Analysis of Human Faces using a Measurement-Based Skin Reflectance Model

Tim Weyrich * Wojciech Matusik † Hanspeter Pfister † Bernd Bickel * Craig Donner ‡ Chien Tu †
Janet McAndless † Jinho Lee † Addy Ngan § Henrik Wann Jensen ‡ Markus Gross *

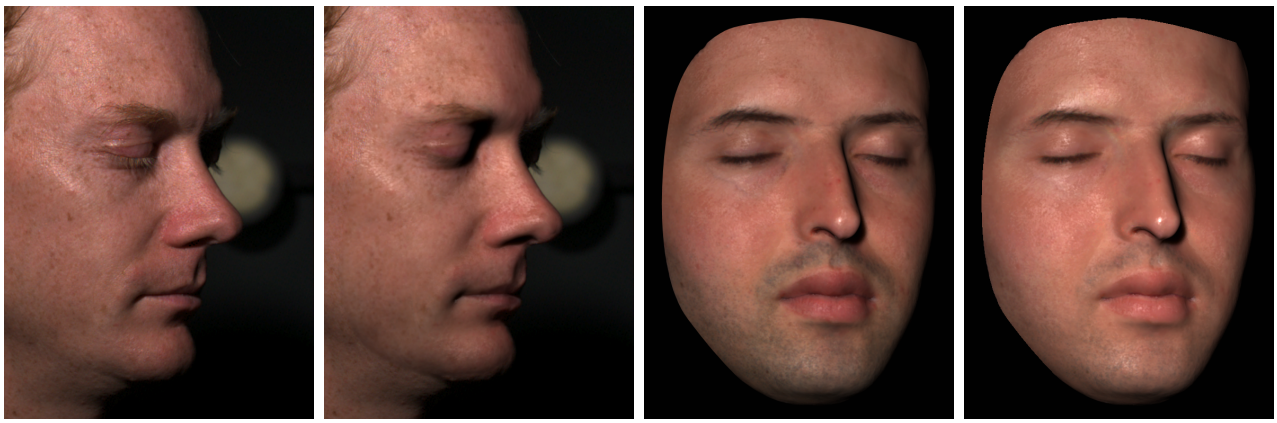

Photograph Rendering Original Model Appearance Change

Figure 1: *Photograph compared to a face rendered using our skin reflectance model. The rendered image was composited on top of the photograph. Right: Changing the albedo and BRDF using statistics of measured model parameters from a sample population.*

Abstract

We have measured 3D face geometry, skin reflectance, and subsurface scattering using custom-built devices for 149 subjects of varying age, gender, and race. We developed a novel skin reflectance model whose parameters can be estimated from measurements. The model decomposes the large amount of measured skin data into a spatially-varying analytic BRDF, a diffuse albedo map, and diffuse subsurface scattering. Our model is intuitive, physically plausible, and – since we do not use the original measured data – easy to edit as well. High-quality renderings come close to reproducing real photographs. The analysis of the model parameters for our sample population reveals variations according to subject age, gender, skin type, and external factors (e.g., sweat, cold, or makeup). Using our statistics, a user can edit the overall appearance of a face (e.g., changing skin type and age) or change small-scale features using texture synthesis (e.g., adding moles and freckles). We are making the collected statistics publicly available to the research community for applications in face synthesis and analysis.

Keywords: Face Modeling, Reflection Models, Data-Driven Models

*ETH Zürich, Email: [weyrich,bickel,grossm]@inf.ethz.ch
†MERL, Email: [matusik,pfister]@merl.com
‡UCSD, Email: [cdonner,henrik]@cs.ucsd.edu
§MIT, Email: addy@graphics.csail.mit.edu

1 Introduction

One of the most difficult computer graphics challenges is creating digital faces that are indistinguishable from real ones. The process of capturing an actor's performance and likeness has accurately been named "Digital Face Cloning." Digital face cloning has many applications in movies, games, medicine, cosmetics, computer vision, biometrics, and virtual reality. Despite recent progress, however, it remains a challenging problem because humans are incredibly sensitive to the subtleties of real faces.

In this paper, we focus on modeling skin reflectance of human faces, an important aspect of face appearance. Our primary goal is to accurately capture and reproduce the face reflectance of a person. We use custom-built devices to measure 3D face geometry, light reflection, and subsurface scattering. Our reflectance model consists of a spatially-varying analytic reflectance model (Torrance-Sparrow or Blinn-Phong), a diffuse albedo map, and diffuse scattering coefficients. We concisely describe the mathematics of the model and show how its components can be robustly estimated from measured data. The model is compact, intuitive, and easy to use. Images generated from it come very close to real photographs (see Figure 1).

A second goal is to capture and describe the variation of face reflectance for a sample population. We measured the parameters of our reflectance model for 149 people of different race, gender, and age. The result is a unique database of spatially-varying parameters for the Torrance-Sparrow and Blinn-Phong reflectance models, histograms of skin albedo, and diffuse scattering parameters for faces. From this data we computed the distributions of the parameters (most of them are Gaussian) and analyzed them using statistical methods. We publish the key findings of our analysis in this paper, and make the complete statistical data available to the research community. To our knowledge, this is the first study of face reflectance of this kind.

The results of our analysis have applications in realistic face synthesis, face reconstruction from photographs, inverse rendering, and

computer vision. We also demonstrate how an artist can edit faces using the parameter distributions. For example, we can change the skin type of a model by transferring statistics of albedo and reflectance parameters (see Figure 1). We can also use texture synthesis/analysis to transfer small-scale skin features, such as moles and freckles. The combination of these approaches allows an artist to convincingly transfer face appearance between individuals.

2 Previous Work

Properties of human skin have been measured and studied extensively in the computer graphics, biomedical, and computer vision communities. An excellent survey on the physiological and anatomical properties of skin and the state of the art in skin appearance modeling has been published by Igarashi et al. [2005]. In this section we provide an overview of the relevant work in the area of face appearance modeling in computer graphics.

Analytic Face Appearance Models: The simplest reflectance model for skin is the Bidirectional Reflectance Distribution Function (BRDF) [Nicodemus et al. 1977]. Dana et al. [1999] measured the BRDF of skin and made the data available in the Columbia-Utrecht Reflectance and Texture (CuReT) database (http://www.cs.columbia.edu/CAVE/curet/). Because the data was not collected *in vivo* it does not accurately reflect the properties of living skin. Marschner et al. [1999; 2000b] measured spatially uniform BRDFs on the foreheads of four subjects of different gender, race, and age. They observed that the BRDF of skin exhibits strong forward scattering in off-specular direction at grazing angles. In later work they augmented the uniform specular BRDF with a spatially-varying albedo texture [Marschner et al. 2000a]. Their BRDF model was used in the human face project at Walt Disney Studios [Williams 2005]. That project, which had a key influence in the historical development of human face models (e.g., it was influential on the process developed for the *Matrix* sequels), was completed in 2001. Other researchers have estimated parameters of uniform analytic BRDF models from photographs with static [Blanz and Vetter 1999; Paris et al. 2003] or varying illumination [Georghiades 2003; Fuchs et al. 2005b]. Fuchs et al.[2005b] cluster different parts of the face, such as cheek, forehead, and lips, into regions. They estimate a spatially-varying diffuse albedo but assume specularity to be constant within each region. The fundamental limitation of all these models is that they ignore subsurface scattering of light that is largely responsible for the soft appearance of facial skin.

Hanrahan and Krueger [1993] were the first to model subsurface scattering in skin using a diffuse constant and a lighting-dependent single scattering term. They described scattering and absorption parameters of skin using the Henyey-Greenstein function and compute scattering with Monte Carlo simulations. A similar approach was used by Ng and Li [2001]. To avoid expensive Monte Carlo simulations, Stam [2001] developed an analytic approximation to multiple subsurface scattering for skin with a rough surface. All of these models describe skin reflectance using a BRDF and ignore light scattering between different incoming and outgoing locations on the surface.

To describe the full effect of light scattering between two points on the surface one can use the Bidirectional Surface-Scattering Distribution Function (BSSRDF). Krishnaswamy et al. [Krishnaswamy and Baranoski 2004] introduced a BSSRDF model for skin with spectral and biophysically-based parameters. The BSSRDF is eight-dimensional, making it very difficult to measure and represent. Jensen et al. [2001] were the first to propose a practical low-dimensional analytic approximation for the BSSRDF of homogeneous materials using diffusion theory. To estimate the parameters of their model, they measure scattering of light in skin on a forearm *in vivo* using a white focused beam of light and a camera. We use their approximation for subsurface scattering. To estimate the parameters of their model we developed a measurement device that can safely be used in faces (see Section 5.4). Donner and Jensen [2005] have applied a multi-layer diffusion model to the problem of rendering skin. Unfortunately, measuring the parameters of each layer of multi-layered translucent materials is a difficult problem.

