

Acquiring the Reflectance Field of a Human Face

Paul Debevec[†] Tim Hawkins[†] Chris Tchou[†] Haarm-Pieter Duiker[†] Westley Sarokin[†]
and Mark Sagar[‡]

[†]University of California at Berkeley¹ [‡]LifeF/X, Inc.

ABSTRACT

We present a method to acquire the reflectance field of a human face and use these measurements to render the face under arbitrary changes in lighting and viewpoint. We first acquire images of the face from a small set of viewpoints under a dense sampling of incident illumination directions using a light stage. We then construct a reflectance function image for each observed image pixel from its values over the space of illumination directions. From the reflectance functions, we can directly generate images of the face from the original viewpoints in any form of sampled or computed illumination. To change the viewpoint, we use a model of skin reflectance to estimate the appearance of the reflectance functions for novel viewpoints. We demonstrate the technique with synthetic renderings of a person's face under novel illumination and viewpoints.

Categories and subject descriptors: I.2.10 [Artificial Intelligence]: Vision and Scene Understanding - *intensity, color, photometry and thresholding*; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism - *color, shading, shadowing, and texture*; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism - *radiosity*; I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture - *radiometry, reflectance, scanning*; I.4.8 [Image Processing]: Scene Analysis - *photometry, range data, sensor fusion*. **Additional Key Words and Phrases:** facial animation; image-based modeling, rendering, and lighting.

1 Introduction

Creating realistic renderings of human faces has been an endeavor in computer graphics for nearly three decades [28] and remains a subject of current interest. It is a challenging problem due to the complex and individual shape of the face, the subtle and spatially varying reflectance properties of skin, and the complex deformations of the face during movement. Compounding the problem, viewers are extremely sensitive to the appearance of other people's faces.

Recent work has provided solutions to the problems of geometrically modeling and animating faces. 3D photography techniques, such as the Cyberware scanner, can acquire accurate geometric

models of individual faces. Work to animate facial expressions through morphing [2, 4, 29], performance-driven animation [38], motion capture [14], and physics-based simulation [34, 20, 30] has produced examples of realistic facial motion.

An outstanding problem is the lack of a method for capturing the spatially varying reflectance characteristics of the human face. The traditional approach of texture-mapping a photograph of a face onto a geometric model usually fails to appear realistic under changes in lighting, viewpoint, and expression. The problem is that the reflectance properties of the face are complex: skin reflects light both diffusely and specularly, and both of these reflection components are spatially varying. Recently, skin reflectance has been modeled using Monte Carlo simulation [16], and several aggregate reflectance descriptions have been recorded from real people [22], but there has not yet been a method of accurately rendering the complexities of an individual's facial reflectance under arbitrary changes of lighting and viewpoint.

In this paper we develop a method to render faces under arbitrary changes in lighting and viewing direction based on recorded imagery. The central device in our technique is a *light stage* (Fig. 2) which illuminates the subject from a dense sampling of directions of incident illumination. During this time the subject's appearance is recorded from different angles by stationary video cameras.

From this illumination data, we can immediately render the subject's face from the original viewpoints under any incident field of illumination by computing linear combinations of the original images. Because of the additive nature of light [5], this correctly reproduces all of the effects of diffuse and specular reflection as well as interreflections between parts of the face. We demonstrate this technique by rendering faces in various forms of natural illumination captured in real-world environments, and discuss how this process can be performed directly from compressed images.

In the second part of this paper we present a technique to extrapolate a complete reflectance field from the acquired data which allows us to render the face from novel viewpoints. For this acquire a geometric model of the face through structured lighting, which allows us to project the appearance from the original viewpoints onto the geometry to render from novel viewpoints. However, re-rendering directly from such projected images does not reproduce view-dependent reflection from the face; most notably, the specular components need to shift position according to the rendered viewpoint.

To reproduce these view-dependent effects, we use a skin reflectance model to extrapolate the reflectance observed by the cameras to that which would be observed from novel viewpoints. The model is motivated by a set of in-plane reflectance measurements of a patch of skin using polarizers on the light and the camera to separate the reflection components. This model allows us to separate the specular and sub-surface reflection components of the light stage data using chromaticity analysis, and then to transform each reflectance component into how it would appear from a novel viewpoint. Using this technique, we can realistically render the face from arbitrary viewpoints and in arbitrary lighting.

The rest of this paper is organized as follows. In the next section we review related work and discuss the reflectance field. In Sec-

¹Computer Science Division, University of California at Berkeley. Email: {debevec,tsh,ctchou,duiker,wsarokin}@cs.berkeley.edu, msagar@lifefx.com. For more information see <http://www.debevec.org/>

tion 3 we describe the light stage and how we synthesize physically correct images of the subject under arbitrary illumination. In Section 4 we develop a model of skin reflectance and use it to render the face from novel viewpoints under arbitrary illumination. We discuss future work in Section 5 and conclude in Section 6.

2 Background and Related Work

In this section we give an overview of related work in the areas of facial modeling and animation, reflectometry, and image-based modeling and rendering. We conclude with a description of the reflectance field.

Facial Modeling and Animation Since the earliest work in facial modeling and animation [28], generating realistic faces has been a central goal. 3D photography techniques for acquiring facial geometry, such as the laser-triangulation based scanners made by Cyberware, have been a helpful development. Such techniques often also photograph a texture map for the face at the time of the scan, which can be projected onto the face to produce renderings. However, using such texture maps usually falls short of producing photorealistic renderings since the map is illumination-dependent and does not capture directionally varying reflectance properties. Other work estimates facial models directly from images: [11, 29, 3] recover geometry by fitting morphable facial models; [11, 3] use the models to estimate albedo maps but do not consider specular properties. [29] produces view-dependent reflectance under the original illumination conditions through view-dependent texture mapping [10].

Several techniques have been used to animate facial models; [2, 4, 29] blend between images in different expressions to produce intermediate expressions. [38, 14] use the captured facial motion of a real actor to drive the performance of a synthetic one. Physics-based simulation techniques [34, 30, 40, 20] have helped animate the complex deformations of a face in its different expressions.

Reflectometry Reflectometry is the measurement of how materials reflect light, or, more specifically, how they transform incident illumination into radiant illumination. This transformation can be described by the four-dimensional bi-directional reflectance distribution function, or BRDF, of the material measured [25]. Several efforts have been made to represent common BRDFs as parameterized functions called *reflectance models* [35, 6, 37, 27, 19].

Hanrahan and Krueger [16] developed a parameterized model for reflection from layered surfaces due to subsurface scattering, with human skin as a specific case of their model. Their model of skin reflectance was motivated by the optical properties of its surface, epidermal, and dermal layers [36]. Each layer was given several parameters according to its scattering properties and pigmentation, and a Monte Carlo simulation of the paths light might take through the skin surfaces produced renderings exhibiting a variety of qualitatively skin-like reflectance properties. The authors selected the reflectance properties manually, rather than acquiring them from a particular individual. The authors also simulated a uniform layer of oil over the face to produce specular reflection; in our work we acquire a reflectance model that reproduces the varying diffuse and specular properties over the skin.

