A Measurement-Based Skin Reflectance Model for Face

Tim Weyrich, Wojciech Matusik, Hanspeter Pfister, Jinho Lee, Addy Ngan, Henrik Wann Jensen, Markus Gross
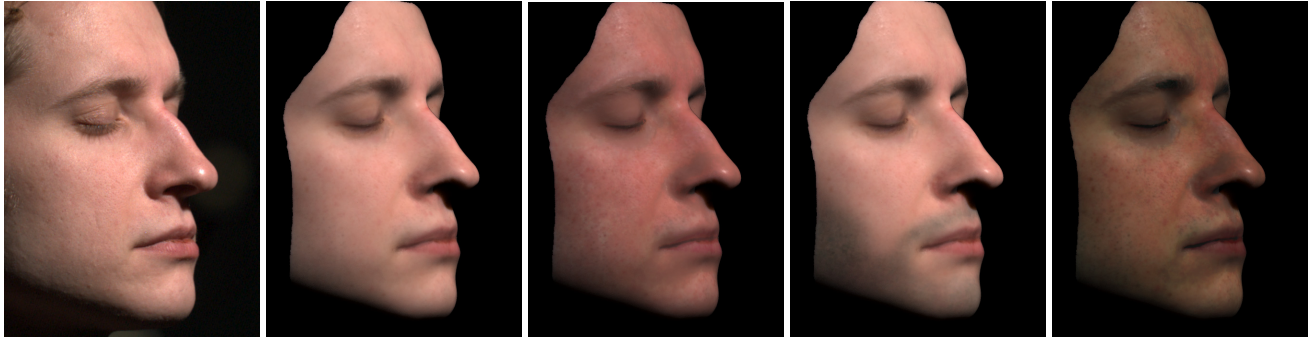
Abstract

We present a novel skin reflectance model for faces and its application to face appearance editing. We decompose the high-dimensional bidirectional scattering surface reflectance distribution function (BSSRDF) of skin into components that can be estimated from measured data. Our model is intuitive, amenable to interactive rendering, and easy to edit. High-quality renderings come close to reproducing real photographs. We have measured 3D face geometry, skin reflectance, and subsurface scattering for a large group of people using custom-built devices and fit the data to our model. The analysis of the reflectance data reveals variations according to subject age, race, gender, and external factors (heat, cold, makeup, etc.) We derive a low-dimensional model using non-negative matrix factorization (NMF) that spans the spaace of skin reflectance in our database. A user can define meaningful parameters in this space - such as race, gender, and age - and change the overall appearance of a person (e.g. making a Caucasian face look more Asian) or change local features (e.g., adding moles, freckles, or hair follicles).

A Measurement-Based Skin Reflectance Model for Face Rendering and Editing

Tim Weyrich [*]   Wojciech Matusik [†]   Hanspeter Pfister [†]   Jinho Lee [†]   Addy Ngan [‡]   Henrik Wann Jensen [§]

Markus Gross [*]

Figure 1: *A face rendered and edited using our skin reflectance model. From left to right: Real photograph; rendering; making the face more tanned; adding hair follicles in the beard area; making the skin darker.*

Abstract

We present a novel skin reflectance model for faces and its application to face appearance editing. We decompose the high-dimensional bidirectional scattering surface reflectance distribution function (BSSRDF) of skin into components that can be estimated from measured data. Our model is intuitive, amenable to interactive rendering, and easy to edit. High-quality renderings come close to reproducing real photographs. We have measured 3D face geometry, skin reflectance, and subsurface scattering for a large group of people using custom-built devices and fit the data to our model. The analysis of the reflectance data reveals variations according to subject age, race, gender, and external factors (heat, cold, makeup, etc.) We derive a low-dimensional model using non-negative matrix factorization (NMF) that spans the space of skin reflectance in our database. A user can define meaningful parameters in this space – such as race, gender, and age – and change the overall appearance of a person (e.g., making a Caucasian face look more Asian) or change local features (e.g., adding moles, freckles, or hair follicles).

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

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

1 Introduction

One of the most difficult computer graphics challenges is creating realistic human faces. Humans have evolved to be incredibly adept at interpreting facial appearance. For example, we can easily distinguish if a person is tired, hot, excited, or sick. Although a lot of effort has been devoted to face modeling in computer graphics, no synthetic face model to date achieves this level of expressiveness and realism.