Image-Based Face Models: Image-based models have provided highly realistic representations for human faces because they implicitly capture effects such as self-shadowing, inter-reflections, and subsurface scattering. Pighin et al. [1998] use view-dependent texture mapping [Debevec et al. 1996] to reproduce faces under static illumination conditions. More recent efforts allow variations in lighting for static faces [Georghiades et al. 1999; Debevec et al. 2000], expressions [Hawkins et al. 2004], and real-time performances [Wenger et al. 2005]. Debevec et al. [2000] present a process for creating realistic, relightable 3D face models by mapping image-based reflectance characteristics onto 3D-scanned geometry. In order to change the viewpoint, they use color-space analysis to separate the image data into specular and diffuse components that can be extrapolated to new viewpoints. While their method does consider the aggregate behavior of subsurface scattering, they do not model a specific diffusion parameter. Thus, unlike our approach, their method cannot produce correct subsurface scattering effects for closeup light sources or high-frequency spatially-varying illumination. Borshukov and Lewis [2003] combine an image-based model, an analytic surface BRDF, and an image-space approximation for subsurface scattering to create highly realistic face models for the movie industry. Sander et al. [2004] developed a variant of this method for real-time face rendering on modern graphics hardware.

Image-based models come at the price of extensive data storage. They are difficult to edit, and the measurements are not related to any physical model of light-skin interaction. It is very difficult to simulate local light sources and to reproduce spatially-varying illumination [Jones et al. 2005], especially if shadows from a distant light disagree with shadows from a local light. An example of this would be a light source close to the nose – even though the skin is illuminated, an image-based method would predict zero reflection since a distant light from that direction would be blocked by the nose. In contrast to image-based methods, we measure a full BRDF at each surface point using a face-scanning dome similar to the LightStage by Debevec et al. [2002], but with more cameras (16). The resulting parametric face reflectance model has lower storage requirements than a full surface reflectance field. In addition it is physically plausible, easy to render, and intuitive to edit.

Skin Color and Texture Models: Biophysical studies show that skin appearance is largely dependent on wavelength [Igarashi et al. 2005]. Angelopoulou et al. [2001] measured skin BRDF as a function of wavelength for a sample population of 22 people and noted that five Gaussian basis functions reproduce the data well. The study was limited to 0° incident and 4° reflection angle.

Tsumura et al. [2003] used Independent Component Analysis (ICA) in color space to decompose images of faces into melanin and hemoglobin layers. To change skin color and transfer skin texture they adapted the texture synthesis method of Heeger and Bergen [1995]. We use a similar approach to transfer texture (see Section 8). Cula et al. [2005; 2004] collected the Bidirectional Texture Function (BTF) for different body regions of a sample population (http://www.caip.rutgers.edu/rutgers_texture/cvg/). They analyzed and classified skin texture using histograms of image-pyramid features [Cula and Dana 2002]. We use a similar analysis technique for albedo maps (see Section 7).

3 Skin Reflectance Model

The BSSRDF is a function $S(x_i, \omega_i, x_o, \omega_o)$, where ω_i is the direction of the incident illumination at point x_i, and ω_o is the observation direction of radiance emitted at point x_o. The outgoing radiance $L(x_o, \omega_o)$ at a surface point x_o in direction ω_o can be computed by integrating the contribution of incoming radiance $L(x_i, \omega_i)$ for all incident directions Ω over the surface A:

$$L(x_o, \omega_o) = \int_A \int_\Omega S(x_i, \omega_i, x_o, \omega_o) L(x_i, \omega_i)(N \cdot \omega_i) d\omega_i dA(x_i). \quad (1)$$

We will use the notation dx_i for $(N \cdot \omega_i) dA(x_i)$.

We separate this integral into a specular BRDF term and a diffuse subsurface scattering term:

$$L(x_o, \omega_o) = L_{specular}(x_o, \omega_o) + L_{diffuse}(x_o, \omega_o), \quad (2)$$

with:

$$L_{specular}(x_o, \omega_o) = \int_\Omega f_s(x_o, \omega_o, \omega_i) L(x_o, \omega_o)(N \cdot \omega_i) d\omega_i, \quad (3)$$

$$L_{diffuse}(x_o, \omega_o) = \int_A \int_\Omega S_d(x_i, \omega_i, x_o, \omega_o) L(x_i, \omega_i) d\omega_i dx_i. \quad (4)$$

$S_d(x_i, \omega_i, x_o, \omega_o)$ is called the diffuse BSSRDF and $f_s(x_o, \omega_o, \omega_i)$ is a surface BRDF. As observed in [Jensen and Buhler 2002], multiple scattering dominates in the case of skin and can be modeled as a purely diffuse subsurface scattering term and a specular term that ignores single scattering (see Figure 2).

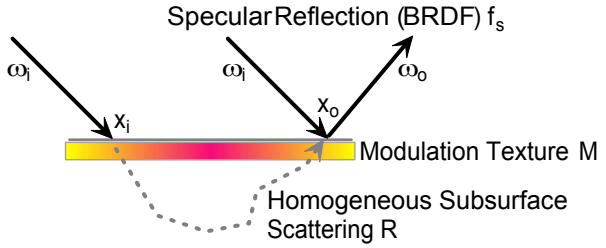

Figure 2: *Spatially-varying skin reflectance can be explained by a specular BRDF f_s at the air-oil interface, and a diffuse reflectance component due to subsurface scattering. We model the diffuse reflectance with a spatially-varying modulation texture M and a homogeneous subsurface scattering component R_d.*

The diffuse BSSRDF S_d is very difficult to acquire directly. Consequently, it is often decomposed into a product of lower-dimensional functions. A commonly used decomposition for S_d is:

$$S_d(x_i, \omega_i, x_o, \omega_o) \approx \frac{1}{\pi} f_i(x_i, \omega_i) R_d(x_i, x_o) f_o(x_o, \omega_o). \quad (5)$$

Here $R_d(x_i, x_o)$ describes the diffuse subsurface reflectance of light entering at point x_i and exiting at point x_o. The diffuse surface reflectance/transmittance functions $f_i(x_i, \omega_i)$ and $f_o(x_o, \omega_o)$ further attenuate the light at the entrance and exit points based on the incoming and outgoing light directions, respectively.

The complete formula for our skin model is:

$$L(x_o, \omega_o) = \int_\Omega f_s(x_o, \omega_o, \omega_i) L(x_o, \omega_i)(N \cdot \omega_i) d\omega_i + \int_A \int_\Omega F_{ti} R_d(\|x_o - x_i\|) M(x_o) F_{to} L(x_i, \omega_i) d\omega_i dx_i. \quad (6)$$

Like Hanrahan and Krueger [1993] and Jensen et al. [2001], we assume that skin is mostly smooth and use transmissive Fresnel functions $F_{ti} = F_t(\eta, \omega_i)$ and $F_{to} = F_t(\eta, \omega_o)$. We use $\eta \approx 1.38$ as the relative index of refraction between skin and air for the Fresnel terms [Tuchin 2000].

We model the diffuse reflectance using a homogeneous subsurface reflectance term $R_d(\|x_o - x_i\|)$ and a spatially-varying *modulation texture* $M(x_o)$ [Fuchs et al. 2005a]. Intuitively, this modulation texture represents a spatially-varying absorptive film of zero thickness on top of homogeneously scattering material (see Figure 2). To model homogeneous subsurface scattering we use the dipole diffusion approximation by Jensen et al. [2001]. It simplifies $R_d(x_i, x_o)$ to a rotationally-symmetric function $R_d(\|x_o - x_i\|)$, or $R_d(r)$ with radius r. Jensen and Buhler [Jensen and Buhler 2002] develop fast rendering methods for this approximation, and Donner and Jensen [2005] extend it to a multipole approximation for multiple layers of translucent materials.

Our formulation in Equation (6) was motivated by the skin model shown in Figure 2. Goesele et al. [2004] assume $f_i(x_i, \omega_i)$ and $f_o(x_o, \omega_o)$ to be constant. To model inhomogeneous translucent materials, they exhaustively measure $R_d(x_i, x_o)$ and tabulate the results. In follow-up work [Fuchs et al. 2005a] they fit a sum of Gaussians to the tabulated values of $R_d(x_i, x_o)$ to model anisotropic scattering. However, acquiring $R_d(x_i, x_o)$ directly requires a substantial experimental setup (controlled laser illumination from different angles) that is not applicable for human faces.

Donner and Jensen [2005] use spatially-varying diffuse transmission functions $\rho_{dt}(x_i, \omega_i) = 1 - \int_\Omega f_s(x_i, \omega_i, \omega_o)(\omega_o \cdot n) d\omega_o$ (similarly for $\rho_{dt}(x_o, \omega_o)$). They integrate a BRDF at the surface over the hemisphere of directions and tabulate the results to obtain the diffuse transmission functions. To capture materials with homogeneous subsurface scattering and complex surface mesostructure (so-called quasi-homogeneous materials), Tong et al. [2005] measure and tabulate a spatially-varying exit function $f_o(x_o, \omega_o)$ and a uniform entrance function $f_i(\omega_i)$. We now give an overview of how we measure and estimate the parameters of our model.