Much work has been done to estimate reflectance properties of surfaces based on images taken under known lighting. [37] and [17] presented techniques and apparatus for measuring anisotropic reflectance of material samples; [7] applied reflectometry techniques to the domain of textured objects. In our work, we leverage being able to separate reflection into diffuse and specular components. This separation can be done through colorspace analysis [31] as well as a combined analysis of the color and polarization of the reflected light [24]; in our work we make use of both color and polarization. [32] used object geometry and varying light directions to derive diffuse and specular parameters for a coffee mug; [41]

used an inverse radiosity method to account for mutual illumination in estimating spatially varying diffuse and piecewise constant specular properties within a room.

Marschner, Westin, Lafortune, Torrance, and Greenberg [22] recently acquired the first experimental measurements of living human facial reflectance in the visible spectrum. The authors photographed the forehead of their subjects under constant point-source illumination and twenty viewing directions, and used the curvature of the forehead to obtain a dense set of BRDF samples. From these measurements, they derived a non-parametric isotropic BRDF representing the average reflectance properties of the surface of the forehead. In our work, we have chosen the goal of reproducing the spatially varying reflectance properties across the surface of the face; as a result, we sacrifice the generality of measuring a full BRDF at each surface point and use models of specular and diffuse reflectance to extrapolate the appearance to novel viewpoints.

Image-Based Modeling and Rendering In our work we leverage several principles explored in recent work in image-based modeling and rendering. [26, 15] showed how correct views of a scene under different lighting conditions can be created by summing images of the scene under a set of basis lighting conditions; [39] applied such a technique to create light fields [21, 13] with controllable illumination. [42] showed that by illuminating a shiny or refractive object with a set of coded lighting patterns, it could be correctly composited over an arbitrary background by determining the direction and spread of the reflected and refracted rays. [8] presented a technique for capturing images of real-world illumination and using this lighting to illuminate synthetic objects; in this paper we use such image-based lighting to illuminate real faces.

2.1 Definition of the Reflectance Field

The light field [12, 21], plenoptic function [1], and lumigraph [13] all describe the presence of light within space. Ignoring wavelength and fixing time, this is a five dimensional function of the form $P = P(x, y, z, \theta, \phi)$. The function represents the radiance leaving point (x, y, z) in the direction (θ, ϕ) .

[21, 13] observed that when the viewer is moving within unoccluded space, the light field can be described by a four-dimensional function. We can characterize this function as $P' = P'(u, v, \theta, \phi)$, where (u, v) is a point on a closed surface A and (θ, ϕ) is a direction as before. A light field parameterized in this form induces a five-dimensional light field in the space outside of A : if we follow the ray beginning at (x, y, z) in the direction of (θ, ϕ) until it intersects A at (u, v) , we have $P(x, y, z, \theta, \phi) = P'(u, v, \theta, \phi)$. In an example from [21] A was chosen to be a cube surrounding the object; in an example from [13] A was chosen to be the visual hull of the object. We can also consider the viewer to be inside of A observing illumination arriving from outside of A as shown in [21].

Images generated from a light field can have any viewing position and direction, but they always show the scene under the same lighting. In general, each field of incident illumination on A will induce a different field of radiant illumination from A . We can represent the radiant light field from A under every possible incident field of illumination as an eight-dimensional *reflectance field*:

$$R = R(R_i; R_r) = R(u_i, v_i, \theta_i, \phi_i; u_r, v_r, \theta_r, \phi_r) \quad (1)$$

Here, $R_i(u_i, v_i, \theta_i, \phi_i)$ represents the incident light field arriving at A and $R_r(u_r, v_r, \theta_r, \phi_r)$ represents the radiant light field leaving A (see Figure 1(a)). Except that we do not presume A to be coincident with a physical surface, the reflectance field is equivalent to the bidirectional scattering-surface reflectance distribution function S , or BSSRDF, described in Nicodemus et al. [25]. Paraphrasing [25], this function “provides a way of quantitatively expressing the connection between reflected flux leaving (u_r, v_r) in

a given direction and the flux incident at (u_i, v_i) in another given direction.”

In this work we are interested in acquiring reflectance fields of real objects, in particular human faces. A direct method to acquire the reflectance field of a real object would be to acquire a set of light fields of an object $R_r(u_r, v_r, \theta_r, \phi_r)$ for a dense sampling of incident beams of illumination from direction (θ_i, ϕ_i) arriving at the surface A at (u_i, v_i) . However, recording a four dimensional light field for every possible incident ray of light would require a ponderous amount of acquisition time and storage. Instead, in this work we acquire only *non-local reflectance fields* where the incident illumination field originates far away from A so that $R_i(u_i, v_i, \theta_i, \phi_i) = R_i(u'_i, v'_i, \theta_i, \phi_i)$ for all u_i, v_i, u'_i, v'_i . Thus a non-local reflectance field can be represented as $R' = R'(\theta_i, \phi_i; u_r, v_r, \theta_r, \phi_r)$. This reduces the representation to six dimensions, and is useful for representing objects which are some distance from the rest of the scene. In Section 3.4 we discuss using a non-local reflectance field to produce local illumination effects.

In this work we extrapolate the complete field of radiant illumination from data acquired from a sparse set of camera positions (Section 3) and choose the surface A to be a scanned model of the face (Figure 1(b)), yielding a surface reflectance field analogous to a surface light field [23]. A model of skin reflectance properties is used to synthesize views from arbitrary viewpoints (Section 4).




Figure 1: **The Reflectance Field** (a) describes how a volume of space enclosed by a surface A transforms an incident field of illumination $R_i(u_i, v_i, \theta_i, \phi_i)$ into a radiant field of illumination $R_r(u_r, v_r, \theta_r, \phi_r)$. In this paper, we acquire a non-local reflectance field (b) in which the incident illumination consists solely of directional illumination (θ_i, ϕ_i) . We choose A to be coincident with the surface of the face, yielding a surface reflectance field which allows us to extrapolate the radiant light field $R_r(u_r, v_r, \theta_r, \phi_r)$ from a sparse set of viewpoints.

3 Re-illuminating Faces

The goal of our work is to capture models of faces that can be rendered realistically under any illumination, from any angle, and, eventually, with any sort of animated expression. The data that we use to derive our models is a sparse set of viewpoints taken under a dense set of lighting directions. In this section, we describe the acquisition process, how we transform each facial pixel location into a reflectance function, and how we use this representation to render the face from the original viewpoints under any novel form of illumination. In the following section we will describe how to render the face from new viewpoints.

3.1 The Light Stage

The light stage used to acquire the set of images is shown in Fig. 2. The subject sits in a chair which has a headrest to help keep his or her head still during the capture process. Two digital video cameras view the head from a distance of approximately three meters; each captures a view of the left or the right side of the face. A spotlight, calibrated to produce an even field of illumination across the




Figure 2: **The Light Stage** consists of a two-axis rotation system and a directional light source. The outer black bar θ is rotated about the central vertical axis and the inner bar ϕ is lowered one step for each θ rotation. Video cameras placed outside the stage record the face’s appearance from the left and right under this complete set of illumination directions, taking slightly over a minute to record. The axes are operated manually by cords and an electronic audio signal triggered by the θ axis registers the video to the illumination directions. The inset shows a long-exposure photograph of the light stage in operation.

subject’s head, is affixed at a radius of 1.5 meters on a two-axis rotation mechanism that positions the light at any azimuth θ and any inclination ϕ . In operation, the light is spun about the θ axis continuously at approximately 25 rpm and lowered along the ϕ axis by $\frac{180}{32}$ degrees per revolution of θ (the cord controlling the ϕ axis is marked at these increments). The cameras, which are calibrated for their flat-field response and intensity response curve, capture frames continuously at 30 frames per second which yields 64 divisions of θ and 32 divisions of ϕ in approximately one minute, during which our subjects are usually capable of remaining still. A future version could employ high-speed cameras running at 250 to 1000 frames per second to lower the capture time to a few seconds. Some source images acquired with the apparatus are shown in Fig. 5.

