

Practical Modeling and Acquisition of Layered Facial Reflectance

Abhijeet Ghosh

Tim Hawkins

Pieter Peers

Sune Frederiksen

Paul Debevec

University of Southern California
Institute for Creative Technologies




Figure 1: Layers of skin reflectance which are modeled by our technique and used to render faces for novel viewpoints and lighting.

Abstract

We present a practical method for modeling layered facial reflectance consisting of specular reflectance, single scattering, and shallow and deep subsurface scattering. We estimate parameters of appropriate reflectance models for each of these layers from just 20 photographs recorded in a few seconds from a single viewpoint. We extract spatially-varying specular reflectance and single-scattering parameters from polarization-difference images under spherical and point source illumination. Next, we employ direct-indirect separation to decompose the remaining multiple scattering observed under cross-polarization into shallow and deep scattering components to model the light transport through multiple layers of skin. Finally, we match appropriate diffusion models to the extracted shallow and deep scattering components for different regions on the face. We validate our technique by comparing renderings of subjects to reference photographs recorded from novel viewpoints and under novel illumination conditions.

1 Introduction

Realistically reproducing the appearance of the human face from novel viewpoints and under novel complex illumination remains a challenging problem in computer graphics due to the complexity of human facial reflectance and our keen eye for its subtleties. The appearance of the face under given lighting conditions is the result of complex light interactions with a complex, inhomogeneous material [Igarashi et al. 2007; Tuchin 2007]. Realistic facial reflectance requires a model consisting of spatially-varying specular and diffuse reflectance which reproduces the effects of light scatter-

ing through multiple layers of translucent tissue [Donner and Jensen 2005; Donner and Jensen 2006].

Advances in the field of 3D scanning and reflectance measurement have enabled significant strides in the rendering of realistic faces. However, while existing methods for accurately modeling the appearance of human skin are able to achieve impressive results, it is not clear how to practically acquire the necessary parameters for these models to accurately reproduce the facial appearance of live subjects. Existing acquisition techniques are either very data intensive, or they extrapolate parameters from a small exemplar skin patch to cover the whole face, or they make simplifications to the skin reflectance model.

This paper aims to develop a practical appearance model which is in addition easy to incorporate in existing rendering systems. The detail in the facial appearance model should be such that full-screen close-ups can be faithfully reproduced. Additionally, working with live subjects requires fast acquisition to avoid registration problems, temporal changes in the appearance (e.g., due to sweat or blood flow), and to enable capture of facial appearance of natural expressions which are difficult to hold for more than a few seconds.

To achieve these goals, we model facial skin reflectance as a combination of different layers: specular reflectance, single scattering, and shallow and deep multiple scattering (Fig. 1). A suitable reflectance or scattering model is selected for each layer, and parameters are obtained using a single high-resolution still camera to capture a small set of 20 photographs under environmental and projected lighting conditions. For each reflectance component, we estimate or infer high-frequency details such as albedo and normals *per pixel* based on the environmental illumination patterns, while modeling lower-frequency BRDF and scattering behavior *per region* based on the projected patterns. This allows for fast acquisition and straightforward processing, while achieving a high level of realism in the resulting models. Although previous research has captured and modeled some of these individual components, no existing system has acquired and modeled all of these reflectance components together of a live subject. We demonstrate the effectiveness of our technique with both qualitative visual comparisons as well as quantitative validation of extracted model parameters against those available in the literature.

In summary, the principal contributions of this work are:

e-mail: {ghosh, debevec}@ict.usc.edu

1. A practical method for estimating specular reflectance and explicit modeling of single scattering of a subject from a few lighting conditions.
2. A practical method for estimating scattering parameters for a *data-driven* multi-layer diffusion model of a subject from a small set of photographs.
3. An integrated system for capturing detailed facial reflectance at high resolution in just 20 photographs, recorded in a few seconds.

2 Related Work

Our work builds on a range of related work in acquiring and modeling the shape and reflectance of complex objects including human faces. While a wide body of work has measured and modeled diffuse and specular reflectance properties of objects, we focus our discussion on measuring and modeling the reflectance properties of human faces due to their particular combination of specular and multi-layer subsurface reflectance and the special requirements of measuring these characteristics for real faces.

Modeling Skin with BRDFs Marschner et al. [1999] use an image-based technique to obtain the aggregate BRDF of a human forehead from photographs taken under multiple lighting directions. Marschner et al. [2000] create facial renderings by modulating the diffuse component of such a BRDF with the diffuse albedo map estimated from multiple cross-polarized photographs of the face. Georghiades et al. [1999] built models of facial shape and reflectance from a small number of unknown point-source lighting directions using an enhanced version of *photometric stereo* [Woodham 1978]. These works assume a Lambertian reflection model, and ignore specular reflection. To account for specular reflections, Georghiades [2003] extend [Georghiadis et al. 1999] to estimate a single Torrance-Sparrow specular lobe across the entire face. However, they note that the lack of spatially-varying specular behavior limits the technique’s ability to model the observed data, which limits the realism of the renderings. Reflectance Sharing [Zickler et al. 2006] trades spatial resolution for angular reflectance information to estimate spatially-varying BRDFs from a small number of photographs of a face. All of these methods model skin reflectance solely using BRDF models, omitting the subsurface scattering behavior of skin. As in [Zickler et al. 2006], we estimate spatially-varying specular reflectance parameters, but augment this with high fidelity normal estimates and also include single scattering and subsurface scattering models.