4 Parameter Estimation Overview

A block diagram of our processing pipeline is shown in Figure 3. We first capture the 3D geometry of the face using a commercial 3D

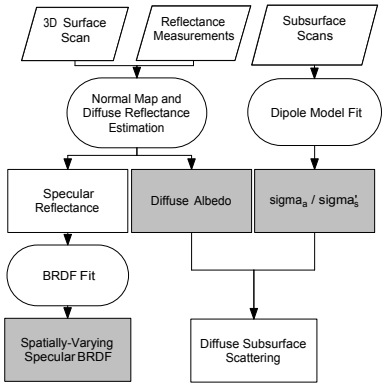

Figure 3: *A block diagram of our data processing pipeline. Blocks in grey are the parameters of our skin reflectance model.*

face scanner. Then we measure the total skin reflectance $L(x_o, \omega_o)$ in a calibrated face-scanning dome by taking photographs from different viewpoints under varying illumination. We refine the geometry and compute a high-resolution face mesh and a normal map using photometric stereo methods. To map between 3D face space and 2D (u, v) space we use the conformal parameterization of Sheffer et al. [2005]. The data is densely interpolated using push-pull interpolation into texture maps of 2048×2048 resolution.

The reflectance measurements are registered to the 3D geometry and combined into lumitexels [Lensch et al. 2001]. Each lumitexel stores samples of the skin BRDF $f_r(x_o, \omega_i, \omega_o)$. We estimate the diffuse reflectance $\alpha(x_o)$ of f_r using a stable minimum of lumitexel samples. Note that this diffuse reflectance encompasses the effects of the diffuse subsurface scattering $R_d(r)$ and the modulation texture $M(x_o)$. We store $\alpha(x_o)$ in a diffuse albedo map and subtract $\frac{1}{\pi}F_t(\eta,\omega_i)\alpha(x_o)F_t(\eta,\omega_o)$ from $f_r(x_o,\omega_i,\omega_o)$. To the remaining specular reflectance $f_s(x_o,\omega_i,\omega_o)$ we fit the analytic BRDF function.

The diffuse subsurface reflectance $R_d(r)$ and the modulation texture $M(x_o)$ cannot be measured directly on human subjects. Instead, we measure their combined effect $R(r)$ at a few locations in the face using a special contact device. To our measurements $R(r)$ we fit the analytic dipole approximation of Jensen et al. [2001] and store the reduced scattering coefficient σ_s' and the absorption coefficient σ_a. Similar to Donner and Jensen [2005] we derive that:

$$R_d(r)M(x_o) = R(r)\frac{\alpha(x_o)}{\alpha_{avg}}, \tag{7}$$

where α_{avg} is the average albedo. We use the albedo map $\alpha(x_o)$, the average of all albedo map values α_{avg}, and the dipole approximation for $R(r)$ to compute $R_d(r)M(x_o)$ in Equation (6).

In summary, the parameters for each face are the spatially-varying coefficients of the specular BRDF f_s, an albedo map with diffuse reflectance values $\alpha(x_o)$, and the values for σ_s' and σ_a. Section 6 describes how this data is used for rendering. We now describe each of these measurement and processing steps in more detail.

5 Measurement System

Figure 4 shows a photograph of our face-scanning dome. The subject sits in a chair with a headrest to keep the head still during the capture process. The chair is surrounded by 16 cameras and 150 LED light sources that are mounted on a geodesic dome. The system sequentially turns on each light while simultaneously capturing images with all 16 cameras. We capture high-dynamic range (HDR) images [Debevec and Malik 1997] by immediately repeating the capture sequence with two different exposure settings. The complete sequence takes about 25 seconds for the two passes through all 150 light sources (limited by the frame rate of the cameras). To minimize the risk of light-induced seizures we ask all subjects to close their eyes. We report more details about the system and its calibration procedure in [Weyrich et al. 2005].

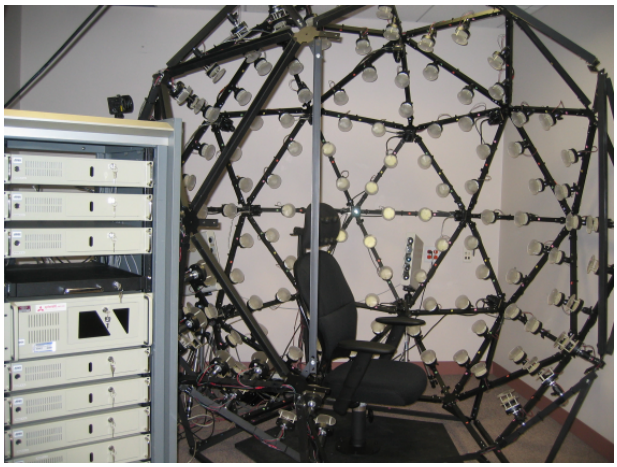

Figure 4: *The face-scanning dome consists of 16 digital cameras, 150 LED light sources, and a commercial 3D face-scanning system.*

A commercial face-scanning system from 3QTech (www.3dmd.com) is placed behind openings of the dome. Using two structured light projectors and four cameras, it captures the complete 3D face geometry in less than one second. The output mesh contains about 40,000 triangles and resolves features as small as 1 mm. We clean the output mesh by manually cropping non-facial areas and fixing non-manifold issues and degenerate triangles. The cleaned mesh is refined using Loop-subdivision [Loop 1987] to obtain a high-resolution mesh with about 700,000 to 900,000 vertices. The subdivision implicitly removes noise.

Next, we compute a lumitexel [Lensch et al. 2001] at each vertex from the image reflectance samples. The lumitexel radiance values $L(x_o,\omega_o)$ are normalized by the irradiance $E(x_o,\omega_i)$ to obtain BRDF sample values:

$$f_r(x_o,\omega_i,\omega_o) = \frac{L(x_o,\omega_o)}{E(x_o,\omega_i)} = \frac{L(x_o,\omega_o)}{\int_\Omega L(\omega_i)(\omega_i\cdot n)d\omega_i}. \tag{8}$$

We calibrate the radiance $L(\omega_i)$ of each light source using a white ideal diffuse reflector (Fluorilon) and explicitly model the beam spread by a bi-variate 2nd degree polynomial [Weyrich et al. 2005]. We use shadow mapping to determine the visibility of lumitexel samples for each light source and camera. On average, a lumitexel contains about 900 BRDF samples per color channel, with many lumitexels having up to 1,200 samples. The numbers vary depending on the lumitexel visibility. All processing is performed on RGB data except where noted otherwise.

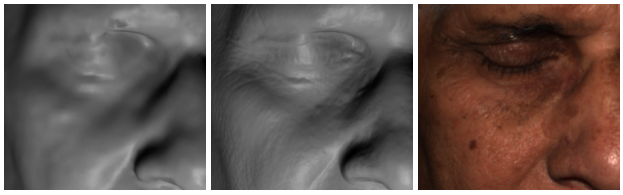

Figure 5: *Facial detail shown in closeups. Left: Measured geometry before correction. Middle: Corrected high-quality geometry. No normal map was used; the surface detail is due to the actual geometry of the model. Right: Comparison photo.*

5.1 Geometry Processing

We estimate normals at each lumitexel from the reflectance data using photometric stereo [Barsky and Petrou 2001]. We compute a normal map for each of the camera views. Using the method by Nehab et al. [2005] we correct low-frequency bias of the normals. For each camera, we also compute the cosine between it and each normal and store the result as a blending weight in a texture. To compute the final *diffuse normal map*, we average all normals at each vertex weighted by their blending weights. To avoid seams at visibility boundaries between cameras, we slightly smooth the blending texture and the diffuse normal map using a Gaussian filter. Finally, we compute the high-resolution geometry by applying the method of Nehab et al. Figure 5 compares the measured 3QTech geometry to the final high-resolution mesh.

Figure 6 (a) shows a rendering of a face using the Torrance-Sparrow BRDF model and the diffuse normal map. The rendering does not accurately reproduce the specular highlights that are visible in the input photo in Figure 6 (d). For example, compare the detailed highlights on the cheeks in (d) with the blurred highlight in (a). These highlights are due to micro-geometry, such as pores or very fine wrinkles, which is not accurately captured in the face geometry or diffuse normal map.

