedings of SIGGRAPH 2000*, 121–130.

COOK, R. L., AND TORRANCE, K. E. 1981. A reflectance model for computer graphics. *Computer Graphics (Proceedings of SIGGRAPH 81) 15*, 3 (August), 307–316.

CURLESS, B., AND LEVOY, M. 1996. Better optical triangulation through spacetime analysis. In *Proceedings of SIGGRAPH 95*, International Conference on Computer Vision, 987–994.

DANA, K. J., GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1997. Reflectance and texture of real-world surfaces. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, 151–157.

DEBEVEC, P. E., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH 97*, 369–378.

DEBEVEC, P. E., TAYLOR, C. J., AND MALIK, J. 1996. Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach. In *SIGGRAPH 96*, 11–20.

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.

GAT, N. 1998. Real-time multi- and hyper-spectral imaging for remote sensing and machine vision: an overview. In *Proc. 1998 ASAE Annual International Mtg.*

HALSTEAD, M., BARSKY, B. A., KLEIN, S., AND MANDELL, R. 1996. Reconstructing curved surfaces from specular reflection patterns using spline surface fitting of normals. In *Proceedings of SIGGRAPH 96*, Computer Graphics Proceedings, Annual Conference Series, 335–342.

HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. *Proceedings of SIGGRAPH 93* (August), 165–174.

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.

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., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2001. Image-based reconstruction of spatially varying materials. In *Rendering Techniques 2001: 12th Eurographics Workshop on Rendering*, 103–114.

LEVOY, M., PULLI, K., CURLESS, B., RUSINKIEWICZ, S., KOLLER, D., PEREIRA, L., GINZTON, 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.

MALZBENDER, T., GELB, D., AND WOLTERS, H. 2001. Polynomial texture maps. *Proceedings of SIGGRAPH 2001* (August), 519–528.

MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., TORRANCE, K. E., AND GREENBERG, D. P. 1999. Image-based BRDF measurement including human skin. *Eurographics Rendering Workshop 1999* (June).

MARSCHNER, S. 1998. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University.

MCALLISTER, D. K. 2002. *A Generalized Surface Appearance Representation for Computer Graphics*. PhD thesis, University of North Carolina at Chapel Hill.

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.

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).

NISHITA, T., OKAMURA, I., AND NAKAMAE, E. 1985. Shading models for point and linear sources. *ACM Transactions on Graphics 4*, 2 (April), 124–146.

OREN, M., AND NAYAR, S. K. 1994. Generalization of Lambert's reflectance model. *Proceedings of SIGGRAPH 94* (July), 239–246.

POULIN, P., AND AMANATIDES, J. 1991. Shading and shadowing with linear light sources. *Computers & Graphics 15*, 2, 259–265.

POULIN, P., AND FOURNIER, A. 1990. A model for anisotropic reflection. In *Computer Graphics (Proceedings of SIGGRAPH 90)*, vol. 24, 273–282.

RAMAMOORTHI, R., AND HANRAHAN, P. 2001. A signal-processing framework for inverse rendering. In *Proceedings of ACM SIGGRAPH 2001*, ACM Press / ACM SIGGRAPH, Computer Graphics Proceedings, Annual Conference Series, 117–128. ISBN 1-58113-292-1.

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.

TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory for off-specular reflection from roughened surfaces. *Journal of Optical Society of America 57*, 9.

WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *SIGGRAPH 92*, 265–272.

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.

ZONGKER, D. E., WERNER, D. M., CURLESS, B., AND SALESIN, D. H. 1999. Environment matting and compositing. *Proceedings of SIGGRAPH 99* (August), 205–214.