Modeling Subsurface Scattering Modeling subsurface scattering behavior is important to create the soft, semi-translucent appearance of skin. Without subsurface scattering, renderings of skin look too harsh. Hanrahan and Krueger [1993] use a Monte-Carlo simulation to develop local reflectance models for the single and multiple scattering components of human skin and other layered tissues. Jensen et al. [2001] introduced a practical dipole model to simulate scattering behavior, and show how to infer parameters from the observation of the spread of a small white beam of light incident on a patch of skin. Donner and Jensen [2005] extend the dipole model to simulate transmission through and reflection from multiple layers, yielding a more accurate skin rendering model. More recently, Donner and Jensen [2006] presented an easily parameterized, spectrally-accurate version of the multi-layer model. These works mostly focus on practically modeling subsurface scattering for rendering. However, they do not deal with obtaining spatially-varying parameters for the dipole model or the multi-layer models. Specialized techniques, such as [Goesele et al. 2004; Tong et al. 2005; Peers et al. 2006; Wang et al. 2008], can acquire and model a

wide variety of subsurface scattering materials, including skin, but are limited to planar samples only, or have acquisition times that are impractically long for human subjects. The presented method is specifically designed to minimize the number of photographs (and thus acquisition time) from which multi-layer scattering parameters can be estimated. In concurrent work, Donner et al. [2008] estimate multi-layer scattering parameters driven by chromophore concentrations. These concentrations are measured using photographs of a small patch of skin at nine different wavelengths.

Realistic Face Scanning Debevec et al. [2000] use a dense sphere of incident lighting directions to record specular and subsurface *reflectance functions* of a face at relatively high angular resolution. However, the model is data-intensive in both acquisition and storage. Additionally, inclusion in existing rendering systems requires significant effort. Fuchs et al. [2005] use a smaller number of photographs and lighting directions, at the cost of sacrificing continuously-varying specular reflectance. Tariq et al. [2006] use a set of approximately forty phase-shifted video projector lines to estimate per-pixel scattering parameters for faces. However, their acquisition times were as long as a minute, and they did not model the specular reflectance of skin. Weyrich et al. [2006] use a dense sphere of lighting directions and sixteen cameras to model the per-pixel specular BRDF and diffuse albedo of faces. In addition, they use a custom subsurface scattering measurement probe to obtain scattering parameters for skin. While the obtained appearance model yields impressive results, it still requires a minute to complete a full capture consisting of thousands of images. In contrast, our method estimates a more expressive facial reflectance model from just 20 photographs captured from a single viewpoint. As a result our method is less data intensive, can be implemented in high resolution at a relatively low cost, and avoids the task of building reflectance datasets from images from multiple viewpoints.

3 Acquisition

Before discussing our skin reflectance model in Sec. 4, we will first overview our measurement setup, calibration process, and 3D scanning system.

Setup Our setup consists of an LED sphere with approximately 150 individually controllable lights. Each light is covered with a linear polarizer in the pattern of [Ma et al. 2007]. Additionally, a vertically polarized LCD video projector is aimed towards the center of the sphere. A stereo pair of radiometrically calibrated 10-Megapixel Canon 1D Mark III digital SLR cameras are placed on opposite sides of the projector. The right camera, used only for geometry measurement, is horizontally polarized while the left camera switches between horizontal and vertical polarization through a mechanical actuator.

Calibration The purpose of using polarized illumination is to tune out specular reflections on the subject. For this we need to align the linear polarizers on the sphere such that specular highlights are invisible through a horizontally polarized camera. This can be easily achieved by placing a dielectric spherical reflector (i.e., plastic ball) in the middle of the LED sphere, and rotating each polarizer until no highlight is visible through the left camera.

A challenge for reflectance measurement is that we have two different illuminants in our setup: the LCD projector, and the white LEDs. To compensate for the differences in emitted spectra, we measure the responses of 24 ColorChecker squares and 10 corresponding skin patches on different subjects. Using SVD, we compute a 3×3 color matrix that transforms the observed photographs to a common illuminant color space. The skin colors did not match well when using only the ColorChecker samples. Including the skin




Figure 2: Modeled skin reflectance components.

samples provides a much closer match between the different color spaces. A similar color calibration is performed for additional illuminants used to generate the reference images in the results in this paper.

In addition, we subtract a reference black level photograph of the subject from every recorded photograph under projected illumination to compensate for the black level illumination from the projector.

Geometry Acquisition Accurate 3D geometry of a subject is required to faithfully model the subject’s skin reflectance. In this paper we use the method of Ma et al. [2007] to obtain geometry from stereo correspondence and *specular* normals. For this we capture four projected color fringe patterns for 3D stereo reconstruction, and eight photographs of the subject under four different gradient illumination conditions and two polarization directions. However, alternative methods that can measure detailed facial geometry with accurate surface normals could also be used for this purpose.

In addition to these twelve photographs, eight more photographs are recorded to infer the appropriate reflectance and scattering models (Sec. 4). These eight photographs are:

- A black level reference for the video projector (1 image).
- A cross-polarized grid of black dots projected from the front to measure subsurface scattering parameters (1 image).
- A pair of cross- and parallel-polarized front-lit (i.e., full-on projector pattern) images to model specular and diffuse reflectance (2 images).
- Four phase-shifted stripe patterns to separate shallow and deep scattering (4 images).

Recording these 20 photographs takes just 5 seconds with our current setup, with the major limiting factor being the frame rate of the digital SLR cameras. Using faster high resolution cameras could reduce acquisition times to under a second.

4 Skin Reflectance Model

We approximate skin reflectance as a combination of four phenomena: specular reflection, single scattering, shallow multiple scattering, and deep multiple scattering (Fig. 2). We design illumination conditions to measure each of these components as directly and independently as possible. We fit our image-based measurements to different reflectance models, each of which is chosen according to the type of phenomena being modeled. We later create renderings by summing the contributions of these four components, modulating the light received by the scattering components by appropriate transmittance terms. In order to model these reflectance effects from a limited set of photographs, we model some aspects of reflectance per pixel (e.g., albedos and surface normals), some aspects per region (e.g., specular roughness and scattering parameters), and some aspects for the entire face (e.g., the angular dependence of the scattering components).



Figure 3: Measured Reflectance Components (a) Polarization difference image under spherical illumination, used for estimating specular albedo. (b) Cross-polarized image under spherical illumination, used to measure total scattered albedo. (c) Polarization difference image under directional illumination, used for estimating the specular lobe shape per region. The image also includes some polarization preserving non-specular backscattering (which we model as mostly single-scattering), which can be seen to pick up color from the melanin in the epidermis. (d) Cross-polarized image under directional illumination, showing multiple scattering. (e) “direct” component of (d), showing shallow scattering. (f) “indirect” component of (d), showing deeply scattered light. Note that $(d) = (e) + (f)$ and that $(c) + (d)$ produces a typical front-lit photograph.