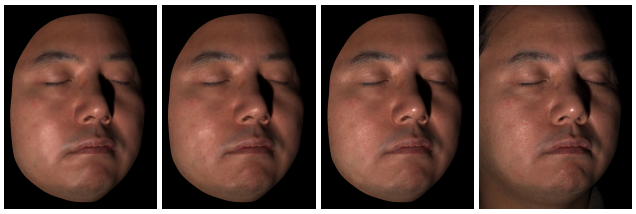

Figure 6: *Torrance-Sparrow model and (from left to right) (a) only diffuse normal map; (b) one diffuse and one global micro-normal map for specular reflection; (c) one diffuse and per-view micro-normal maps; (d) input photograph. Method (c) most accurately reproduces the specular highlights of the input image.*

To improve the result, we use *micro-normals* to render the specular component. Debevec et al. [Debevec et al. 2000] used micro-normals for a Torrance-Sparrow model on their specular term. They estimated improved normals using a specular reflection image obtained using cross-polarized lighting. For each camera viewpoint, we choose the sample with maximum value in each lumitexel. The half-way vector between the view direction and the sample's light source direction is the micro-normal for this view. We either store these normals in per-camera micro-normal maps, or generate a single micro-normal map using the blending weights from before. Figures 6 (b) and (c) show renderings with single or per-view micro-normal maps, respectively. Figure 6 (b) is better than (a), but only Figure 6 (c) accurately reproduces the highlights of the input image. The results in this paper were computed using per-view micro-normals. For between-camera views and animations we interpolate the micro-normals with angular interpolation [Debevec et al. 1996].

5.2 Diffuse Albedo Estimation

We estimate the total diffuse reflectance α at each surface point from the lumitexel data. We assume that surface reflectance vanishes for at least some of the observation angles, revealing pure diffuse reflectance. Accordingly, the total diffuse reflectance is the maximum α that is compliant with the observations, i.e., $(1/\pi)\alpha$ does not exceed the Fresnel-normalized BRDF samples $f_r/(F_t(\eta,\omega_i)F_t(\eta,\omega_o))$. We compute α as the minimum ratio between this Fresnel-normalized f_r and the unit diffuse reflectance:

$$\alpha(x_o) = \min_i \frac{\pi f_r(x_o,\omega_o,\omega_i)}{F_t(\eta,\omega_i)F_t(\eta,\omega_o)}. \tag{9}$$

In order to reduce outliers and the influence of motion artifacts, we determine a stable minimum by penalizing grazing observations and discarding the k smallest values. The α values for each surface point are re-parameterized, interpolated, and stored as the diffuse albedo map.

5.3 BRDF Fit

The specular reflection is computed at each lumitexel by subtracting the total diffuse reflectance:

$$f_s(x_o,\omega_o,\omega_i) = f_r(x_o,\omega_o,\omega_i) - \frac{\alpha(x_o)}{\pi} F_t(\eta,\omega_i)F_t(\eta,\omega_o). \tag{10}$$

Highlights on dielectric materials like skin are of the same color as the light source (white, in our case). Consequently, we convert the specular reflectance to grey scale to increase the stability of the BRDF fit. Unlike previous work [Marschner et al. 1999; Lensch et al. 2001; Fuchs et al. 2005b] we do not need to cluster reflectance data because our acquisition system collects enough samples for a spatially-varying BRDF fit at each vertex. The data for lumitexels with a badly conditioned fit is interpolated during creation of the BRDF coefficient texture.

We fit the Blinn-Phong [Blinn 1977] and Torrance-Sparrow [Torrance and Sparrow 1967] isotropic BRDF models to the data. The Blinn-Phong model is the most widely used analytic reflectance model due to its simplicity:

$$f_{sBP} = \rho_s \frac{n+2}{2\pi} \cos^n \delta. \tag{11}$$

Here δ is the vector between the normal N and the half-way vector H, ρ_s is the scaling coefficient, and n is the specular exponent. The factor $(n+2)/2\pi$ is for energy normalization so that the cosine lobe always integrates to one. This assures that 100% of the incident energy is reflected for $\rho_s = 1$.

We also use the physically-based Torrance-Sparrow model:

$$f_{sTS} = \rho_s \frac{1}{\pi} \frac{DG}{(N\cdot\omega_i)(N\cdot\omega_o)} F_r(\omega_o\cdot H), \tag{12}$$

with:

$$G = \min\{1, \frac{2(N\cdot H)(N\cdot\omega_o)}{(\omega_o\cdot H)}, \frac{2(N\cdot H)(N\cdot\omega_i)}{(\omega_o\cdot H)}\}, \tag{13}$$
$$D = \frac{1}{m^2\cos^4\delta} \exp^{-[(\tan\delta)/m]^2}.$$

G is the geometry term, D is the Beckmann micro-facet distribution, and F_r is the reflective Fresnel term. The free variables are the scaling coefficient ρ_s and the roughness parameter m.

For the fit we use the fitting procedure by Ngan et al. [2005]. The objective function of the optimization is the mean squared error between the measured BRDF f_s and the target model f_{sM} (either f_{sBP} or f_{sTS}) with parameter vector $\mathbf{p}$:

$$E(\mathbf{p}) = \sqrt{\frac{\sum w[f_s(\omega_i,\omega_o)\cos\theta_i - f_{sM}(\omega_i,\omega_o;\mathbf{p})\cos\theta_i]^2}{\sum w}}. \tag{14}$$

The sum is over the non-empty samples of the lumitexel, and θ_i is the elevation angle of incident direction. The weight w is a correction term that allows us to ignore data with ω_i or ω_o greater than 80°. Measurements close to extreme grazing angles are generally unreliable due to imprecisions in the registration between 3D geometry and camera images. We apply constrained nonlinear optimization based on sequential quadratic programming over the specular lobe parameters to minimize the error metric.

To assess the quality of our fit, we compare renderings from camera viewpoints to the actual photographs. The camera and light source calibration data were used to produce identical conditions in the renderings. Figure 7 shows a comparison of the original and synthesized images with the same illumination and different BRDF models. The spatially-varying Blinn-Phong model (c) overestimates the specular component and does not match the photograph as closely as the spatially-varying Torrance-Sparrow model (b). Nevertheless, it produces more accurate results compared to the spatially uniform Torrance-Sparrow BRDF (d).

Another popular choice for BRDF fitting is the Lafortune model [Lafortune et al. 1997] because it models off-specular reflections while still being computationally simple. Because the non-linear Lafortune model has more parameters than the Torrance-Sparrow and Blinn-Phong models, fitting of the model becomes more demanding, even for a single-lobe approximation. A realistic employment of the Lafortune model would use at least three lobes. This fit has been reported by [Ngan et al. 2005] to be very unstable compared to other reflectance models. In addition, [Lawrence et al. 2004] explain that "Often, a nonlinear optimizer has difficulty fitting more than two lobes of a Lafortune model without careful user

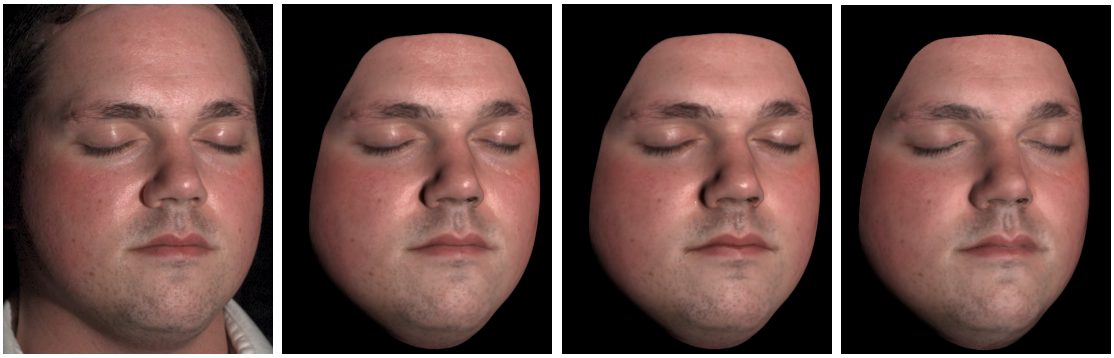

(a) (b) (c) (d)

Figure 7: *(a) Input photograph compared to our face model with (b) spatially-varying Torrance-Sparrow; (c) spatially-varying Blinn-Phong; and (d) uniform Torrance-Sparrow BRDF models.*

intervention." Given the high quality that we could achieve with the Torrance-Sparrow model, we decided not to further investigate the Lafortune model.

5.4 Measuring Translucency

Our subsurface measurement device is an image-based version of a fiber optic spectrometer with a linear array of optical fiber detectors [Nickell et al. 2000] (see Figure 8). Similarly, more sophisticated devices called Laser Doppler Flowmetry (LDF) probes are used to measure micro-circulatory blood flow and diffuse reflectance in skin [Larsson et al. 2003]. The advantage of our device over LDF or the measurement procedure in [Jensen et al. 2001] is that it can be safely used in faces. Light from a white LED is coupled to a source fiber. The alignment of the fibers is linear to minimize sensor size. A sensor head holds the source fiber and 28 detection fibers. A digital camera records the light collected by the detector fibers. The camera and detector fibers are encased in a light-proof box with air cooling to minimize imager noise. We capture 23 images bracketed by 2/3 f-stops to compute an HDR image of the detector fibers. The total acquisition time is about 88 seconds.