In this paper, we focus on modeling skin reflectance of human faces, an important aspect of face appearance. It varies for different people (e.g., due to race or gender) and even varies for the same person throughout the course of a day (e.g., hot vs. cold skin). A realistic skin reflectance model should be able to accommodate these variations. It should also allow a graphic artist to change the appearance of skin based on easy to interpret parameters (e.g., race, gender, or age). The model needs to easily connect to measurements of real faces for the creation of virtual doubles. Images generated from the model – ideally in real-time – need to look photorealistic from arbitrary viewpoints. And the model should allow easy modification or transfer of skin appearance.

To achieve these goals we have developed a novel skin reflectance model whose components can be robustly estimated from measured data. Our model is accurate, compact, and intuitive to edit. It can be used in interactive and offline rendering systems and generates results that come close to reproducing real photographs. We use custom-built devices to measure in-vivo light reflection and subsurface scattering of a large and diverse group of people. [1] Our data ranges across age (13 to 74 years old), gender, race, and external factors (e.g., cosmetics, cold, and sweat). We fit our model to the measured data and compute a low-dimensional face reflectance space using non-negative matrix factorization (NMF) [Lee and Seung 1999]. User-defined parameters – such as gender, race, or tan – allow us guide the interpolation of reflectance data to change the appearance of a face overall or locally.

---

[1]At the time of submission we measured over 85 subjects. The face-scanning project is ongoing and the database continues to grow.

2 Previous Work

Properties of human skin have been measured and studied extensively in the biomedical, cosmetics, and computer vision communities. In this section we provide an overview of the relevant work in the area of computer graphics and image synthesis.

Analytic Skin Reflectance Models: Analytic reflectance models are attractive because of their computational efficiency. Hanrahan and Krueger [1993] modeled single-scattering of light in skin composed of multiple smoothly-bounded internal layers. Ng and Li [2001] extended this model by adding an oil layer to the skin surface. Stam [2001] developed an analytic approximation to multiple subsurface scattering in skin with a rough surface. More recent work [Krishnaswamy and Baranoski 2004] proposes a biophysically-based multi-layer model for image synthesis with biologically meaningful parameters.

Several skin modeling approaches use analytic bidirectional surface reflectance functions (BRDFs) [Blanz and Vetter 1999; Debevec et al. 2000; Haro et al. 2001; Paris et al. 2003; Tsumura et al. 2003; Fuchs et al. 2005]. The BRDF parameters can be estimated from reflectance measurements using non-linear optimization. Although a BRDF describes local light transport at each surface point, it ignores subsurface scattering, which is largely responsible for the appearance of skin.

Jensen et al. [2001; 2002] propose an analytic model for the bidirectional surface-scattering distribution function (BSSRDF). The BSSRDF describes the full effect that incident light at a point has on the reflected light from a surface patch around that point. The BSSRDF is eight-dimensional, assuming a two-dimensional parameterization of the surface. Because dense sampling of an eight-dimensional function is challenging, we subdivide the BSSRDF into components that can be more easily measured (see Section 3).

Non-parametric Skin Reflectance Models: Instead of fitting an analytic BRDF model, Marschner et al. [1999] estimate a non-parametric BRDF of a human face by combining reflectance samples from different points on the surface. They later extended this work by adding a detailed albedo texture [Marschner et al. 2000]. They observe that the BRDF of skin is quite unusual and exhibits strong forward scattering at grazing angles that is uncorrelated with the specular direction.

We use the data-driven BRDF model of Matusik et al. [Matusik et al. 2003] to estimate a non-parametric surface BRDF at each surface point. We found that this introduces less error than imposing the behavior of a particular analytic BRDF model. More importantly, it does not require non-linear optimization and leads to a more robust fitting procedure.

Image-based Face Modeling: Image-based methods have provided highly realistic representations for human faces. They easily capture effects such as self-shadowing, inter-reflections, and sub-surface scattering [Pighin et al. 1998]. Recent efforts allow variations in lighting [Georghiades et al. 1999; Debevec et al. 2000], viewpoint, and expression [Hawkins et al. 2004]. Cula et al. [2005; 2004] collected a database containing more than 3500 skin texture images that were taken under various illumination and viewing conditions [Rutgers]. However, the memory requirements for image-based models are large. The measurement procedures are inefficient and assume non-local low-frequency lighting. Pure image-based representations are also inherently difficult to edit and modify.