The remainder of this section is organized as follows. Sec. 4.1 introduces the specular and single scattering model. We show how polarization can be used to isolate these phenomena from multiple subsurface scattering, and detail which data is required to fit appropriate reflectance models. Sec. 4.2 further separates the multiple subsurface scattering into *deep*, and *shallow* scattering.

4.1 Specular Reflection and Single Scattering

We leverage the polarization properties of skin to extract specular reflectance and single scattering. Both phenomena generally maintain the polarization of light [Morgan and Ridgway 2000]. Multiple scattering phenomena, on the other hand, generally depolarizes light [Tuchin 2007]. We therefore acquire data under polarized spherical and front-lit illumination, and record parallel- and cross-polarized images of each lighting condition. The cross-polarized images only include depolarized reflected light (i.e., due to multiple scattering events), whereas the parallel-polarized images contain both polarized as well as depolarized reflected light. Computing the difference between the corresponding parallel- and cross-polarized images yields an image exhibiting only polarized reflected light, i.e., specular reflected and *some* non-specular reflected light which maintains polarization. The latter component is dominated by single scattering, because the probability of de-polarization of light increases exponentially with each additional scattering event [Morgan and Ridgway 2000]. We therefore treat any observed polarization preserving non-specular reflection as the result of single scattering events. The polarization-difference images in Figs. 3 (a) and (c) show specular reflections and single scattering on a face under spherical and directional illumination respectively. Figs. 3 (b) and




Figure 4: Estimating Per-region Specular BRDFs (a) Face segmentation into regions. (b) A front-lit rendering of the spatially-varying specular reflectance. (c) A front-lit rendering with both the spatially-varying specular reflectance and modeled single scattering. (d) Front-lit polarization difference image with specular reflection and single scattering. **Graph:** Extracted specular distributions per region.

(d) show the effects of multiple scattered illumination under the same lighting conditions.

We will now discuss the appropriate reflectance models, and fitting procedures we use for specular reflectance and single scattering.

Specular Reflection As noted in [Debevec et al. 2000; Georghides 2003], the spatially varying specular behavior of skin is important for reproducing facial appearance realistically. In order to minimize the number of measurements, a per-pixel estimation of the specular lobe and albedo is not practical. Therefore, we will estimate specular albedo per-pixel and (as in [Weyrich et al. 2006]) extract separate specular roughness distributions for different regions of the face corresponding to the forehead, eyelids, nose, cheekbone, lips, and lower cheek regions (Fig. 4(a)).

We model the specular roughness distributions over a region using a microfacet BRDF model. In previous work, dense measurements have been fitted to analytic BRDF models such as the Torrance-Sparrow or Blinn-Phong model [Debevec et al. 2000; Fuchs et al. 2005; Weyrich et al. 2006]. However, to keep the number of measurement small, we only use backscattering measurements from a *single* photograph under point source illumination (i.e., a full-on projector pattern) to estimate per-region microfacet distributions for the Torrance-Sparrow [1967] model:

$$\rho(\hat{k}_1, \hat{k}_2) = \frac{c \cdot p(\hat{h}) F(r_o, \hat{k}_2 \cdot \hat{h}) G}{(\hat{k}_1 \cdot \hat{n})(\hat{k}_2 \cdot \hat{n})}, \quad (1)$$

where \hat{k}_1 is the incident light direction, \hat{k}_2 is the viewing direction, c is a normalization constant (corresponding to specular intensity in our case), $p(\hat{h})$ is the normalized distribution, $F(r_o, \hat{k} \cdot \hat{h})$ is the Fresnel reflectance term based on Snell’s laws of reflection, and G is the geometric shadowing and masking term based on V-shaped grooves.

Similar to [Debevec et al. 2000], we replace the Gaussian distribution in the original Torrance-Sparrow model with a data-driven distribution term derived directly from the observed backscattering data. We extract this data-driven distribution in a manner similar to the procedure discussed in [Ashikhmin 2006], where the effects of the Fresnel term and the geometric term are assumed to be minimal in the backscattering direction, and the distribution-based BRDF model simplifies to a function that is proportional to the distribution $p(\hat{h})$:

$$\rho(\hat{k}, \hat{k}) = \frac{c \cdot r_o \cdot p(\hat{h})}{2(\hat{k} \cdot \hat{n}) - (\hat{k} \cdot \hat{n})^2}. \quad (2)$$

This distribution can then be directly tabulated, without requiring any numerical optimizations, from the observed data using Equation 2.

We use the polarization-difference image of the face lit from the front to observe the backscattered specular reflection (in addition to single scattering) (Fig. 4(d)). We trade spatial resolution across the face for angular resolution in order to densely sample a distribution $p(\hat{h})$ per region from a single photograph. To eliminate the effects of single scattering, we isolate the regions where specular reflection dominates by considering only pixels above a certain brightness threshold and whose surface normals lie within a cone of 45° from the viewing direction for constructing the specular distributions. The argument for a 45° threshold is that the specular lobes we have observed for faces are much sharper than 45° , and single scattering is predominately directed forward in skin. The observed single scattering is therefore dominated by the specular reflection, and hence can be directly used to estimate the specular lobes. Similar angular and intensity separation methods are commonly used in the tissue optics literature [Morgan and Ridgway 2000].

The specular intensity c is unknown at this point, and is required to extract the specular distributions. We therefore bootstrap the estimation process by (initially) assuming a per-region constant specular intensity. Next, we tabulate the observed reflectance values against the halfway vectors corresponding to the normal direction. The graph in Fig. 4 plots distributions obtained for different facial regions. As expected, the measured specular lobe shape differs for the different regions.