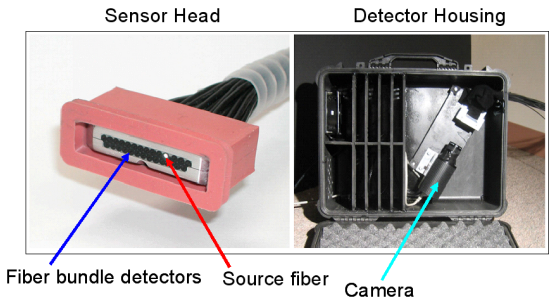

Figure 8: *Left: A picture of the sensor head with linear fiber array. The source fiber is lit. Right: The fiber array leads to a camera in a light-proof box. The box is cooled to minimize imaging sensor noise.*

Figure 9 shows the sensor head placed on a face. We have chosen to measure three points where the sensor head can be placed reliably: forehead, cheek, and below the chin. For hygienic reasons we do not measure lips. We found that pressure variations on the skin caused by the mechanical movement of the sensor head influence the results. To maintain constant pressure between skin and sensor head we attached a silicone membrane connected to a suction pump. This greatly improves the repeatability of the measurements.

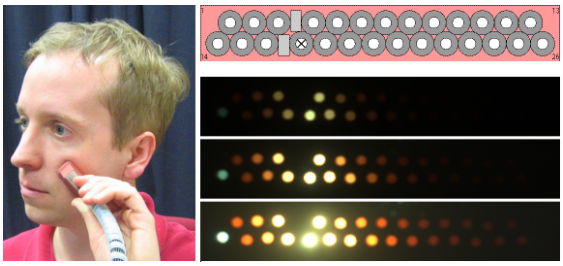

Figure 9: *Left: The sensor head placed on a face. Top: Sensor fiber layout. The source fiber is denoted by a cross. Bottom: An HDR image of the detector fibers displayed with three different exposure values.*

For more details on the subsurface device and calibration procedure see [Weyrich et al. 2005].

Previous work in diffuse reflectometry [Nickell et al. 2000] suggests that some areas of the human body exhibit anisotropic subsurface scattering (e.g., the abdomen). We measured two-dimensional subsurface scattering on the abdomen, cheek, and forehead for a few subjects. We verified the presence of significant anisotropy in the abdominal region (see Figure 10). However, the plots show that the diffuse subsurface scattering of facial skin can be well approximated with an isotropic scattering model. Consequently, we measure only a one-dimensional profile and assume rotational symmetry.

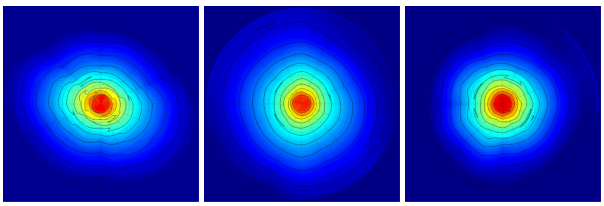

Figure 10: *Subsurface scattering curves for abdomen, cheek, and forehead measured along 16 1D profiles.*

We fit the analytic BSSRDF model of Jensen et al. [2001] to the data points of each subsurface measurement, providing us with the reduced scattering coefficient σ_s' and absorption coefficient σ_a. From the measured σ_a and σ_s' data we can derive the effective transport coefficient:

$$\sigma_{tr} = \sqrt{3\sigma_a(\sigma_a + \sigma_s')} \approx 1/\ell_d \,. \tag{15}$$

ℓ_d is the diffuse mean free path of photons (mm) and provides a measure of skin translucency.

Table 1 shows the mean and variance of σ_{tr} for a sample population of various people (see Section 7). The measurement points on

σ_{tr}	Cheek		Forehead		Neck	
(mm^{-1})	**Mean**	**Var.**	**Mean**	**Var.**	**Mean**	**Var.**
red	0.5572	0.1727	0.5443	0.0756	0.6911	0.2351
green	0.9751	0.2089	0.9831	0.1696	1.2488	0.3686
blue	1.5494	0.1881	1.5499	0.2607	1.9159	0.4230

Table 1: *Mean and variance of σ_{tr} in our dataset.*

cheek and forehead are quite similar in translucency. The measurement point on the neck underneath the chin shows a rather different mean, but also higher variance. This is probably due to measurement noise, as the sensor head is hard to place there.

Figure 11 shows closeups of the subjects with minimum $[1.0904, 0.6376, 0.6035]$ and maximum $[2.8106, 1.2607, 0.6590]$ translucency values ℓ_d in our dataset. There are subtle differences

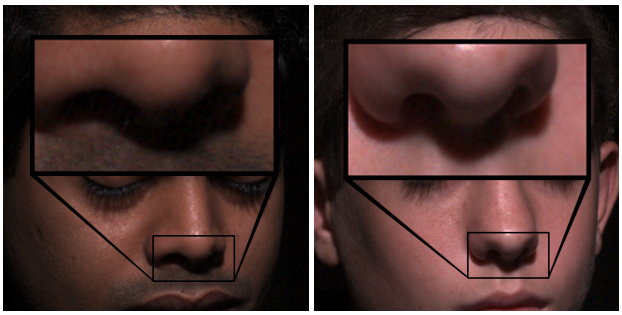

Figure 11: *Photographs of subjects with minimum (left) and maximum (right) translucency values in our dataset. The differences at shadow boundaries are subtle.*

visible at shadow boundaries. Figure 12 shows closeups computed with our model using the same minimum and maximum translucency values on one of our face models. Note that the model is

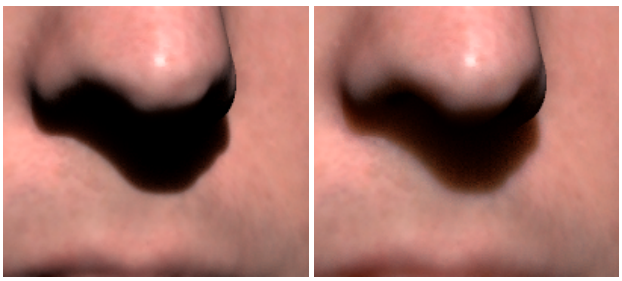

Figure 12: *Synthetic images with minimum (left) and maximum (right) translucency values.*

capable of reproducing the subtle differences of Figure 11. Translucency values ℓ_d do not vary much between measurement points and between individuals, which allows us to approximate σ_{tr} with an average value $\sigma_{trAvg} = [0.5508, 0.9791, 1.5497]$ for all subjects. This is the mean over all subjects of the measurements for the forehead and cheek, ignoring the neck.

As observed by [Larsson et al. 2003], it is difficult to reliably measure the spatially-varying absorption coefficient σ_a with a contact device. However, using σ_{trAvg} we can estimate spatially uniform scattering parameters σ_a and σ_s' from skin albedo *without* any direct measurements of subsurface scattering. We use the diffuse BRDF approximation by Jensen et al. [2001]:

$$S_{BRDF}(x_o, \omega_i, \omega_o) = \frac{\alpha'}{2\pi}\left(1 + \exp^{-\frac{4}{3}\frac{1+F_r}{1-F_r}\sqrt{3(1-\alpha')}}\right)\exp^{-\sqrt{3(1-\alpha')}}. \tag{16}$$

This approximation is valid for planar, homogeneously scattering materials under uniform directional illumination. It is quite accurate for the average total diffuse reflectance α_{avg}, which is the average of all albedo map values $\alpha(x_o)$. Using $S_{BRDF} = \alpha_{avg}$, we compute the apparent albedo α' by inverting Equation (16) using a secant root finder. We derive the model parameters σ_s' and σ_a from α' and σ_{trAvg} using [Jensen and Buhler 2002] $\sigma_s' = \alpha'\sigma_t'$ and $\sigma_a = \sigma_t' - \sigma_s'$, with $\sigma_t' = \sigma_{trAvg}/\sqrt{3(1-\alpha')}$. These diffuse scattering parameters are not spatially varying because we use α_{avg} to compute them. Instead, we capture the spatial variation of diffuse reflectance using the modulation texture $M(x_o)$.