Borshukov and Lewis [2003] combine an image-based model, an analytic surface BRDF, and an approximation of subsurface scattering to create highly realistic face images for the movie *Matrix Reloaded*. Sander et al. [2004] developed a variant of this method for real-time skin rendering on modern graphics hardware.

An interesting image-based method was presented by Tsumura et al. [2003], who use independent component analysis (ICA) to decompose images of faces into layers (melanin and hemoglobin). Their method is capable of re-synthesizing new images while adding effects like tanning or aging.

3 Skin Reflectance Model

Overall skin reflectance can be described as the sum of specular reflection on the skin surface (air-oil interface) and diffuse reflection due to subsurface scattering (see Figure 2). Diffuse subsurface
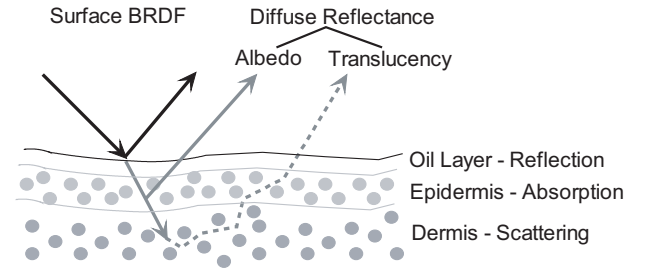
Figure 2: *Skin reflectance can be explained by a specular (BRDF) component at the air-oil interface, and a diffuse reflectance component due to subsurface scattering. Most of the high-frequency spatial color variation in human skin is due to the epidermal layer, whereas strong light scattering in the dermal layer is a more slowly varying effect. We model the first (high-frequency) effect with an albedo map and the second (low-frequency) light transport with a translucency map.*

scattering is due to absorption and light scattering in the epidermal and dermal skin layers. The epidermis scatters light strongly and contains melanin (along the interface to the dermis layer), which is highly absorbing. This absorption component is a local effect with high spatial variation across the face due to hair follicles, sweat glands, freckles, dimples, etc. The dermis/blood layer is highly scattering in the red channel and strongly absorbing in the green and blue channels (mainly due to haemoglobin). The dermal light scattering is a non-local, slowly varying effect.

We model the light that is immediately reflected from the oil-skin layer with a spatially-varying *surface BRDF* and divide diffuse subsurface reflectance into two components: A diffuse *albedo map* that captures high-frequency color variations due to epidermal absorption and scattering, and a *translucency map* that captures low-frequency absorption and scattering in the dermal layer. Fine-scale face geometry is represented by a normal map.

More formally, we denote the BSSRDF as $S(x_i, \omega_i, x_o, \omega_o)$, 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. Similarly, we use $f_s(x_i, \omega_i, \omega_o)$ for the surface BRDF. The relative contributions of the surface BRDF and the diffuse reflectance due to subsurface scattering are modulated by Fresnel coefficients:

$$S_{skin} = F_t(\eta, \omega_i) S(x_i, \omega_i, x_o, \omega_o) F_t(\eta, \omega_o), \quad (1)$$
$$f_{skin} = F_r(\eta, \omega_i) f_s, \quad (2)$$

where F_r and $F_t = (1 - F_r)$ are the Fresnel coefficient at the air-skin boundary for both the incoming and outgoing radiance, and η is the relative index of refraction between skin and air (≈ 1.3). We model the subsurface scattering term S using the dipole diffusion approximation [Jensen et al. 2001], while the specular BRDF component f_s is modeled using a data-driven approach. The parameters of both the BSSRDF and the BRDF are estimated from measurements as described in the following sections.