Finally, we need to infer a per-pixel specular intensity c . We observe that the polarization-difference image under constant spherical illumination (e.g., Fig. 3(a)) is dominated by the specular reflection for all pixels, unlike front-lit illuminated pixels where single scattering can dominate for pixels facing away from the view (and light) direction. This polarization-difference image under spherical illumination is taken to encode the specular intensity at each pixel modulated by view-dependent Fresnel reflectance. Note that this illumination condition is also one of the gradient patterns used for computing the surface normals [Ma et al. 2007], and thus no additional photograph needs to be recorded. From this we can estimate the specular intensity using the previously extracted distributions, and factor out Fresnel reflectance effects, assuming a constant index of refraction of 1.38 for skin as in [Donner and Jensen 2005]. Formally, let the observed intensity in the polarization-difference image under constant hemispherical illumination for a given pixel be c' , for a fixed viewing direction \hat{k}_2 , then the following holds: $c' = \int \rho(\hat{k}_1, \hat{k}_2)(\hat{k}_1 \cdot \hat{n}) d\omega$. By dividing c' by the (numerically) hemispherically integrated BRDF (assuming $c = 1.0$, and including Fresnel reflectance) the best-fit specular intensity c is obtained. To further refine the estimation of the specular distribution $p(\hat{h})$ and specular intensity c , we could iteratively alternate between estimating $p(\hat{h})$ and c . However, we found that a single pass yields accurate results.

A rendering of the obtained specular component under directional illumination from the front can be seen in Fig. 4(b). This rendering closely follows the observed specular reflectance in Fig. 4(d). Note that the differences between both are due to the single scattering included in the polarization-difference photograph.

Single Scattering We model the remaining single scattering component with the 1st order single scattering BRDF model of Hanrahan and Krueger [1993]:

$$\rho_{\text{single scatter}}(\hat{k}'_1, \hat{k}'_2) = \alpha \cdot T_{dt} \cdot p(\cos \theta) \frac{1}{\hat{n} \cdot \hat{k}'_1 + \hat{n} \cdot \hat{k}'_2}, \quad (3)$$

where α is the scattering albedo, T_{dt} is the transmittance term, and p is the Henyey-Greenstein scattering phase function given as $p(\cos \theta) = \frac{1-g^2}{4\pi(1-g+2g\cos\theta)^{3/2}}$, with θ being the angle between incident \hat{k}'_1 and scattered \hat{k}'_2 directions, and g the mean cosine of the scattering angle.

Similar to the specular lobe fits, the Henyey-Greenstein function is fitted to match the observed backscattering in the polarization-difference image under directional illumination. We assume that the observed single scattering is mainly due to the top layer of skin, and set the index of refraction of this layer to 1.38 as in [Donner and Jensen 2005]. Furthermore, we use the observed polarization-difference image under uniform spherical illumination minus the specular intensity c as the albedo α for the single scattering fit. Our choice of employing the polarization-difference image as a basis for the single scattering albedo is more data-driven than strictly physically-based, given that we model any polarization preserving non-specular backscatter as single scattering and we do not observe any texture variation in the observed single scattering.

Given that the Torrance-Sparrow BRDF models a rough specular surface, we replace the Fresnel equations for transmission in a smooth surface with diffuse transmission T_{dt} due to the rough specular surface [Ashikhmin et al. 2000; Donner and Jensen 2005]: $T_{dt} = \rho_{dt}(x, \omega_i)\rho_{dt}(x, \omega_o)$, where:

$$\rho_{dt}(x, \omega_o) = 1.0 - \int \rho_{\text{specular}}(x, \hat{k}_1, \hat{k}_2)(\hat{n}_s \cdot \hat{k}_1) d\omega. \quad (4)$$

As with the specular reflectance, we leverage the polarization-difference image under constant hemispherical illumination which encodes this per-pixel integral. To facilitate computations, we build a look-up table for average diffuse transmittance values across the face. This reduces fitting the observed single scattering to the above BRDF model to a simple search for the best channel-wise g values that minimize the RMS error of the fit to the observed data. Given the slowly varying nature of the data, we found that using a single set of channel-wise g values across the entire face is sufficient. A front-lit rendering of the combined single scattering and specular component is shown in 4(c), which closely matches the reference photograph in Fig. 4(d).

4.2 Modeling Multiple Scattering Components

Multiple subsurface scattering of light in skin is an important phenomena that contributes significantly to its soft appearance [Jensen et al. 2001]. Without subsurface scattering, renderings of skin look too harsh. However, modeling skin as a single homogeneous scattering media results in a too soft or “waxy” appearance. Modeling skin as a multi-layer subsurface scattering medium [Hanrahan and Krueger 1993; Donner and Jensen 2005] represents the structure of skin much better, and yields more realistic results.



Figure 5: Separated Multiple Scattering Layers (a) Separated shallow scattering (direct) component. (b) Separated deep scattering (indirect) component. Deep scattering exhibits more saturated coloring and a greater amount of light diffusion than the shallow scattering component.

A possible physically-based model for the appearance of skin is to represent it as a two layer subsurface scattering medium (Fig. 2). In such a case, the top layer corresponds to the epidermal layer, which is a scattering layer with a thickness of approximately 0.5mm , with a color that is mostly determined by the melanin content. In contrast, the bottom layer corresponds to the dermis, which is a (relatively) thick layer with a reddish hue due to blood. However, measuring the scattering properties of these two layers exactly is a difficult problem. We will therefore use an approximate *data-driven* two-layer model, where the interface between both layers corresponds only approximately to the interface between the different skin layers. We denote the two scattering layers as *shallow* and *deep* to emphasize that we do not precisely associate them with specific anatomical skin layers.