6 Rendering

We implemented our reflectance model using a high-quality Monte Carlo offline ray tracer. The direct surface reflectance was evaluated using the Torrance-Sparrow BRDF model. The subsurface reflectance was computed using the diffusion dipole method as described in [Jensen et al. 2001] and modulated by the modulation texture $M(x_o)$. $M(x_o)$ can be derived from Equation (7) as described in Section 4. Each image took approximately five minutes to render on a single Intel Xeon 3.06GHz workstation. Figure 13 shows comparisons between real and synthetic images for different faces and different viewpoints. The synthetic and real images look very similar, but not absolutely identical. The main limitation is a lack of sharpness, both in the texture and in geometry, mainly because accuracy in geometric calibration and alignment remain issues. The intensity and shape of the specular highlights in the synthetic images is sometimes underestimated. The shape of the specular highlights – especially on the forehead – is greatly affected by fine wrinkles and small pores. Several sources of error (measured 3D geometry, motion artifacts, calibration errors, camera noise, Bayer interpolation, errors in the photometric stereo estimations, etc.) prevent us from capturing the micro-normal geometry with 100% accuracy. Much higher-resolution and higher-quality face scans can be obtained by creating a plaster mold of the face and scanning it with a high-precision laser system. For example, the XYZRGB 3D scanning process originally developed by the Canadian National Research Center contributed significantly to the realism of the Disney project [Williams 2005] and the *Matrix* sequels. However, it would be prohibitive to scan more than one hundred subjects this way. It would also be difficult to correspond the 3D geometry with the image data from the face scanning dome due to registration errors, subtle changes in expressions, or motion artifacts. In addition, the molding compound may lead to sagging of facial features [Williams 2005]. We believe that our system strikes a good tradeoff between speed, convenience, and high image quality.

Facial hair is currently not explicitly represented in our model. Although the overall reflectance and geometry of the bearded person in Figure 13 has been captured, the beard is lacking some details that are visible in the photograph. We avoided scanning people with full beards or mustaches.

We model subsurface scattering using a single layer. However, skin consists of several layers that have different scattering and absorption properties [Tuchin 2000]. The multi-layer dipole model of Donner et al. [2005] would provide a better approximation, but it is not clear how to measure its parameters *in vivo* for people of different skin types.

The dipole diffusion approximation overestimates the total fluence in the superficial areas of skin, making the skin appear slightly



Figure 13: *Comparison of real photographs (first and third row) to our model (second and last row). All photographs were cropped according to the 3D model to remove distracting features.*

Skin Type	Skin Color	Sun Exposure Reaction	Subjects (M/F)
I	Very white	Always burn	N/A
II	White	Usually burn	8 / 6
III	White to olive	Sometimes burn	49 / 18
IV	Brown	Rarely burn	40 / 8
V	Dark brown	Very rarely burn	13 / 2
VI	Black	Never burn	4 / 1

Table 2: *The Fitzpatrick skin type system and the number of subjects per skin type.*

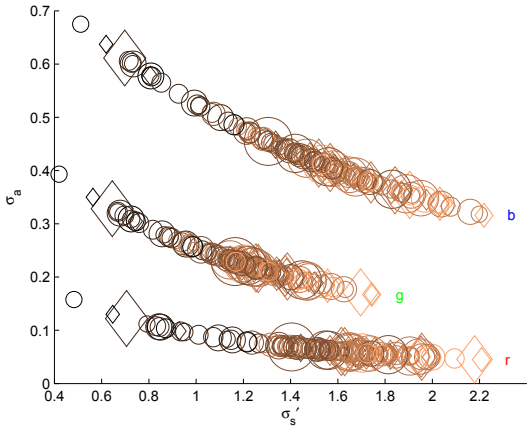

Figure 14: *Plot of σ_a [mm] and σ_s' [mm] for all subjects. Color encodes skin type (light = II, black = VI); shape encodes gender ($\diamond$ = F, $\circ$ = M); marker size encodes age (larger = older). Skin absorption is highest for skin with low scattering and decreases for more highly scattering skin.*

more opaque than actual skin [Igarashi et al. 2005]. In addition, our measurements are confined to RGB data. For example, the parameters σ_a and σ_s' for skin are wavelength dependent [Igarashi et al. 2005]. Krishnaswamy and Baranoski [2004] developed a multispectral skin model and emphasize the importance of spectral rendering in the context of skin.

7 Face Reflectance Analysis

We measured face reflectance for 149 subjects (114 male / 35 female) of different age and race. The distribution by age is (M/F): 20s or younger (61/18), 30s (37/10), 40s (10/2), 50s or older (6/5). Because race is an ambiguous term (e.g., some Caucasians are darker than others), we classify the data by skin type according to the Fitzpatrick system [Fitzpatrick 1988]. Since we measured only two individuals of skin type I, we group types I and II together. Table 2 lists the number of subjects per skin type.

Analysis of Absorption and Scattering Parameters: We already discussed the measured subsurface scattering data in Section 5.4. The absorption coefficient σ_a and the reduced scattering coefficient σ_s' represent the amounts of light attenuation caused by absorption and multiple scattering, respectively. We estimate these parameters from the average albedo α_{avg} and the average translucency σ_{trAvg}. Figure 14 shows the distribution of absorption and reduced scattering values for all subjects. Skin type is encoded by color (light = II, black = VI), gender by marker type ($\diamond$ = F, $\circ$ = M), and age by marker size (larger = older). Subjects with skin type V and VI have higher absorption and lower scattering coefficients than subjects with skin type II or III. There seems to be no correlation with age. As expected, the absorption coefficient σ_a is highest

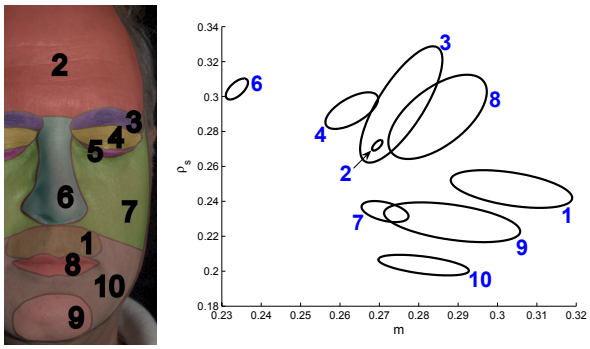

Figure 15: *Left: 10 face regions. Right: Variation of Torrance-Sparrow parameters per face region averaged over all subjects. The center of the ellipse indicates the mean. The axis of each ellipse shows the directions of most variation based on a PCA.*

for the blue and green channels. In addition, the range of σ_s' is much smaller for the green channel. These effects are due to the absorption of blue and green light by blood in the deeper levels of the skin. The average absorption σ_a decreases in people with higher scattering coefficient σ_s'. The relationship is near-exponential since we are using the diffuse BRDF approximation by Jensen et al. (Equation (16)).

Analysis of Torrance-Sparrow BRDF Parameters: To study the variations of the specular reflectance we analyzed the physically-based Torrance-Sparrow model parameters. Similar to the work by Fuchs et al. [Fuchs et al. 2005b], we distinguish the 10 different face regions shown in Figure 15 (left). For each region, we randomly chose a few thousand vertices by marking them manually. We collected the BRDF parameters for the vertices and averaged them per region. Figure 15 (right) visualizes the resulting parameter distributions, averaged over all subjects. To create this visualization, we computed a $2 \times n$ matrix for each region, where n is the total number of samples in that region for all subjects. We then performed a Principal Component Analysis (PCA) on each matrix. The resulting eigenvectors and eigenvalues are the directions and size of the ellipse axis, respectively. The centers of the ellipses have been translated to the mean values of the BRDF parameters for that region.

As shown in Figure 15, the vermilion zone (or red zone) of the lips (8) exhibits a large parameter variation. Lips have smooth, glabrous skin, i.e., they have no follicles or small hairs. The epidermis is thin, almost transparent, and the blood vessels of the dermis are close to the surface, giving raise to the red appearance of the lips. The reflectance variations are most likely due to circulation differences in these blood vessels and the highly folded dermis, which results in the ridged appearance of the lips. Big parameter variations can also be observed in regions 1 (mustache), 3 (eyebrows), 9 (chin), and 10 (beard). Male subjects contain facial hair in these regions, which leads to noisy normal estimations and widely varying parameter distributions. Smoother facial areas, such as the forehead (2), eyelids (4), nose (6), and cheeks (7), show much smaller parameter variance. The highest specular scaling coefficients ρ_s appear for nose (6), eyebrows (3), eyelids (4), lips (8), and forehead (2). The nose and forehead especially are areas of high sebum secretion, making them appear more specular. Not surprisingly, the nose (6) shows the smallest values for the roughness parameter m because it lacks small facial wrinkles.