3.2 Constructing Reflectance Functions

For each pixel location (x, y) in each camera, we observe that location on the face illuminated from 64×32 directions of θ and ϕ . From each pixel we form a slice of the reflectance field called




Figure 3: Reflectance Functions for a Face This mosaic is formed from the reflectance functions of a 15×44 sampling of pixels from the original 480×720 image data. Each 64×32 reflectance function consists of the corresponding pixel location's appearance under two thousand lighting directions distributed throughout the sphere. The inset shows the same view of the face under a combination of three lighting directions. The functions have been brightened by a factor of four from the original data.

a reflectance function $R_{xy}(\theta, \phi)$ corresponding to the ray through that pixel. Note that we are using the term “reflectance” loosely as true reflectance divides out the effect of the foreshortening of incident light. However, since the surface normal is unknown, we do not make this correction. If we let the pixel value at location (x, y) in the image with illumination direction (θ, ϕ) be represented as $L_{\theta, \phi}(x, y)$, then we have simply:

$$R_{xy}(\theta, \phi) = L_{\theta, \phi}(x, y) \quad (2)$$

Fig. 3 shows a mosaic of reflectance functions for a particular viewpoint of the face. Four of these mosaics are examined in detail in Fig. 4. The reflectance functions exhibit and encode the effects of diffuse reflection, specular reflection, self-shadowing, translucency, mutual illumination, and subsurface scattering.

3.3 Re-illuminating the Face




Figure 4: **A Sampling of Facial Reflectance Functions** The above reflectance functions appear in the mosaic of Fig. 3. The middle of each function corresponds to the pixel being illuminated from the direction of the camera; as one moves within the reflectance function the light direction moves in the same manner. Reflectance function (a) is taken from the forehead toward the right of the image, and exhibits a noticeable specular lobe as well as an unoccluded diffuse lobe. (b) from the right of the underside of the jaw exhibits a weaker specular component and some self-shadowing at lower lighting angles caused by the shoulder blocking the light source. (c) from the subject’s cheek to the right and below the nose exhibits a mild specular reflection and shadowing due to the nose in the upper left. (d) sampled from a pixel inside the pinna of the ear exhibits illumination from diffuse reflection and from light scattering through the tissue when illuminated from behind. Each function exhibits a thin black curve in its lower half where the phi axis bar occasionally obscures the view of the face, and a bright spot due to lens flare where the light points into the camera. These regions appear in the same places across images and are ignored in the lighting analysis.

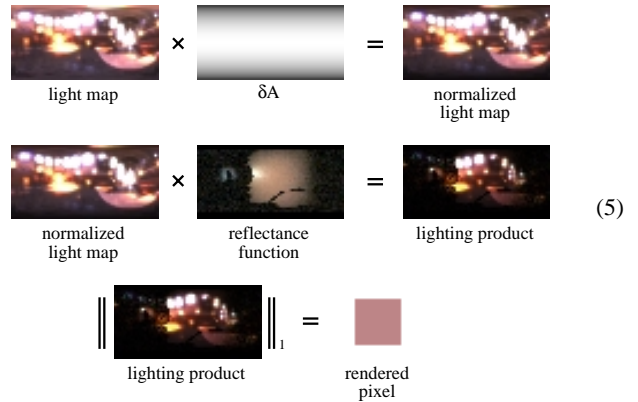
Suppose that we wish to generate an image of the face in a novel form of illumination. Since each $R_{xy}(\theta, \phi)$ represents how much light is reflected toward the camera by pixel (x, y) as a result of illumination from direction (θ, ϕ) , and since light is additive, we can compute an image of the face $\hat{L}(x, y)$ under any combination of the original light sources $L_i(\theta, \phi)$ as follows:

$$\hat{L}(x, y) = \sum_{\theta, \phi} R_{xy}(\theta, \phi) L_i(\theta, \phi) \quad (3)$$

Each color channel is computed separately using the above equation. Since the light sources densely sample the viewing sphere, we can represent any form of sampled incident illumination using this basis. In this case, it is necessary to consider the solid angle δA covered by each of the original illumination directions:

$$\hat{L}(x, y) = \sum_{\theta, \phi} R_{xy}(\theta, \phi) L_i(\theta, \phi) \delta A(\theta, \phi) \quad (4)$$

For our data, $\delta A(\theta, \phi) = \sin \phi$; the light stage records more samples per solid angle near the poles than at the equator. Equation 5 shows the computation of Equation 4 graphically. First, the map of incident illumination (filtered down to the 64×32 (θ, ϕ) space) is normalized by the map of $\delta A(\theta, \phi)$. Then, the resulting map is multiplied by the pixel’s reflectance function. Finally, the pixel values of this product are summed to compute the re-illuminated pixel value. These equations assume the light stage’s light source is white and has unit radiance; in practice we normalize the reflectance functions based on the light source color. Figure 6 shows a face synthetically illuminated with several forms of sampled and synthetic illumination using this technique.



Writing the re-illumination equation of Equation 4 as the sum of the product of two 64×32 images allows us to gain efficiency in both storage and computation using the techniques presented by Smith and Rowe [33] by computing the product directly on JPEG-compressed versions of the images. This can reduce both storage and computation by a factor of twenty while maintaining good image quality.

3.4 Discussion

Since each rendered image can also be represented as a linear combination of the original images, all of the proper effects of non-diffuse reflectance, mutual illumination, translucency, and subsurface scattering are preserved, as noted in [26]. The 64×32 set of illumination directions used is somewhat coarse; however, the reflectance functions are generally not aliased at this resolution, which implies that when the light maps are also properly filtered down to 64×32 there will be no aliasing in the resulting renderings. The place where the reflectance functions do become aliased is where there is self-shadowing; the expected result of this is that one would see somewhat stairstepped shadows in harsh lighting situations. Such effects could be smoothed by using an area light source to illuminate the subject.

Since this technique captures slices of a non-local reflectance field, it does not tell us how to render a person under dappled light or in partial shadow. A technique that will in many cases produce reasonable results is to illuminate different pixels of the face using different models of incident illumination; however, this will no longer produce physically valid images because changes to the indirect illumination are not considered. As an example, consider rendering a face with a shaft of light hitting just below the eye. In reality, the light below the eye would throw indirect illumination on the underside of the brow and the side of the nose; this technique would not capture this effect.




Figure 5: **Light Stage Images** Above are five of the 2048 images taken by one camera during a run of the light stage. The pixel values of each location on the face under the 2048 illumination directions are combined to produce the mosaic images in Fig. 3. Below each image is the impulse light map that would generate it.