To measure the per-pixel ratio between both layers, we observe that the shallow layer scatters light much less than the deep layer. Recently, Nayar et al. [2006] presented a method to separate a photograph into direct and indirect components using high frequency illumination patterns. In scattering materials, the frequency of the illumination patterns determines which part of scattered light is classified as direct, and which part as indirect. We make the observation that selecting the frequency of the patterns to be on the order of the thickness of the epidermis, separates the reflectance into an image containing deep scattering only, and an image containing only shallow scattering. We use *four* phase-shifted high-frequency patterns of 1.2mm -wide stripes from a video projector. Computing a per-pixel *max* and *min* over the four images yields the the direct/shallow scattering image ($\text{max} - \text{min}$), and indirect/deep scattering image ($2 \times \text{min}$). Furthermore, we employ cross-polarization to eliminate specular reflections and single scattering. Separated components are shown in Figs. 3 (e) and (f), and Fig. 5. The shallow scattering shows relatively little color saturation relative to the deep scattering, and the deep scattering exhibits less distinct texture detail. This corresponds to the thesis that the direct component approximately corresponds to the shallow scattering of light in the epidermis while the indirect component approximately corresponds to light which has scattered more deeply within the dermis.

The proposed two layer subsurface scattering model *sums* the contributions of the shallow and deep scattering layers, due to the way the deep and shallow scattering layers are separated. This differs from [Donner and Jensen 2005] in which the individual layers’ contributions are convolved according to the Kubelka-Munk theory. In this respect, our two-layer model is more data-driven in nature than physically-based.

Formally, the multiple subsurface scattering of light in skin can be




Figure 6: Measurement of Per-region Scattering Parameters (a) Dot pattern used to observe the scattering profiles (depicted in(d)). (b) Subject under full illumination. (c) Zero-crossings computed from subtracting (a) from (b). (e) Fitted deep scattering model versus the observed scattering profile for two different regions. Note that the poor fit close to the peak is because the observed scattering profile also contain shallow scattering effects. However, further from the peak, where deep scattering dominates, a good fit is obtained.

represented as:

$$L_{multiple}(x_o, \omega_o) = \int_A \int_{\Omega} T_{dt} R_d(\|x_o - x_i\|) \cos \theta_i d\omega_i dA(x_i), \quad (5)$$

where ω_i is the direction of incident illumination at point x_i , and ω_o is the observed direction of emitted radiance at point x_o . $R_d(\|x_o - x_i\|)$ describes the diffusion of light entering at a point x_i and exiting at point x_o , and T_{dt} is given according to Equation 4. Our separation technique then further yields:

$$R_d(\|x_o - x_i\|) = R_{deep}(\|x_o - x_i\|) + R_{shallow}(\|x_o - x_i\|). \quad (6)$$

We employ the dipole diffusion model to approximate the deep scattering component $R_{deep}(\|x_o - x_i\|)$ from measured scattering profiles, assuming an infinitely deep dermis. Subsequently, we remove the effects of deep scattering from the measured scattering profiles using the dipole fit, and estimate scattering parameters for the shallow scattering $R_{shallow}(\|x_o - x_i\|)$ using the multipole model. We will discuss the modeling of both layers in detail in the remainder of this subsection.

Deep Scattering We model the deep scattering component using the dipole diffusion model [Jensen et al. 2001]:

$$R_{deep}(\|x_o - x_i\|) = \frac{\alpha'}{4\pi} \left(z_r(\sigma_{tr} + \frac{1}{d_r}) \frac{e^{-\sigma_{tr}d_r}}{d_r^2} + z_v(\sigma_{tr} + \frac{1}{d_v}) \frac{e^{-\sigma_{tr}d_v}}{d_v^2} \right), \quad (7)$$

where z_r (d_r) is the distance of the real source to the surface (x_o), and z_v (d_v) is the distance of the virtual source to the surface (x_o). This requires estimating two model parameters: the reduced albedo α' for x_o , and translucency (diffuse mean free path) $l_d = 1/\sigma_{tr}$. For optically dense materials, the following relation holds for α' :

$$R_{deep} = \frac{\alpha'}{2} \left(1 + e^{-\frac{4}{3}A\sqrt{3(1-\alpha')}} \right) e^{-\sqrt{3(1-\alpha')}}, \quad (8)$$

where R_{deep} is the diffuse albedo, and A is the internal reflection parameter that we compute as $\frac{1+\rho_d}{1-\rho_d}$ with ρ_d the reflectance of a rough specular surface due to hemispherical illumination. We employ the per-pixel R_{deep} values obtained from the separated indirect component (Fig. 5(b)), after factoring in the cosine falloff, to compute per-pixel α' values.

We estimate a per-region (Fig. 4(a)) translucency value l_d across the face from the scattering profiles observed by projecting a (polarized) solid white pattern with black dots on the face (Fig. 6(a)). The projected dots are 6mm in diameter and with 1cm spacing between them, which exceeds the typical scattering distance of light through skin. As in [Tariq et al. 2006], we pre-compute a lookup table of effective diffusion profiles due to such an illumination pattern in order to obtain estimates for l_d in various regions of the face. We prefer to use spatially-varying diffusion parameters instead of a using a modulation texture in our model as it results in a finer-scale control of the subsurface scattering. While this does not achieve the same accuracy to model a heterogeneous medium (i.e., skin) as with fully data-driven methods [Peers et al. 2006], the spatially varying parameters provide a flexible, yet compact, approximation for modeling the observed variation in different regions of the face.

The observed scattering profiles are the combined result of deep and shallow scattering. However, the extent of shallow scattering is much less than that of deep scattering. Therefore, by only considering the inner two-thirds of the projected black dots, the effects of shallow scattering are minimized, and a dipole fit can be computed.

Accurately localizing the dot boundaries is important for model fitting and is complicated by the blurring of the dot edges by the scattering. To localize the dot boundaries, we subtract the dot image from the fully-lit projector image Fig. 6(b), obtaining an image of illuminated blurry dots on a dark background. The zero-crossings of the difference between these negative and positive dot images reliably indicate sharp estimates of the dot boundaries as in Fig. 6(c). To use all of the information within each dot, we perform a radial average of the diffusion profile from the center going outwards to the dot periphery and use data up to two-thirds of the way (a 30 pixel radius) for the fitting process. Results of this fitting process are depicted in Fig. 6(e). As can be seen, the fitted dipole matches the observations closely in the last two-thirds (the fitted region), while exhibiting a larger error on the first third of the scattering profiles (extrapolated region). Finally, we average the translucency estimate from the dots in each region and blur the estimates across region boundaries.