To analyze how the Torrance-Sparrow parameters vary with skin type, age, and gender, we computed the average of the BRDF parameters in each region and stored the data for all subjects in a 23×149 matrix M. The rows of the matrix contain the average BRDF parameters for each face region, the gender $[0 = M, 1 = F]$,

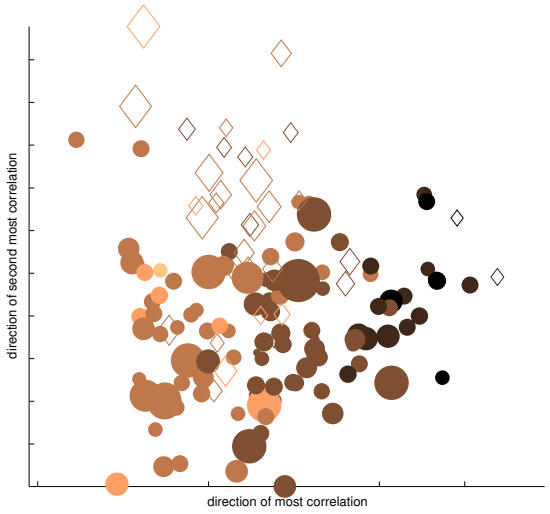

Figure 16: *Canonical correlation analysis of the Torrance-Sparrow BRDF parameters. Color encodes skin type (light = II, black = VI); shape encodes gender (◇ = F, ○ = M); marker size encodes age (larger = older). Skin type is the only reliable predictor of BRDF parameters.*

the age, and the skin type $[II - VI]$. The columns contain the data for each person. To visualize the matrix M we tried linear (PCA) and non-linear dimensionality reduction techniques. However, the data is tightly clustered in a high-dimensional ball that cannot be reduced to fewer dimensions. Instead, we use Canonical Correlation Analysis (CCA) [Green et al. 1966]. CCA finds projections of two multivariate datasets that are maximally correlated. Figure 16 shows the output of CCA for the Torrance-Sparrow model. Here the data is projected (orthogonally) onto the two dimensions that are most strongly correlated with the experimental variables of skin type, gender, and age. The first dimension is best predicted by skin type (color), the second by gender (◇ = F, ○ = M). Age is encoded in marker size. The principal correlations are 0.768 and 0.585 (to mixtures of the experimental variables); the direct correlations are 0.766 for skin type, 0.573 for gender, and 0.537 for age. This indicates that CCA was not able to greatly improve the direct correlations. Maybe not surprisingly, skin type is the only reliable predictor for BRDF parameters.

To analyze variations of BRDF parameters due to external factors, we also measured one male subject under several external conditions (normal, cold, hot, sweaty, with makeup, with lotion, and with powder). As expected, there are noticeable differences between these BRDFs, especially between lotion / hot and cold / powder.

Dissemination: The statistics of our face reflectance analysis of all subjects will be made public to the research community. We are releasing all reflectance data, such as BRDF coefficients, absorption and scattering coefficients, and albedo histograms, for each subject. In addition, we are providing an interactive applet that outputs average parameter values and histograms based on user choice for gender, age, and skin type. This data will allow image synthesis, editing, and analysis of faces without having to acquire and process a large dataset of subjects.

8 Face Editing

Our face reflectance model has intuitive parameters for editing, such as BRDF parameters, average skin color, or detailed skin texture. To change the surface reflectance of a model, we can either edit the parameters of the analytic BRDF model directly, or we can adjust them using the statistics from our database. For each face region, we can match the parameter histogram of the source face with another parameter histogram using the histogram interpolation technique presented in [Matusik et al. 2005]. This allows us to change a face BRDF based on, for example, different external conditions (e.g., from normal to lotion) (see Figure 17, bottom). We can also transfer skin color histograms from one subject to another. This allows us to change, for example, the skin type (e.g., from II to V) (see Figure 1, right).

To transfer skin texture, we use the texture analysis technique of Heeger and Bergen [1995]. We first transform the albedo map of each subject into a common and decorrelated color space using the method discussed in [Heeger and Bergen 1995][Section 3.5] and process each transformed color channel independently. We then compute statistics (histograms) of the original albedo map at full resolution, and of filter responses at different orientations and scales organized as a steerable pyramid. We use seven pyramid levels with four oriented filters, and down-sample the albedo map by a factor of two at each level. Each histogram has 256 bins. The first level of the pyramid is a low-pass and high-pass filter computed on the full-resolution albedo map. The output of the low-pass filter is a histogram of average skin color.

The histograms of all pyramid filter responses (30 total) and the histogram of the original albedo map are concatenated into a $256 \times 31 \times 3 = 23,808$ element vector H. To transfer skin textures, we use the histogram vector H of a source and a target face and apply the histogram-matching technique of Heeger and Bergen [1995]. We can either use the histograms of the whole albedo map or restrict the analysis/synthesis to certain facial regions. Note that Heeger and Bergen start their texture synthesis with a noise image as the source. In our case it makes more sense to start from the source albedo map. To allow for sufficient variation during histogram matching, we add some noise to the source albedo map before we compute its histogram vector H. A similar method was used by Tsumura et al. [2003] for synthesis of 2D skin textures, and by Cula and Dana [2002] to analyze bi-directional texture functions (BTFs) of skin patches. Figure 17 shows transfers of skin texture, BRDF, and skin color for two face models. The translucency values of the source and result models remain the same. The example in Figure 1 (right) shows skin albedo and BRDF transfers. The resulting face appears lighter and shinier.

9 Conclusions and Future Work

In this paper we have proposed a simple and practical skin model. An important feature of our model is that all its parameters can be robustly estimated from measurements. This reduces the large amount of measured data to a manageable size, facilitates editing, and enables face appearance changes. Images from our model come close to reproducing photographs of real faces for arbitrary illumination and pose. We fit our model to data of a large and diverse group of people. The analysis of this data provides insight into the variance of face reflectance parameters based on age, gender, or skin type. We are making the database with all statistics available to the research community for face synthesis and analysis.

As discussed in Section 6, there is room for improvement in our model. For example, it would be interesting to measure wavelength-dependent absorption and scattering parameters. It would also be interesting to compare the results from the diffuse dipole approximation with a full Monte Carlo subsurface scattering simulation. Other important areas that require a different modeling approach are facial hair (eyebrows, eyelashes, mustaches, and beards), hair, ears, eyes, and teeth. Very fine facial hair also leads to asperity scattering and the important "velvet" look of skin near

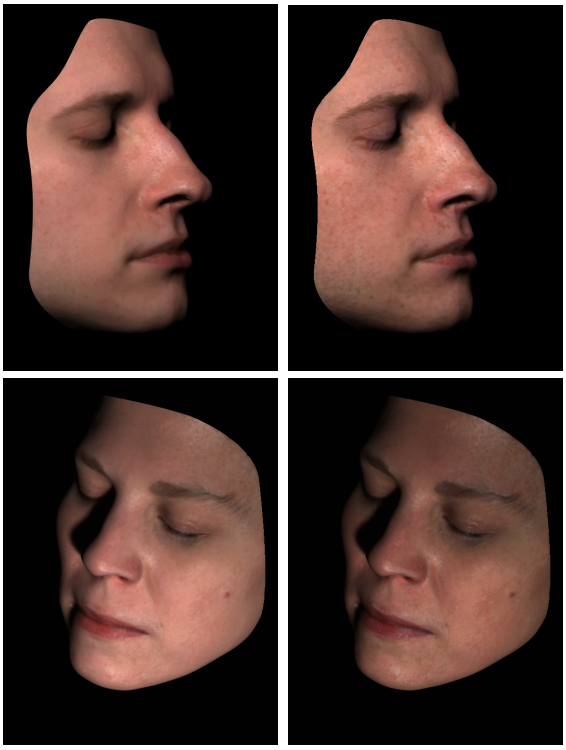

Figure 17: *Face appearance changes. Top: Transferring the freckled skin texture of the subject in Figure 1 (left). Bottom: Changing the skin type (BRDF and albedo) from type III to type IV.*

grazing angles [Koenderink and Pont 2003]. Our model does not take this into account.

We captured face reflectance on static, neutral faces. Equally important are expressions and face performances. For example, it is well known that the blood flow in skin changes based on facial expressions. Our setup has the advantage that such reflectance changes could be captured in real-time.

10 Acknowledgments

We would like to thank Diego Nehab and Bruno Lévy for making their geometry processing software available to us, Agata Lukasinska and Natalie Charton for processing some of the face data, and Jennifer Roderick Pfister for proofreading the paper. John Barnwell, William "Crash" Yerazunis, and Freddy Bürki helped with the construction of the measurement hardware, and Cliff Forlines produced the video. Many thanks also to Joe Marks for his continuing support of this project. Henrik Wann Jensen and Craig Donner were supported by a Sloan Fellowship and the National Science Foundation under Grant No. 0305399.

References

ANGELOPOULOU, E., MOLANA, R., AND DANIILIDIS, K. 2001. Multispectral skin color modeling. In *IEEE Conf. on Comp. Vision and Pattern Rec. (CVPR)*, 635–642.

BARSKY, S., AND PETROU, M. 2001. Colour photometric stereo: simultaneous reconstruction of local gradient and colour of rough textured surfaces. In *Eighth IEEE International Conference on Computer Vision*, vol. 2, 600–605.

BLANZ, V., AND VETTER, T. 1999. A morphable model for the synthesis of 3D faces. *Computer Graphics 33*, Annual Conference Series, 187–194.

BLINN, J. 1977. Models of light reflection for computer synthesized pictures. *Computer Graphics 11*, Annual Conference Series, 192–198.

BORSHUKOV, G., AND LEWIS, J. 2003. Realistic human face rendering for the matix reloaded. In *ACM SIGGRAPH 2003 Conference Abstracts and Applications (Sketch)*.