4 Measurement Procedure Overview

A block diagram of our measurement pipeline is shown in Figure 3. We capture the 3D geometry of the face using a commercial 3D
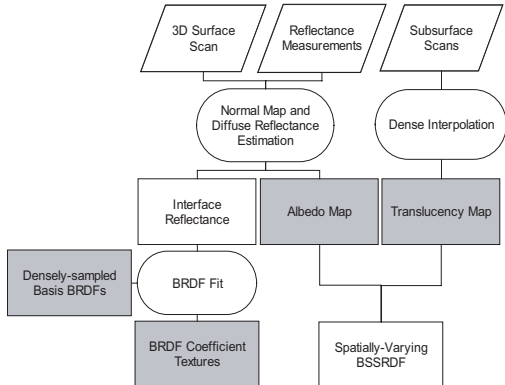
Figure 3: *A block diagram of our data processing pipeline. Blocks in grey are the parameters of our skin reflectance model.*

face scanner. Digital photographs from different viewpoints and with different illumination directions are taken in a calibrated face-scan dome. The data is used to compute a normal map and to estimate the diffuse reflectance at each surface point. We subtract the diffuse reflectance from the measured data and fit a set of densely-measured BRDFs to the remaining surface reflectance. We compress the BRDF basis using NMF to derive a small set of NMF basis BRDFs. We then measure the subsurface scattering of skin at few locations in the face using a special contact device and estimate skin translucency.

To map between 3D face space and texture space we use the area-preserving texture parameterization of Desbrun et al. [2002]. The data is densely interpolated using push-pull interpolation into texture maps of 2048×2048 resolution. The parameters of our reflectance model are the NMF basis BRDFs (typically four), textures with coefficients for the linear combination of basis BRDFs, one albedo map with diffuse reflectance values, and one translucency map. The following sections describe each of these processing steps in more detail.

5 Measuring Skin Reflectance

Figure 4 shows a photograph of our face-scanning dome. The subject sits in a chair with a head rest 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 each light on 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 [Anonymous 2005].[2]

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

---

[2]This technical report has been submitted as supplemental material.
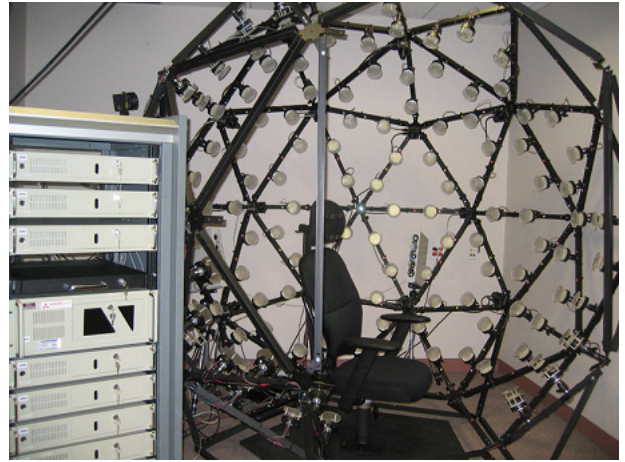
Figure 4: *The face-scanning dome consists of 16 digital cameras, 150 LED light sources, and a commercial 3D face-scanning system.*

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 500,000 to 1 million vertices. The subdivision implicitly removes noise. We store the high-resolution mesh as an unstructured list of point samples (surfels) without connectivity. Each surfel stores the necessary information for image reconstruction using EWA splatting [Zwicker et al. 2002].

Next, we compute a lumitexel [Lensch et al. 2001] at each surfel position from the image reflectance samples. Each observed radiance value $L(\omega_o)$ is normalized by the irradiance, $E_i(\omega_i)$, of the corresponding light source l in order to obtain a BRDF sample value:

$$f_r(\omega_i, \omega_o) = \frac{L(\omega_o)}{E_i(\omega_i)} \tag{3}$$

We calibrated the BRDF measurements using Fluorilon — a material with know properties [Anonymous 2005]. All processing is performed on RGB data except where noted otherwise.

We determine lumitexels in shadow areas by rendering the face from camera and light source viewpoints using EWA splatting [Zwicker et al. 2001]. EWA filtering guarantees that each surfel splat covers at least one pixel. The percentage of visible pixels per splat is approximated by the ratio of total rasterized pixels to effectively blended pixels. If this percentage is above a threshold we mark the surfel as visible from that camera or light source. This process of determining visibility and occlusion is similar to shadow mapping [Williams 1978], but using EWA splatting.

Figure 5 shows a visualization of lumitexels from two different points in the same face. The images give an impression of the hemispherical sampling for each camera viewpoint. On average, a lumitexel contains about 900 reflectance samples per color channel, with many lumitexels having up to 1,200 samples. The numbers vary depending on the occlusions across the face (see Figure 5). In contrast to previous work [Marschner et al. 1999; Lensch et al. 2001; Fuchs et al. 2005], we collect enough samples for a reliable BRDF fit at almost all lumitexels without clustering. The data for lumitexels with a badly conditioned fit is interpolated during creation of the texture maps.