Figure 6: **Face Rendered under Sampled Illumination** Each of the above images shows the face synthetically illuminated with novel lighting, with the corresponding light map shown below. Each image is created by taking the dot product of each pixel’s reflectance function with the light map. The first four illumination environments are light probe measurements acquired from real-world illumination (see [8]) recorded as omnidirectional high dynamic range images; the rightmost lighting environment is a synthetic test case.

A person’s clothing reflects indirect light back onto the face, and our capture technique reproduces the person’s appearance in whatever clothing they were wearing during the capture session. If we need to change the color of the person’s clothing (for example, to place a costume on a virtual actor), we can record the subject twice, once wearing white clothing and once with black clothing. Subtracting the second image from the first yields an image of the indirect illumination from the clothing, which can then be tinted to any desired color and added back in to the image taken with the black clothing; this process is illustrated in Figure 7.

By recording the light stage images in high dynamic range [9] and using the process of environment matting [42], we can apply this technique to translucent and refractive objects and reproduce the appearance of the environment in the background; this process is described in the Appendix.

4 Changing the Viewpoint

In this section we describe our technique to extrapolate complete reflectance fields from the reflectance field slices acquired in Section 3, allowing us to render the face from arbitrary viewpoints as well as under arbitrary illumination. In our capture technique, we

observe the face under a dense set of illumination conditions but from only a small set of viewpoints. To render the face from a novel viewpoint, we must resynthesize the reflectance functions to appear as they would from the new viewpoint.

To accomplish this, we make use of a skin reflectance model which we introduce in Section 4.1. This model is used to guide the shifting and scaling of measured reflectance function values as the viewpoint changes. As such, our technique guarantees that the resynthesized reflectance function will agree exactly with the measured data if the novel viewpoint is the same as the viewpoint for data capture.

The resynthesis technique requires that our reflectance functions be decomposed into specular and diffuse (subsurface) components. Section 4.2 describes this separation process. Section 4.3 describes the re-synthesis of a reflectance function for a new viewpoint. Section 4.4 discusses the technique in the context of shadowing and mutual illumination effects. Section 4.5 explains the method used to produce renderings of the entire face using resynthesized reflectance functions.




Figure 7: **Modeling indirect light from clothing** Indirect reflectance from the subject's clothing can be modeled by recording the subject wearing both white (a) and black (b) clothing (we drape the white clothing on the subject and pull it away to reveal the black clothing.) (a) exhibits indirect lighting on the neck and beneath the chin and nose. Correct renderings of the person wearing any color clothing can be created by adding a tinted version of (a) minus (b) to (b). Using this method, (c) shows the subject with the indirect light she would receive from green clothing.

4.1 Investigating Skin Reflectance

In this section we consider the reflectance properties of skin, and describe our data-driven skin reflectance model. The model is intended to capture the behavior of skin, but could be useful for a wider class of surfaces.

Following [16], we note that the light reflected from the skin can be decomposed into two components: a specular component consisting of light immediately reflected at the index of refraction transition at the air-oil interface (see Figure 8), and a non-Lambertian diffuse component consisting of light transmitted through the air-oil interface that, after some number of subsurface scattering interactions, is transmitted from the oil layer to air.

We first investigated the general behavior of these two components. As shown in Figure 8, light which reflects specularly off the skin will maintain the polarization of the incident light; however, light which emerges from below the surface will have been depolarized by scattering interactions. Taking advantage of this fact, we can separate the reflection components by placing linear polarizers on both the light source and the camera¹. Figure 9 shows separated specular and diffuse reflection components of a face using this technique.




Figure 8: **Skin Reflectance** Light reflecting from skin must have reflected specularly off the surface (a) or at some point entered one or more of the scattering layers (b, c, d). If the incident light is polarized, the specularly reflected light will maintain this polarization; however, light which scatters within the surface becomes depolarized. This allows reflection components to be separated as in Figures 9 and 10.

Using this effect, we carried out an in-plane experiment to measure the specular and diffuse reflectance properties of a small patch

¹In these tests we polarize the light source vertically with respect to the plane of incidence so that the specular reflection does not become attenuated near the Brewster angle.




Figure 9: **Separating diffuse and specular components** can be performed by placing a linear polarizer on both the light source and the camera. (a) Normal image under point-source illumination. (b) Image of diffuse reflectance obtained by placing a vertical polarizer on the light source and a horizontal polarizer on the camera, blocking specularly reflected light. (c) Image of accentuated specular reflectance obtained by placing both polarizers vertically (half the diffusely reflected light is blocked relative to the specularly reflected light). (d) Difference of (c) and (b) yielding the specular component. The images have been scaled to appear consistent in brightness.

of skin on a person's forehead. Figure 10 shows how we adapted the light stage of Figure 2 for this purpose by placing the ϕ axis in the horizontal position and placing a vertical polarizer on the light source. We rotated the horizontal θ axis continuously while we placed a video camera aimed at our subject's vertically aligned forehead at a sampling of reflected illumination angles. The camera angles we used were $\pm(0, 22.5, 45, 60, 75, 82.5, 86.25, 89)$ degrees relative to the forehead's surface normal in order to more densely sample the illumination at grazing angles. At 89 degrees the skin area was very foreshortened so we were not able to say with certainty that the measurement we took originated only from the target area. We performed the experiment twice: once with the camera polarizer placed horizontally to block specular reflection, and once with the camera polarizer placed vertically to accentuate it. The average intensity and color of the reflected light from a 2×5 pixel area on the forehead was recorded in this set of configurations.




Figure 10: **Reflectometry Experiment** In this experiment, the diffuse and specular reflectance of an area of skin on the subject's forehead was recorded from sixty-four illumination directions for each of fifteen camera positions. Polarizers on the light and camera were used to separate the reflection components.

We noted two trends in the acquired reflectance data (Figure 11). First, the specular component becomes much stronger for large values of θ_i or θ_r , and exhibits off-specular reflection. To accommodate this behavior in our model, we use the microfacet-based framework introduced by Torrance and Sparrow [35]. This framework assumes geometric optics and models specular lobes as surface (Fresnel) reflection from microfacets having a Gaussian distribution of surface normals. Shadowing and masking effects between the microfacets are computed under the assumption that the microfacets form V-shaped grooves. Our model differs only in that we do not




Figure 11: **Reflectometry Results** The left image shows the measured diffuse (sub-surface) component of the skin patch obtained from the experiment in Fig. 10 for incident illumination angle θ_i and viewing direction θ_r . θ_r is nonuniformly spaced at angles of $\pm(0, 22.5, 45, 60, 75, 82.5, 86.25, 89)$ degrees. Invalid measurements from the light source blocking the camera’s view are set to black. The right image shows the corresponding data for accentuated specular reflectance.

assume that the microfacet normal distribution is Gaussian; since we have measurements of the specular component for dense incident directions, we simply take the microfacet normal distribution directly from the observed data. This allows the measured specular lobe to be reproduced exactly if the viewpoint is unchanged.

The second trend in the data is a desaturation of the diffuse component for large values of θ_i and θ_r . To accommodate this, we make a minor deviation from pure Lambertian behavior, allowing the saturation of the diffuse chromaticity to ramp between two values as θ_i and θ_r vary.