Shallow Scattering Most of the first third of the scattering profiles observed under the black dot pattern is the result of both shallow and deep scattering. The deep scattering is estimated from the inner two-thirds, which we presume to be negligibly influenced by the shallow scattering. Fig. 6(e) illustrates this effect clearly. Using the estimated deep scattering dipole model, we can remove the effects of deep scattering from the observed scattering profiles, and fit an appropriate scattering model to the residual. We model shallow scattering in the top epidermal layer of skin with the multipole diffusion model [Donner and Jensen 2005]:

$$R_{shallow}(\|x_o - x_i\|) = \frac{\alpha'}{4\pi} \sum_{i=-n}^n \left(\frac{z_{r,i}(1 + \sigma_{tr}d_{r,i})e^{-\sigma_{tr}d_{r,i}}}{d_{r,i}^3} - \frac{z_{v,i}(1 + \sigma_{tr}d_{v,i})e^{-\sigma_{tr}d_{v,i}}}{d_{v,i}^3} \right).$$

We employ a similar fitting process to the deep scattering fit where an additional lookup table is employed for the residual profile using the shallow scattering albedo observed from the separated *direct* component (Fig. 5(a)). In our implementation, we use the multipole model with five dipoles and assume a layer depth of 0.5mm , which is roughly half the width of the projected separation patterns, for obtaining such a fit. We also assume an index of refraction of 1.38 for the top layer of skin. To further simplify the multipole fitting, we assume that there is no change in the index of refraction between the shallow and deep scattering layers.

5 Results

In this section, we present results rendered with our layered facial reflectance model and the corresponding fits obtained from the acquired data. To visualize the results, we modified the popular PBRT ray tracer [Pharr and Humphreys 2004] to support our facial reflectance model. To render subsurface scattering, we employ photon mapping [Jensen 2001], and added the dipole [Jensen et al. 2001] and multipole diffusion [Donner and Jensen 2005] models as a shader in PBRT. We modify the photon deposition phase to include the cosine of the incident photons and modulate by the transmittance at incidence. During the rendering phase, we switch off one-bounce gathering and use the spatially-varying dipole and multipole kernels respectively for density estimation with further modulation by the transmittance at exitance. We believe our facial reflectance model can also be easily incorporated in production rendering pipelines [Hery 2003].

Fig. 1 shows the layers of facial reflectance which comprise our renderings. The image second from the right shows an offline rendering of the face under novel illumination and viewpoint which is the composition of the layers modulated by the corresponding transmittance terms. At the right is a validation photograph from the side which was not used for reflectance modeling. Despite the significant change in viewpoint and relative lighting direction, the rendering closely resembles the photograph, including the spatially-varying specular and subsurface reflectance. Because our setup uses a single camera for reflectance modeling, some texture stretching can be observed at the sides of the nose. Without correction, the lips and parts of the eyelids will appear darker in the diffuse albedo than in the reference photograph, because the albedo is computed from images under full-on spherical illumination which includes partial occlusion from the lips and nose respectively. We correct the estimate of the diffuse albedo using an inverse simulation.

Fig. 7 illustrates the benefit of our layered model for acquired reflectance data with offline renderings of a female subject. Here, we qualitatively compare the layered rendering with a traditional rendering with acquired data including spatially-varying specular reflectance + single layer subsurface scattering. For the single layer rendering, we extract dipole diffusion parameters from the projected dot patterns similar to the fitting process for the deep scattering layer. Despite both methods using measured data from the same setup, the rendering with our layered reflectance model with additional single scattering and shallow and deep multiple scattering (e) looks much more skin-like compared to rendering with the traditional model for measured data (c), and is a closer match to the validation photograph (b). The deep multiple scattering is fit from observations that modulate incident irradiance by the absorption and transmittance of the shallow scattering layer. Hence, first order effects of interactions (reflectance and transmittance) between the shallow and deep scattering layers are automatically included in the estimated parameters of deep multiple scattering. While the employed dipole model may not fit the resultant scattering profiles perfectly, it better models the combined properties of the shallow and deep scattering layers, and reproduces the subtleties of skin ap-

Technique	Region	Translucency (l_d)		
		R	G	B
Deep Scattering	forehead	2.0293	0.9592	0.7179
	cheek	2.8926	0.9683	0.7607
Single Layer	forehead	1.8701	0.8793	0.6479
	cheek	2.3528	0.8788	0.6664
Weyrich et al. [2006]	cheek	1.8155	1.0213	0.6453
Jensen et al. [2001]	forearm	3.6733	1.3665	0.6827

Table 1: Measured Dipole Diffusion Translucency Parameters *Our measurements are consistent with measurements previously reported in the literature for faces.*

pearance better than a single layer model. The individual layers are shown in (a-b), and (f-i). Fig. 7(c) is the result of combining the single layer subsurface scattering component (a) and the specular layer (b) (+2 f-stops). Fig. 7(e) is the result of combining the four layers in our model: deep multiple scattering (f), shallow multiple scattering (g) (+2 f-stops), single scattering (h) (+5 f-stops), and the specular reflectance (i) (+2 f-stops). Note how the deep multiple scattering (f) contains less texture detail than the single layer approximation (a), which in turn contains less detail than the shallow multiple scattering layer (g).

Table 1 lists some of the dipole diffusion parameter fits we obtained from our measurements for the female subject and corresponding values reported in the literature as a means of quantitative validation of our technique. As can be seen, our estimated diffusion parameters are closer to those reported by Weyrich et al. [2006] for faces than those reported by Jensen et al. [2001] who measured the scattering on a skin patch on the forearm which is most likely more translucent than facial skin. In order to compare our extracted specular distributions for the Torrance-Sparrow model to those reported in the literature, we fit the raw data to a Gaussian distribution with roughness parameter m . The obtained region-wise fits of m for the female subject (nose = 0.2, eyes = 0.25, forehead = 0.3, cheeks = 0.325) are very similar to those reported by Weyrich et al. [2006]. We also estimated the per-channel single scattering Henyey-Greenstein phase function parameter g to be between 0.63 – 0.7 compared to 0.75 reported in [Hanrahan and Krueger 1993]. Our slightly lower values for g can be potentially attributed to the approximation of some amount of polarization preserving multiple scattering as single scattering in our model.