6 Estimating Normals and Diffuse Albedo

We estimate normals at each lumitexel from the reflectance data. For each camera viewpoint, we determine the direction of maximum reflectance by interpolation between the directions of the four
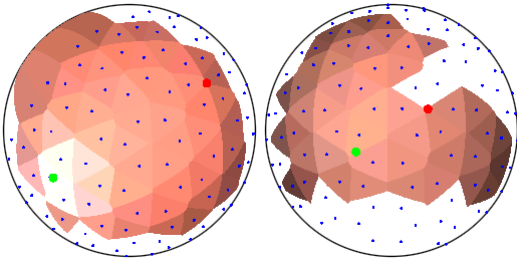
Figure 5: *Visualization of two skin lumitexels for different viewpoints. Only the samples for a single camera are shown. The red and green dots are the intersection of the view vector and reflection vector with the hemisphere, respectively. Reflectance values for each light position (blue dots) are used to shade the hemispherical triangulation. Note that the reflection direction corresponds well with the specular highlight on the left. The right lumitexel is more diffuse, and some light positions are occluded by the nose.*
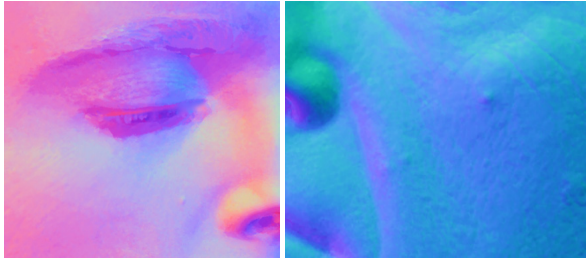
Figure 6: *Facial detail shown in closeups of two normal maps.*

brightest measurements. The half-way vector between this direction and the viewing vector is the normal estimate for this viewpoint. We then remove outliers from the set of normals from all viewpoints and interpolate the final surfel normal. This normal estimation procedure is stable and leads to good results considering the precision of our measurements (see Figure 6). However, real-world reflections are off-specular, i.e., they are not necessarily aligned with the mirror direction. To improve the normal estimation further we could use photometric stereo methods [Goldman et al. 2003].

To separate specular surface reflectance from diffuse subsurface reflectance at each surfel we use the diffuse BRDF approximation of the BSSRDF by Jensen et al. [2001]:

$$S_{brdf}(x, \omega_i, \omega_o) = \frac{R_d}{\pi},\qquad(4)$$

where R_d is the diffuse reflectance:

$$R_d = \frac{\alpha'}{2}\left(1 + e^{-\frac{4}{3}\frac{1+F_r}{1-F_r}\sqrt{3(1-\alpha')}}\right)e^{-\sqrt{3(1-\alpha')}}.\qquad(5)$$

R_d depends only on the reduced albedo α' and the Fresnel terms. The BRDF approximation is equivalent to the full BSSRDF for a semi-infinite plane of homogeneous, almost opaque material under uniform incident illumination. Intuitively, under these assumptions we cannot distinguish if photons enter the surface, scatter, and re-emit as diffuse light, or if they are immediately reflected at a point from a diffuse surface. We found this approximation to be relatively accurate for the diffuse reflectance component of our skin model. Any remaining error will show up as a diffuse component in the surface BRDF.

We estimate R_d at each surface point from the lumitexel data. Based on the assumption that we observe pure diffuse reflectance for at least some of the observation angles,

We are looking for the maximum value R_d that is less than or equal to the minimum of the observed BRDF values f_r. This is

done by determining the minimum ratio between f_r and the unit diffusion reflectance:

$$R_d = \min_i \frac{\pi f_r(\omega_i, \omega_o)}{F_t(\eta, \omega_i)F_t(\eta, \omega_o)}.\qquad(6)$$

Note that we divide the observed BRDF by the Fresnel coefficients in accordance with Equation (2) as we compute the diffuse component using the BSSRDF. 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 R_d values for each surface point are re-parameterized and interpolated into the albedo map.