Representing chromaticities as unit RGB vectors, we model the diffuse chromaticity as:

$$\text{normalize}(\vec{d}_0 + f(\theta_i, \theta_r)(\vec{d}_0 - \vec{s})) \quad (6)$$

where \vec{d}_0 is a representative diffuse chromaticity, \vec{s} is the light source chromaticity, and $f(\theta_i, \theta_r)$ is given by:

$$f(\theta_i, \theta_r) = \alpha_0(\cos \theta_i \cos \theta_r) + \alpha_1(1 - \cos \theta_i \cos \theta_r) \quad (7)$$

We recover the parameters α_0 and α_1 directly from our data for each reflectance function. This correction to the diffuse chromaticity is used for the color space separation of diffuse and specular components described in Section 4.2, and also in our reflectance function resynthesis technique described in Section 4.3.

In addition to this experiment, we also performed Monte Carlo simulations of subsurface scattering similar to those in [16]. We used two scattering layers, both with strong forward scattering, and with the lower layer having significant absorption of shorter wavelengths to simulate the presence of blood in the dermis. These simulations yielded a variation in the chromaticity of the diffuse component similar to that observed in our data.

4.2 Separating Specular and Subsurface Components

We begin by separating the specular and subsurface (diffuse) components for each pixel’s reflectance function. While we could perform this step using the polarization approach of Section 4.1, this would require two passes of the lighting rig (one for diffuse only and one that includes specular) or additional cameras. Furthermore, one of the polarizers would have to rotate in a non-trivial pattern to maintain the proper relative orientations of the polarizers when ϕ is non-horizontal. Instead, we use a color space analysis technique related to [31].

For a reflectance function RGB value $R_x y(\theta, \phi)$, we can write R as a linear combination of its diffuse color \vec{d} and its specular color \vec{s} . In reality, due to noise, interreflections, and translucency, there will also be an error component \vec{e} :

$$R = \mu_d \vec{d} + \mu_s \vec{s} + \mu_e \vec{e}$$

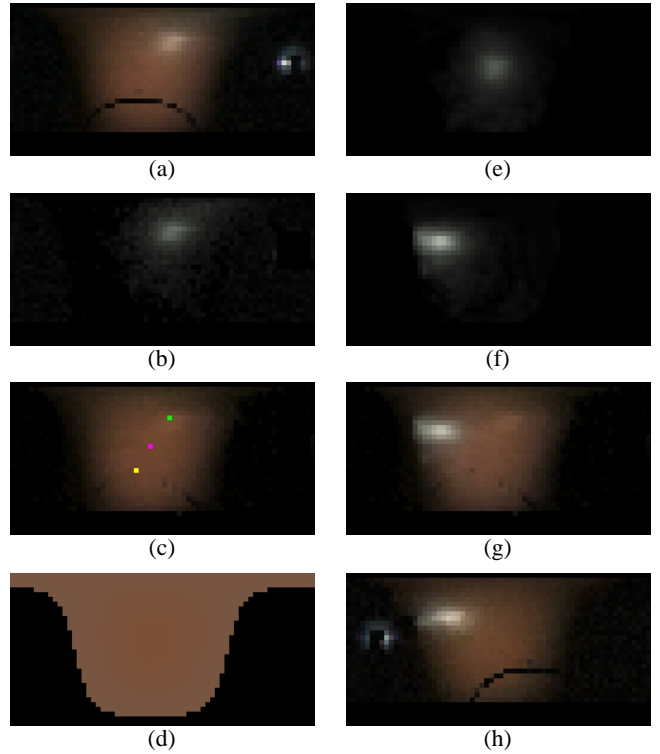


Figure 12: **Analyzing and Resynthesizing Reflectance Functions** Reflectance functions (a) can be decomposed into specular (b) and diffuse (c) components using colorspace analysis based on a model of the variation in diffuse chromaticity (d). We compute a surface normal \vec{n} based on the diffuse component (magenta dot in (c)), and a normal \vec{n}_s (coincident with \vec{n} in this case) based on the maximum (green dot) of the specular component and the known viewing direction (yellow dot). We demonstrate the resynthesis of reflectance functions for new viewpoints by resynthesizing (a), which was captured by the left camera, from the viewpoint of the right camera. We first transform the specular component to a representation independent of the original viewpoint (essentially a microfacet normal distribution) as shown in (e), then transform (e) in accordance with the new viewpoint to produce (f). The diffuse component is chrominance-shifted for the new viewpoint and added to the transformed specular component to produce the new reflectance function (g). For comparison, (h) shows the actual reflectance function (with lens flare spot and ϕ -bar shadow) from the second camera.

We choose $\vec{e} = \vec{d} \times \vec{s}$ and determine values for μ_d , μ_s , and μ_e by inverting the resulting 3×3 matrix. To form the final separation, we compute $S = \max(\mu_s, 0)\vec{s}$ and $D = R - S$ so that the sum of D and S yields the original reflectance function R .

This analysis assumes that the specular and diffuse colors are known. While we can assume that the specular component is the same color as the incident light, the diffuse color presents a more difficult problem, because it changes not only from pixel to pixel, but also within each reflectance function, as described in Section 4.1. To achieve an accurate separation, we must first estimate the diffuse chromaticity ramp.

Since we assume the diffuse chromaticity is a function f of θ_i and θ_r , we must first estimate the surface normal. For this we perform an initial rough color space separation based on a uniform diffuse chromaticity \vec{d}_0 . We derive this diffuse chromaticity by computing the median of the red-green and green-blue ratios over reflectance function values falling in a certain brightness range. We then perform a diffuse-specular separation and fit a Lambertian lobe to the diffuse component, using a coarse-to-fine direct search. This fitting yields an estimate of the surface normal.

We then find the parameters α_0 and α_1 which give the best fit to the observed chromaticities in the original unseparated reflectance function, again using a coarse-to-fine direct search. Knowing the viewpoint and the surface normal, we downweight values near the mirror angle to prevent the color ramp from being biased by strong specularities. The final separation into diffuse and specular components is computed using the fitted model of diffuse chromaticity as shown in Fig. 12.

We use the final separated diffuse component to recompute the surface normal \vec{n} , as seen in Fig. 14(b). For visualization purposes, we can also compute an estimate of the diffuse albedo $\vec{\rho}_d$ and total specular energy ρ_s , which are shown in Fig. 14(c) and (d).

4.3 Transforming Reflectance Functions to Novel Viewpoints

The process of resynthesizing a reflectance function for a novel viewpoint is illustrated in Fig. 12. The resynthesis algorithm takes the following input:

1. The diffuse reflectance function $D(\theta, \phi)$
2. The specular reflectance function $S(\theta, \phi)$
3. The surface normal \vec{n}
4. The index of refraction for surface (specular) reflection
5. The diffuse chromaticity ramp parameters α_0 and α_1
6. The original and novel view direction vectors \vec{v}_0 and \vec{v}_n

The diffuse and specular reflectance functions may optionally be transformed to a representation that does not depend on the original viewing direction, for example by transforming the functions to the form they would have if $\vec{v} = \vec{n}$. In this case, the resynthesis no longer requires the original view direction. An example of this for the specular component is shown in Fig. 12(e).