CULA, O., AND DANA, K. 2002. Image-based skin analysis. In *Texture 2002*, The Second International Workshop on Texture Analysis and Synthesis, 35–42.

CULA, O., DANA, K., MURPHY, F., AND RAO, B. 2004. Bidirectional imaging and modeling of skin texture. *IEEE Transactions on Biomedical Engineering 51*, 12 (Dec.), 2148–2159.

CULA, O., DANA, K., MURPHY, F., AND RAO, B. 2005. Skin texture modeling. *International Journal of Computer Vision 62*, 1–2 (April–May), 97–119.

DANA, K. J., VAN GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and texture of real world surfaces. *ACM Transactions on Graphics 1*, 18, 1–34.

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

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

DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *Computer Graphics*, SIGGRAPH 2000 Proceedings, 145–156.

DEBEVEC, P., WENGER, A., TCHOU, C., GARDNER, A., WAESE, J., AND HAWKINS, T. 2002. A lighting reproduction approach to live-action compositing. *ACM Transactions on Graphics (SIGGRAPH 2002) 21*, 3 (July), 547–556.

DONNER, C., AND JENSEN, H. W. 2005. Light diffusion in multi-layered translucent materials. *ACM Transactions on Graphics 24*, 3 (July), 1032–1039.

FITZPATRICK, T. 1988. The validity and practicality of sun-reactive skin types i through vi. *Arch. Dermatology 124*, 6 (June), 869–871.

FUCHS, C., GOESELE, M., CHEN, T., AND SEIDEL, H.-P. 2005. An empirical model for heterogeneous translucent objects. Tech. Rep. MPI-I-2005-4-006, Max-Planck Institute für Informatik, May.

FUCHS, M., , BLANZ, V., LENSCH, H., AND SEIDEL, H.-P. 2005. Reflectance from images: A model-based approach for human faces. Research Report MPI-I-2005-4-001, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany. Accepted for publication in IEEE TVCG.

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., LANG, J., FUCHS, C., AND SEIDEL, H. 2004. Disco – acquisition of translucent objects. *ACM Transactions on Graphics 24*, 3, 835–844.

GREEN, P., HALBERT, M., AND ROBINSON, P. 1966. Canonical analysis: An exposition and illustrative application. *Marketing Research 3* (Feb.), 32–39.

HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. In *Computer Graphics*, SIGGRAPH 93 Proceedings, 165–174.

HAWKINS, T., WENGER, A., TCHOU, C., GORANSSON, F., AND DEBEVEC, P. 2004. Animatable facial reflectance fields. In *Rendering Techniques '04 (Proceedings of the Second Eurographics Symposium on Rendering)*.

HEEGER, D., AND BERGEN, J. 1995. Pyramid-based texture analysis/synthesis. In *Proceedings of SIGGRAPH 95*, Computer Graphics Proceedings, Annual Conference Series, 229–238.

IGARASHI, T., NISHINO, K., AND NAYAR, S. 2005. The appearance of human skin. Tech. Rep. CUCS-024-05, Department of Computer Science, Columbia University, June.

JENSEN, H. W., AND BUHLER, J. 2002. A rapid hierarchical rendering technique for translucent materials. In *Computer Graphics*, SIGGRAPH 2002 Proceedings, 576–581.

JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Computer Graphics*, SIGGRAPH 2001 Proceedings, 511–518.

JONES, A., GARDNER, A., BOLAS, M., MCDOWALL, I., AND DEBEVEC, P. 2005. Performance geometry capture for spatially varying relighting. In *ACM SIGGRAPH 2005 Conference Abstracts and Applications (Sketch)*.

KOENDERINK, J., AND PONT, S. 2003. The secret of velvety skin. *Machine Vision and Application*, 14, 260–268. Special Issue on Human Modeling, Analysis, and Synthesis.

KRISHNASWAMY, A., AND BARANOSKI, G. 2004. A biophysically-based spectral model of light interaction with human skin. *Computer Graphics Forum 23*, 3 (Sept.), 331–340.

LAFORTUNE, E., FOO, S.-C., TORRANCE, K., AND GREENBERG, D. 1997. Non-linear approximation of reflectance functions. *Computer Graphics 31*, Annual Conference Series, 117–126.

LARSSON, M., NILSSON, H., AND STRÖMBERG, T. 2003. In vivo determination of local skin optical properties and photon path length by use of spatially resolved diffuse reflectance with applications in laser doppler flowmetry. *Applied Optics 42*, 1 (Jan.), 124–134.

LAWRENCE, J., RUSINKIEWICZ, S., AND RAMAMOORTHI, R. 2004. Efficient brdf importance sampling using a factored representation. *ACM Transactions on Graphics 23*, 3, 496–505.

LENSCH, H., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2001. Image-based reconstruction of spatially varying materials. In *Proceedings of the 12th Eurographics Workshop on Rendering*, 104–115.

LOOP, C. 1987. *Smooth Subdivision Surfaces based on Triangles*. Master's thesis, Department of Mathematics, University of Utah.

MARSCHNER, S., WESTIN, S., LAFORTUNE, E., TORRANCE, K., AND GREENBERG, D. 1999. Image-based BRDF measurement including human skin. In *Proceedings of the 10th Eurographics Workshop on Rendering*, 139–152.

MARSCHNER, S., GUNETER, B., AND RAGHUPATHY, S. 2000. Modeling and rendering for realistic facial animation. In *11th Eurographics Workshop on Rendering*, 231–242.

MARSCHNER, S., WESTIN, S., LAFORTUNE, E., AND TORRANCE, K. 2000. Image-based measurement of the Bidirectional Reflectance Distribution Function. *Applied Optics 39*, 16 (June), 2592–2600.

MATUSIK, W., ZWICKER, M., AND DURAND, F. 2005. Texture design using a simplicial complex of morphable textures. *ACM Transactions on Graphics 24*, 3, 787–794.

NEHAB, D., RUSINKIEWICZ, S., DAVIS, J., AND RAMAMOORTHI, R. 2005. Efficiently combining positions and normals for precise 3d geometry. *ACM Transactions on Graphics 24*, 3, 536–543.

NG, C., AND LI, L. 2001. A multi-layered reflection model of natural human skin. In *Computer Graphics International*, 249256.

NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of brdf models. In *Eurographics Symposium on Rendering*, 117–126.

NICKELL, S., HERMANN, M., ESSENPREIS, M., FARRELL, T. J., KRAMER, U., AND PATTERSON, M. S. 2000. Anisotropy of light propagation in human skin. *Phys. Med. Biol. 45*, 2873–2886.

NICODEMUS, F., RICHMOND, J., HSIA, J., GINSBERG, I., AND LIMPERIS, T. 1977. Geometric considerations and nomenclature for reflectance. Monograph 160, National Bureau of Standards (US), October.

PARIS, S., SILLION, F., AND QUAN, L. 2003. Lightweight face relighting. In *Proceedings of Pacific Graphics*, 41–50.

PIGHIN, F., HECKER, J., LISCHINSKI, D., SZELISKI, R., AND SALESIN, D. 1998. Synthesizing realistic facial expressions from photographs. In *Computer Graphics*, vol. 32 of *SIGGRAPH 98 Proceedings*, 75–84.

SANDER, P., GOSSELIN, D., AND MITCHELL, J. 2004. Real-time skin rendering on graphics hardware. SIGGRAPH 2004 Sketch.

SHEFFER, A., LEVY, B., MOGILNITSKY, M., AND BOGOMYAKOV, A. 2005. ABF++: Fast and robust angle based flattening. *ACM Transactions on Graphics 24*, 2, 311–330.

STAM, J. 2001. An illumination model for a skin layer bounded by rough surfaces. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, Springer, Wien, Vienna, Austria, 39–52.

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, 1054–1061.

TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America 57*, 1104–1114.

TSUMURA, N., OJIMA, N., SATO, K., SHIRAISHI, M., SHIMIZU, H., NABESHIMA, H., AKAZAKI, S., HORI, K., AND MIYAKE, Y. 2003. Image-based skin color and texture analysis/synthesis by extracting hemoglobin and melanin information in the skin. *ACM Transactions on Graphics (SIGGRAPH 2003) 22*, 3, 770–779.

TUCHIN, V. 2000. *Tissue Optics: Light Scattering Methods and Instruments for Medical Diagnosis*. SPIE Press.

WENGER, A., GARDNER, A., TCHOU, C., UNGER, J., HAWKINS, T., AND DEBEVEC, P. 2005. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Transactios on Graphics 24*, 3, 756–764.

WEYRICH, T., MATUSIK, W., PFISTER, H., NGAN, A., AND GROSS, M. 2005. Measuring skin reflectance and subsurface scattering. Tech. Rep. TR2005-046, Mitsubishi Electric Research Laboratories (MERL).

WILLIAMS, L. 2005. Case study: The gemini man. In *SIGGRAPH 2005 Course 'Digital Face Cloning'*.