Fig. 8 shows rendering results from five acquired face models. The top row of Fig. 8 shows the ability of our model to reproduce the original front-lit illumination condition used for reflectance modeling for two subjects. A greenish tint near the top of the photographs results from uneven color in the cross-polarized video projector used as the illuminant. The corresponding renderings do not exhibit this effect since their albedo texture is derived from the spherical LED illumination. The middle row of Fig. 8 shows two side-by-side renderings of a male subject with light skin. The left pair shows the subject from the original left camera viewpoint but under novel illumination from an additional point light source. The right camera shows the subject from a novel viewpoint, illuminated from the frontal video projector. Both renderings substantially reproduce the subject’s appearance.

We have also implemented a real-time rendering approach with our acquired reflectance data which leverages hybrid normal maps [Ma et al. 2007] together with a local shading model that includes the inferred specular reflectance and single scattering, and which approximates subsurface scattering by a diffuse BRDF model. Results of this real-time rendering can be seen in the final row of Fig. 8(i), where a male subject with dark skin is rendered with from novel viewpoint together with a validation photograph. While the real-time renderings with this technique lack some of the subtle subsur-



Figure 7: Comparison Renderings with Measured Data (a) The single layer subsurface scattering component which together with the specular layer (b) (+2 f -stops) comprises rendering result with existing techniques (c). (d) Validation photograph. (e) Rendering with the modeled layered reflectance as described in this paper. (f-i) The four layers in our model that form (e): deep multiple scattering, shallow multiple scattering (+2 f -stops), single scattering (+5 f -stops), and specular reflectance (+2 f -stops). The layered rendering with measured data (e) yields a more realistic skin appearance compared to the rendering with specular + single layer subsurface scattering (c).

face scattering effects, we believe that such a rendering technique with hybrid normal maps could be useful for interactive applications. The technique of d'Eon et al. [2007] could potentially also be used with our data, albeit at a higher computation cost, to simulate a fuller range of subsurface scattering effects in real-time.

Finally, the female subject is rendered in a smiling pose with makeup from novel viewpoint in Fig. 8(k) together with a validation photograph. The female subject could be captured in a smiling pose due to the short five-second capture process. It would be difficult to keep a steady expression for longer acquisition times. Our data-driven facial reflectance model is also flexible enough to model such altered skin reflectance.

Limitations: In general, our renderings bear a close resemblance to the original photographs, successfully reproducing the appearance of a wide variety of skin tones and textures. However, due to the simplicity of our model, not all effects are modeled with equal accuracy. Subtle differences can arise due to differences in the specular roughness and diffuse reflectance within facial regions. Back-lighting effects are potentially not reproduced correctly with our model because the dipole diffusion model used for deep scattering is known to underestimate layered transmittance [Donner and Jensen 2005].

While the simplicity of our model introduces some limitations, it is also makes it a practical method that can be easily implemented in existing rendering systems. Additionally, because our model can be inferred from a few photographs and requires no physical contact device to measure scattering properties, it is more robust to changes due to subject motion or blood flow, and is able to capture the facial appearance of people in natural facial expressions that are hard to maintain for more than a few seconds.

6 Conclusion and Future Work

We have presented a practical method for measuring and modeling the appearance of a face from just twenty photographs captured from a single viewpoint under environmental and projected illumination. Key to the technique is the separation of facial appearance into different layers representing specular reflectance, single

scattering, and shallow and deep multiple scattering. Each layer is modeled by an appropriate model that can easily be incorporated in an existing rendering system. Our method is the first practical system that measures single scattering and spatially-varying multi-layer scattering parameters from a live subject. We have shown that the obtained parameters are quantitatively similar to those reported in the literature, and that resulting renderings are qualitatively a close match to reference photographs.

The presented system, due to its short acquisition times, enables new possibilities for analyzing time-varying effects of facial reflectance. For example one could monitor the change in skin reflectance due to blood flow or sweat, or examine the effects of facial animation on the appearance of skin.

For future work, we would like to investigate other methods than [Nayar et al. 2006] for separating the subsurface scattering layers of skin such as color space separation from the different color channels of our images. Furthermore, relating these separated layers more precisely to the multi-layer model of Donner and Jensen [2005] is another interesting avenue for future work. Finally, our model is currently missing a model for asperity scattering [Koenderink and Pont 2003], which is the result of scattering in a thin layer of facial hair and can play a significant role in reproducing the ‘velvety’ or ‘peachy’ facial appearance of many subjects.

Acknowledgments

We would like to thank Tomas Pereira and Wan-Chun Ma for their valuable input, Cynthia Richards, David Price, and Krishna Mamidibathula for sitting as subjects, Marko Vukovic, Jay Busch, and Jen-Yuan Chiang for their assistance in data processing, Bill Swartout, Randy Hill, Randolph Hall, and Max Nikias for their support and assistance with this work. We also thank Wolfgang Heidrich and our anonymous reviewers for their valuable suggestions for improving the paper. This work was sponsored by the University of Southern California Office of the Provost and the U.S. Army Research, Development, and Engineering Command (RDE-COM). The content of the information does not necessarily reflect the position or the policy of the US Government, and no official endorsement should be inferred.