To synthesize a reflectance function from a novel viewpoint, we separately synthesize the diffuse and specular components. A sample in a specular reflectance function represents a specular response to a light source in the corresponding direction. If the view direction is known, we may consider this specular response to be a measure of the proportion of microfacets with normals oriented within some solid angle of the halfway vector between the view direction and the sample's light source direction sample. To compute a specular reflectance function from a new view direction \vec{v}_n , we compute for each light source direction \vec{l}_p the halfway vector:

$$\vec{H} = \text{normalize}(\vec{v}_n + \vec{l}_p)$$

We then find the light source direction \vec{l}_q that would have responded to microfacets near \vec{H} from the original view direction \vec{v}_0 :

$$\vec{l}_q = 2(\vec{H} \cdot \vec{v}_0)\vec{H} - \vec{v}_0$$

Letting ω_i specify a direction of incoming radiance, the Torrance-Sparrow model relates the observed radiance L to the microfacet normal distribution P as follows:

$$L_{\vec{v}} = \int \frac{PL\omega_i GF}{4 \cos \theta_r} d\omega_i \quad (8)$$

where G is a geometric attenuation factor and F is the Fresnel reflectivity. G depends on \vec{v} , \vec{l} , and \vec{n} . The expression for G is somewhat complicated, and we refer the interested reader to [35]. F is given by the Fresnel equation for unpolarized light, which can be computed from \vec{v} and \vec{l} .

Considering all quantities in (8) to be constant over the small solid angle Ω subtended by our light source, we have:

$$L_{\vec{v}} = \frac{PL_{-\vec{l}}\Omega GF}{4(\vec{v} \cdot \vec{n})}$$

Assuming the light source presents a constant $L_{-\vec{l}}\Omega$ as it moves, and recalling that the light direction \vec{l}_q is chosen to sample the same point in the microfacet normal distribution as \vec{l}_p , we can compute the new sample radiance $L_{\vec{v}_n}$ due to a light at \vec{l}_p as a function of the original radiance sample $L_{\vec{v}_0}$ due to a light at \vec{l}_q :

$$L_{\vec{v}_n} = L_{\vec{v}_0} \frac{G(\vec{v}_n, \vec{l}_p, \vec{n})F(\vec{v}_n, \vec{l}_p)(\vec{v}_0 \cdot \vec{n})}{G(\vec{v}_0, \vec{l}_q, \vec{n})F(\vec{v}_0, \vec{l}_q)(\vec{v}_n \cdot \vec{n})} \quad (9)$$

Fig. 12(f) shows a specular reflectance function synthesized using (9) for a view direction 80 degrees from the original view.

For the diffuse component we apply our diffuse chrominance ramp correction to each value in the diffuse reflectance function, first inverting the chrominance shift due to the original view direction and then applying the chrominance shift for the new view direction. The chrominance shift is computed with the recovered parameters α_0 and α_1 as in (6), using the actual sample chromaticity in place of \vec{d}_0 .

A final synthesized reflectance function consisting of the resynthesized diffuse and specular components is shown in Fig. 12(g), and is consistent with an actual reflectance function acquired from the novel viewpoint in Fig. 12(h).

4.4 Considering Shadowing and Interreflection

Since our geometry is presumed to be non-convex, we expect reflectance functions in areas not on the convex hull to exhibit global illumination effects such as shadows and interreflections. To deal with such areas, we compute a shadow map for each reflectance function. This could be done using our geometric model, but since the geometry is incomplete we instead compute the shadow map using brightness thresholding on the original reflectance function. This is demonstrated in Figure 13. We then do the analysis of Section 4.2 on the reflectance function modulated by the shadow map. This will give good results when the direct light dominates the indirect light over the non-shadowed portion of the reflectance function, a good assumption for most areas of the face.

When synthesizing a new specular reflectance function, the shadow map is used to prevent a specular lobe from appearing in shadowed directions. The converse of this effect is that when a specularity is shadowed in our original data, we are unable to recover the specular lobe. This problem could be reduced by using more cameras.

An advantage of our synthesis technique is that diffuse interreflections, and in fact all light paths terminating with a diffuse reflection, are left intact in the diffuse reflectance function and are thus reproduced without the necessity of performing the difficult steps of inverse and forward global illumination.

4.5 Creating Renderings

With the ability to resynthesize reflectance functions for new view directions, it is straightforward to render the face in arbitrary illumination from arbitrary viewpoints. We first use the technique of Section 3 to render a view of the face in the novel lighting using the modified reflectance functions. Although geometrically from the original point of view, the face is shaded as if it were viewed from the novel point of view. We then project this image onto a geometric model of the face (see Fig. 14(e)) and view the model from the novel viewpoint, yielding a rendering in which the illumination and viewpoint are consistent. In our work we use two original viewpoints, one for the left and one for the right of the face, and blend the results over the narrow region of overlap (with more cameras, view-dependent texture mapping could be used to blend between viewpoints as in [10, 29]). Renderings made with this technique are shown in Figs. 14(f),(g) and (h), and comparisons with actual photographs are shown in Fig. 15.



Figure 14: **Analyzing Reflectance and Changing the Viewpoint** (a) An original light stage image taken by the left camera. (b) Recovered surface normals n_d derived from the fitted diffuse reflectance lobe for each pixel; the RGB value for each pixel encodes the X, Y, and Z direction of each normal. (c) Estimated diffuse albedo ρ_d . Although not used by our rendering algorithm, such data could be used in a traditional rendering system. (d) Estimated specular energy ρ_s , also of potential use in a traditional rendering system. (e) Face geometry recovered using structured lighting. (f) Face rendered from a novel viewpoint under synthetic directional illumination. (g,h) Face rendered from a novel viewpoint under the two sampled lighting environments used in the second two renderings of Fig. 6.

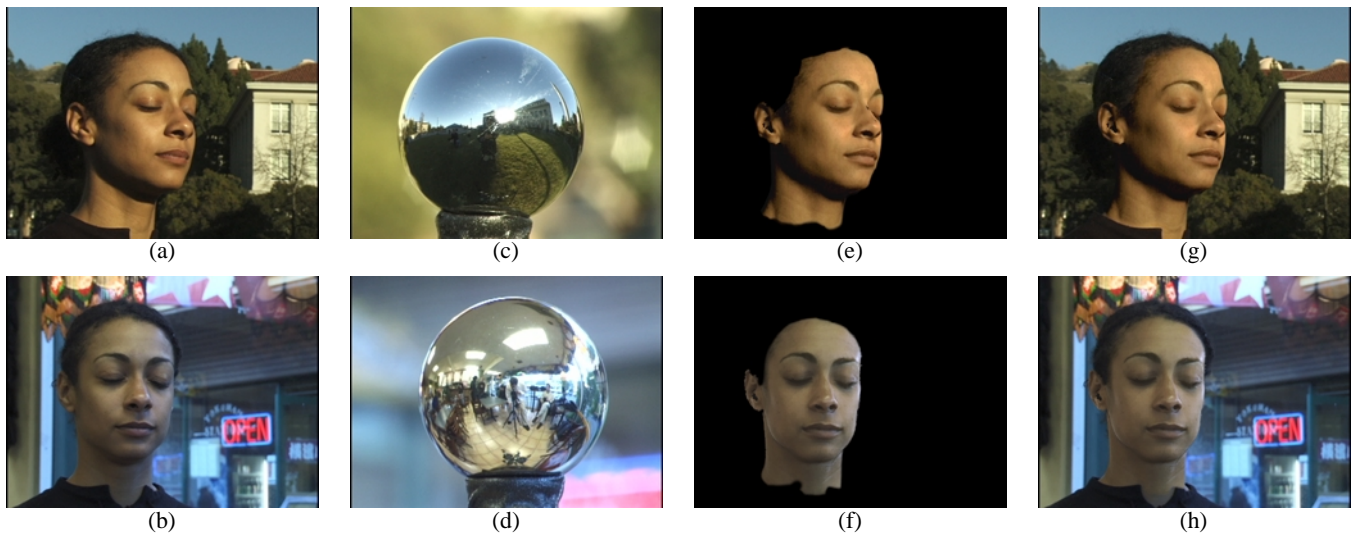


Figure 15: **Matching to Real-World Illumination** (a,b) Actual photographs of the subject in two different environments. (c,d) Images of a light probe placed in the position of the subject's head in the same environments. (e,f) Synthetic renderings of the face matched to the photographed viewpoints and illuminated by the captured lighting. (g,h) Renderings of the synthetic faces (e,f) composited over the original faces (a,b); the hair and shoulders come from the original photographs and are not produced using our techniques. The first environment is outdoors in sunlight; the second is indoors with mixed lighting coming from windows, incandescent lamps, and fluorescent ceiling fixtures.

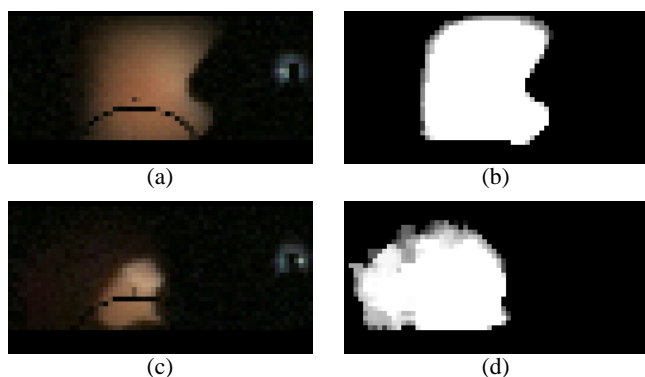


Figure 13: **Reflectance Function Shadow Maps** The reflectance function of a point near the nose (a) and the corresponding shadow map (b) computed using brightness thresholding. (c) shows a point in the ear which receives strong indirect illumination, causing the non-shadowed region in (d) to be overestimated. This causes some error in the diffuse-specular separation and the diffuse albedo to be underestimated in the ear as seen in Fig. 14(c).

5 Discussion and Future work

The work we have done suggests a number of avenues for improvements and extensions. First, we currently extrapolate reflectance functions using data from single viewpoints. Employing additional cameras to record reflectance functions for each location on the face would improve the results since less extrapolation of the data would be required. Using the polarization technique of Fig. 9 to directly record specular and subsurface reflectance functions could also improve the renderings, especially for subjects with pale skin.

A second avenue of future work is to animate our recovered facial models. For this, there already exist effective methods for animating geometrically detailed facial models such as [29], [14], and [34]. For these purposes, it will also be necessary to model and animate the eyes, hair, and inner mouth; reflectometry methods for obtaining models of such structures would need to be substantially different from our current techniques.

We would also like to investigate real-time rendering methods for our facial models. While the fixed-viewpoint re-illumination presented in Section 3 can be done interactively, synthesizing new viewpoints takes several minutes on current workstations. Some recent work has presented methods of using graphics hardware to render complex reflectance properties [18]; we would like to investigate employing such methods to create renderings at interactive rates. We also note that the storage required for a reflectance field could be substantially reduced by compressing the source data both in (u, v) space as well as (θ, ϕ) space to exploit similarities amongst neighboring reflectance functions.

Real skin has temporally varying reflectance properties depending on temperature, humidity, mood, health, and age. The surface blood content can change significantly as the face contorts and contracts, which alters its coloration. Future work could characterize these effects and integrate them into a facial animation system; part the acquisition process could be to capture the reflectance field of a person in a variety different expressions.

Lastly, the data capture techniques could be improved in a number of ways. High-definition television cameras would acquire nearly eight times as many pixels of the face, allowing the pixel size to be small enough to detect illumination variations from individual skin pores, which would increase the skin-like quality of the renderings. One could also pursue faster capture by using high-speed video cameras running at 250 or 1000 frames per second, allowing full reflectance capture in just a few seconds and perhaps, with more advanced techniques, in real time.

6 Conclusion

In this paper we have presented a practical technique for acquiring the reflectance field of a human face using standard video equipment and a relatively simple lighting apparatus. The method allows the face to be rendered under arbitrary illumination conditions including image-based illumination. The general technique of modeling facial reflectance from dense illumination directions, sparse viewpoints, and recovered geometry suggests several areas for future work, such as fitting to more general reflectance models and combining this work with facial animation techniques. It is our hope that the work we have presented in this paper will help encourage continued investigations into realistic facial rendering.


Acknowledgements

We would like to thank Shawn Brixey and UC Berkeley's Digital Media/New Genre program for use of their laboratory space, as well as Bill Buxton and Alias Wavefront for use of the Maya modeling software, and Larry Rowe, Jessica Vallot (seen in Fig. 14), Patrick Wilson, Melanie Levine, Eric Paulos, Christine Waggoner, Holly Cim, Eliza Ra, Bryan Musson, David Altenau, Marc Levoy, Maryann Simmons, Henrik Wann Jensen, Don Greenberg, Pat Hanrahan, Chris Bregler, Michael Naimark, Steve Marschner, Kevin Binkert, and the Berkeley Millennium Project for helping make this work possible. We would also like to acknowledge the Cornell's 1999 Workshop on Rendering, Perception, and Measurement for helping encourage this line of research. Thanks also to the anonymous reviewers for their insightful suggestions for this work. This work was sponsored by grants from Interactive Pictures Corporation, the Digital Media Innovation Program, and ONR/BMDO 3DDI MURI grant FDN00014-96-1-1200.