References

- ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. S. 2000. A microfacet-based BRDF generator. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, 65–74.
- ASHIKHMIN, M., 2006. Distribution-based BRDFs. <http://jesper.kalliope.org/blog/library/dbrdfs.pdf>.
- DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, 145–156.
- D'EON, E., LUEBKE, D., AND ENDERTON, E. 2007. Efficient rendering of human skin. In *Eurographics Symposium on Rendering 2007*.
- DONNER, C., AND JENSEN, H. W. 2005. Light diffusion in multi-layered translucent materials. *ACM Transactions on Graphics* 24, 3 (Aug.), 1032–1039.
- DONNER, C., AND JENSEN, H. W. 2006. A spectral BSSRDF for shading human skin. In *Rendering Techniques 2006: 17th Eurographics Workshop on Rendering*, 409–418.
- DONNER, C., WEYRICH, T., D'EON, E., RAMAMOORTHI, R., AND RUSINKIEWICZ, S. 2008. A layered, heterogeneous reflectance model for acquiring and rendering human skin. In *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*.
- FUCHS, M., BLANZ, V., LENSCH, H., AND SEIDEL, H.-P. 2005. Reflectance from images: a model-based approach for human faces. *IEEE Transactions on Visualization and Computer Graphics* 11, 3 (May/June), 296–305.
- GEORGHIADES, A., BELHUMEUR, P., AND KRIEGMAN, D. 1999. Illumination-based image synthesis: Creating novel images of human faces under differing pose and lighting. In *IEEE Workshop on Multi-View Modeling and Analysis of Visual Scenes*, 47–54.
- GEORGHIADES, A. 2003. Recovering 3-D shape and reflectance from a small number of photographs. In *Rendering Techniques*, 230–240.
- GOESELE, M., LENSCH, H. P. A., LANG, J., FUCHS, C., AND SEIDEL, H.-P. 2004. Disco: acquisition of translucent objects. *ACM Transactions on Graphics* 23, 3 (Aug.), 835–844.
- HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. In *Proceedings of SIGGRAPH 93*, Computer Graphics Proceedings, Annual Conference Series, 165–174.
- HERY, C. 2003. Implementing a skin BSSRDF (or several). In *SIGGRAPH course notes: Renderman Theory and Practices*.
- IGARASHI, T., NISHINO, K., AND NAYAR, S. K. 2007. The appearance of human skin: A survey. In *Foundations and Trends in Computer Graphics and Vision*, vol. 3. 1–95.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, 511–518.
- JENSEN, H. W. 2001. *Realistic Image Synthesis Using Photon Mapping*. A. K. Peters.
- KOENDERINK, J., AND PONT, S. 2003. The secret of velvety skin. *Mach. Vision Appl.* 14, 4, 260–268.
- MA, W.-C., HAWKINS, T., PEERS, P., CHABERT, C.-F., WEISS, M., AND DEBEVEC, P. 2007. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Rendering Techniques*, 183–194.
- MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., TORRANCE, K. E., AND GREENBERG, D. P. 1999. Image-based BRDF measurement including human skin. In *Eurographics Rendering Workshop 1999*.
- MARSCHNER, S., GUENTER, B., AND RAGHUPATHY, S. 2000. Modeling and rendering for realistic facial animation. In *Rendering Techniques 2000: 11th Eurographics Workshop on Rendering*, 231–242.
- MORGAN, S. P., AND RIDGWAY, M. E. 2000. Polarization properties of light backscattered from a two layer scattering medium. *Opt. Express* 7, 12 (Dec), 395–402.
- NAYAR, S. K., KRISHNAN, G., GROSSBERG, M. D., AND RASKAR, R. 2006. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Transactions on Graphics* 25, 3 (July), 935–944.
- PEERS, P., VOM BERGE, K., MATUSIK, W., RAMAMOORTHI, R., LAWRENCE, J., RUSINKIEWICZ, S., AND DUTRÉ, P. 2006. A compact factored representation of heterogeneous subsurface scattering. *ACM Transactions on Graphics* 25, 3 (July), 746–753.
- PHARR, M., AND HUMPHREYS, G. 2004. *Physically Based Rendering*. Morgan Kaufmann, New York.
- TARIQ, S., GARDNER, A., LLAMAS, I., JONES, A., DEBEVEC, P., AND TURK, G. 2006. Efficient estimation of spatially varying subsurface scattering parameters. In *Proc. 11th Int'l Fall Workshop on Vision, Modeling, and Visualization (VMV 2006)*, 165–174.
- TONG, X., WANG, J., LIN, S., GUO, B., AND SHUM, H.-Y. 2005. Modeling and rendering of quasi-homogeneous materials. *ACM Transactions on Graphics* 24, 3 (Aug.), 1054–1061.
- TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory of off-specular reflection from roughened surfaces. *J. Opt. Soc. Am.* 57, 1104–1114.
- TUCHIN, V. 2007. *Tissue Optics: Light Scattering Methods and Instruments for Medical Diagnosis*, 2nd ed. SPIE Press.
- WANG, J., ZHAO, S., TONG, X., LIN, S., LIN, Z., DONG, Y., GUO, B., AND SHUM, H.-Y. 2008. Modeling and rendering of heterogeneous translucent materials using the diffusion equation. *ACM Transactions on Graphics* 27, 1, 1–18.
- WEYRICH, T., MATUSIK, W., PFISTER, H., BICKEL, B., DONNER, C., TU, C., MCANDLESS, J., LEE, J., NGAN, A., JENSEN, H. W., AND GROSS, M. 2006. Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics* 25, 3, 1013–1024.
- WOODHAM, R. J. 1978. Photometric stereo: A reflectance map technique for determining surface orientation from image intensity. In *Proc. SPIE's 22nd Annual Technical Symposium*, vol. 155.
- ZICKLER, T., RAMAMOORTHI, R., ENRIQUE, S., AND BELHUMEUR, P. N. 2006. Reflectance sharing: Predicting appearance from a sparse set of images of a known shape. *IEEE Trans. Pattern Anal. Mach. Intell.* 28, 8, 1287–1302.




Figure 8: Rendering Results and Validation Photographs (a,c) Offline renderings of two subjects under frontal point-source illumination, showing our technique's ability to replicate the appearance shown in the reference photographs in (b,d). (e,g) Offline renderings of a male subject in novel lighting and viewpoint conditions and corresponding validation photographs (f,h). (i) Real-time rendering using hybrid normal maps of a male subject with dark skin rendered from a novel viewpoint and validation photograph (j). (k) Offline rendering of a female subject in a dynamic pose wearing makeup and a validation photograph (l).