References

- [1] ADELSON, E. H., AND BERGEN, J. R. *Computational Models of Visual Processing*. MIT Press, Cambridge, Mass., 1991, ch. 1. The Plenoptic Function and the Elements of Early Vision.
- [2] BEIER, T., AND NEELY, S. Feature-based image metamorphosis. *Computer Graphics (Proceedings of SIGGRAPH 92)* 26, 2 (July 1992), 35–42.
- [3] BLANZ, V., AND VETTER, T. A morphable model for the synthesis of 3d faces. *Proceedings of SIGGRAPH 99* (August 1999), 187–194.
- [4] BREGLER, C., COVELL, M., AND SLANEY, M. Video rewrite: Driving visual speech with audio. *Proceedings of SIGGRAPH 97* (August 1997), 353–360.
- [5] BUSBRIDGE, I. W. *The Mathematics of Radiative Transfer*. Cambridge University Press, Bristol, UK, 1960.
- [6] COOK, R. L., AND TORRANCE, K. E. A reflectance model for computer graphics. *Computer Graphics (Proceedings of SIGGRAPH 81)* 15, 3 (August 1981), 307–316.
- [7] DANA, K. J., GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. Reflectance and texture of real-world surfaces. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.* (1997), pp. 151–157.
- [8] DEBEVEC, P. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH 98* (July 1998).
- [9] DEBEVEC, P. E., AND MALIK, J. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH 97* (August 1997), pp. 369–378.
- [10] DEBEVEC, P. E., YU, Y., AND BORSHUKOV, G. D. Efficient view-dependent image-based rendering with projective texture-mapping. In *9th Eurographics workshop on Rendering* (June 1998), pp. 105–116.
- [11] FUA, P., AND MICCIO, C. From regular images to animated heads: A least squares approach. In *ECCV98* (1998).
- [12] GERSHUN, A. Svetovoe Pole (the Light Field, in English). *Journal of Mathematics and Physics XVIII* (1939), 51–151.
- [13] GORTLER, S. J., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M. F. The Lumigraph. In *SIGGRAPH 96* (1996), pp. 43–54.
- [14] GUENTER, B., GRIMM, C., WOOD, D., MALVAR, H., AND PIGHIN, F. Making faces. *Proceedings of SIGGRAPH 98* (July 1998), 55–66.
- [15] HABERLI, P. Synthetic lighting for photography. Available at <http://www.sgi.com/grafica/synth/index.html>, January 1992.

- [16] HANRAHAN, P., AND KRUEGER, W. Reflection from layered surfaces due to subsurface scattering. *Proceedings of SIGGRAPH 93* (August 1993), 165–174.
- [17] KARNER, K. F., MAYER, H., AND GERVAUTZ, M. An image based measurement system for anisotropic reflection. In *EUROGRAPHICS Annual Conference Proceedings* (1996).
- [18] KAUTZ, J., AND MCCOOL, M. D. Interactive rendering with arbitrary BRDFs using separable approximations. *Eurographics Rendering Workshop 1999* (June 1999).
- [19] LAFORTUNE, E. P. F., FOO, S.-C., TORRANCE, K. E., AND GREENBERG, D. P. Non-linear approximation of reflectance functions. *Proceedings of SIGGRAPH 97* (August 1997), 117–126.
- [20] LEE, Y., TERZOPOULOS, D., AND WATERS, K. Realistic modeling for facial animation. *Proceedings of SIGGRAPH 95* (August 1995), 55–62.
- [21] LEVOY, M., AND HANRAHAN, P. Light field rendering. In *SIGGRAPH 96* (1996), pp. 31–42.
- [22] MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., TORRANCE, K. E., AND GREENBERG, D. P. Image-based BRDF measurement including human skin. *Eurographics Rendering Workshop 1999* (June 1999).
- [23] MILLER, G. S. P., RUBIN, S., AND PONCELEON, D. Lazy decomposition of surface light fields for precomputed global illumination. *Eurographics Rendering Workshop 1998* (June 1998), 281–292.
- [24] NAYAR, S., FANG, X., AND BOULT, T. Separation of reflection components using color and polarization. *IJCV 21*, 3 (February 1997), 163–186.
- [25] NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. Geometric considerations and nomenclature for reflectance.
- [26] NIMEROFF, J. S., SIMONCELLI, E., AND DORSEY, J. Efficient re-rendering of naturally illuminated environments. *Fifth Eurographics Workshop on Rendering* (June 1994), 359–373.
- [27] OREN, M., AND NAYAR, S. K. Generalization of Lambert’s reflectance model. *Proceedings of SIGGRAPH 94* (July 1994), 239–246.
- [28] PARKE, F. I. Computer generated animation of faces. *Proc. ACM annual conf.* (August 1972).
- [29] PIGHIN, F., HECKER, J., LISCHINSKI, D., SZELISKI, R., AND SALESIN, D. H. Synthesizing realistic facial expressions from photographs. *Proceedings of SIGGRAPH 98* (July 1998), 75–84.
- [30] SAGAR, M. A., BULLIVANT, D., MALLINSON, G. D., HUNTER, P. J., AND HUNTER, I. W. A virtual environment and model of the eye for surgical simulation. *Proceedings of SIGGRAPH 94* (July 1994), 205–213.
- [31] SATO, Y., AND IKEUCHI, K. Temporal-color space analysis of reflection. *JOSA-A 11*, 11 (November 1994), 2990–3002.
- [32] SATO, Y., WHEELER, M. D., AND IKEUCHI, K. Object shape and reflectance modeling from observation. In *SIGGRAPH 97* (1997), pp. 379–387.
- [33] SMITH, B., AND ROWE, L. Compressed domain processing of JPEG-encoded images. *Real-Time Imaging 2*, 2 (1996), 3–17.
- [34] TERZOPOULOS, D., AND WATERS, K. Physically-based facial modelling, analysis, and animation. *Journal of Visualization and Computer Animation 1*, 2 (August 1990), 73–80.
- [35] TORRANCE, K. E., AND SPARROW, E. M. Theory for off-specular reflection from roughened surfaces. *Journal of Optical Society of America 57*, 9 (1967).
- [36] VAN GEMERT, M. F. C., JACQUES, S. L., STERENBERG, H. J. C. M., AND STAR, W. M. Skin optics. *IEEE Transactions on Biomedical Engineering 36*, 12 (December 1989), 1146–1154.
- [37] WARD, G. J. Measuring and modeling anisotropic reflection. In *SIGGRAPH 92* (July 1992), pp. 265–272.
- [38] WILLIAMS, L. Performance-driven facial animation. *Computer Graphics (Proceedings of SIGGRAPH 90)* 24, 4 (August 1990), 235–242.
- [39] WONG, T.-T., HENG, P.-A., OR, S.-H., AND NG, W.-Y. Image-based rendering with controllable illumination. *Eurographics Rendering Workshop 1997* (June 1997), 13–22.
- [40] WU, Y., THALMANN, N. M., AND THALMANN, D. A dynamic wrinkle model in facial animation and skin aging. *Journal of Visualization and Computer Animation 6*, 4 (October 1995), 195–206.
- [41] YU, Y., DEBEVEC, P., MALIK, J., AND HAWKINS, T. Inverse global illumination: Recovering reflectance models of real scenes from photographs. *Proceedings of SIGGRAPH 99* (August 1999), 215–224.
- [42] ZONGKER, D. E., WERNER, D. M., CURLESS, B., AND SALESIN, D. H. Environment matting and compositing. *Proceedings of SIGGRAPH 99* (August 1999), 205–214.



Appendix: Combining with Environment Matting

The light stage can be used to relight objects as well as faces. In this experiment we created a scene with diffuse, shiny, refractive, and transmissive objects seen in (a). Because of the sharp specularities, we recorded the scene with a finer angular resolution of 128×64 directions of θ and ϕ and in high dynamic range [9] using five passes of the light stage at different exposure settings. Renderings of the scene in two environments are shown in (c,d). Because high dynamic range imagery was used, the direct appearance of the light source was captured properly, which allows the renderings to reproduce a low-resolution version of the lighting environment in the background. To replace this with a high resolution version of the environment, we captured an environment matte [42] of the scene (b) and computed the contribution of the reflected, refracted, and transmitted light from the background (e,f). We then summed all but the contribution from the background lighting directions to produce (g,h) and added in the light from the environment matte (e,f) to produce a complete rendering of the scene and background (i,j).