7 Computing a Basis for Surface BRDFs

The interface reflection is computed from each lumitexel by subtracting the diffuse reflectance: $R_i = R - R_d$. Instead of fitting an analytic reflectance model to this data [Fuchs et al. 2005] we use the data-driven BRDF model of Matusik et al. [2003]. The goal is to express the data R_i' as a linear combination of densely measured basis BRDFs. We use 59 basis BRDFs of dielectric materials that are related to human skin, such as paints, fabrics, organic materials, and leather. Each basis BRDF consists of over four million samples.

To account for the area of the light source, we convolve each basis BRDF by a disk-shaped filter kernel. This works because each light source spans approximately the same solid angle when viewed from any point on the face.

Let ϑ_l be the angle between illumination direction and surface normal. Rather than normalizing the observed reflectance samples by dividing them with $\cos \vartheta_l$, we multiply the basis BRDFs by this factor. This minimizes the influence of noise and quantization artifacts at grazing angles.

Assume we have n observed reflectance samples $R_i(\omega_i^i, \omega_o^i)$ and m basis BRDFs M^j. We construct an $n \times m$ matrix M where the elements in the i-th row and j-th column are:

$$M_{ij} = \cos \vartheta_l M^j(\omega_i^i, \omega_o^i).\qquad(7)$$

We now solve for the vector x in the system:

$$Mx = R \; s.t. \; x_i \geq 0,\qquad(8)$$

where R is an $n \times 1$ column vector with elements R_i. We use quadratic programming to solve this over-constrained least-squares problem. The resulting x_i are the coefficients of the linear combination of basis BRDFs that best reproduce the lumitexel data. We found that constraining the coefficients to be positive is very important for the stability of the solution. To further improve the results we take the cubic root of the elements in M and R. This transforms the basis BRDFs and reflectance values into a more perceptually appropriate space where highlights are not oversaturated.

Highlights on dielectric materials like skin are of the same color as the light source (white, in our case). Consequently, we can use monochrome basis BRDFs. This reduces the degrees of freedom by a factor of three and increases the stability of the fit.

The 59 basis BRDFs are not specific to human skin – they are capable of representing reflectance of a much wider range of materials [Matusik et al. 2003]. We use dimensionality reduction to discover a smaller linear basis that is specific to human skin. We apply NMF to the vector x of positive coefficients at each surface point independently. Unlike principle-component analysis (PCA), the NMF basis is not orthogonal. However, the original data can be expressed as *positive* linear combinations of the NMF basis vectors. That is, when fitting reflectance data to the reduced NMF basis, we can still perform non-negative least squares optimization. Using

PCA we would have to allow for negative coefficients, leading to over-fitting and visible artifacts like "negative highlights" [Matusik et al. 2003].

To determine the smallest NMF basis that is necessary to represent the data well we plot the reconstruction error as a function of the number of NMF basis BRDFs (see Figure 15). There is a considerable fall-off in the sequential values seen in this plot. The plateau for a single person is reached around four basis BRDFs when the relative reconstruction error does not change much anymore. Figure 7 shows synthesized images of a face using different numbers of NMF basis BRDFs. As expected, the results do not im-
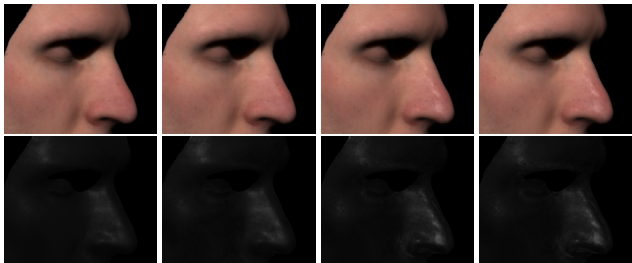

Figure 7: *Synthetic images of a face using, from left to right, 1, 2, 4, and 24 NMF basis BRDFs. Top: Rendered images. Bottom: Surface reflection only. Four NMF basis BRDFs are sufficient to represent the surface reflectance of a single person.*

prove much beyond four NMF basis BRDFs. All renderings shown in this paper (except where noted) were produced using four NMF basis BRDFs. The coefficients for each surface point are stored in four BRDF coefficient textures. In addition, we store the four NMF basis BRDFs per face. We analyze a more general BRDF basis for a large population of people in Section 10.1.

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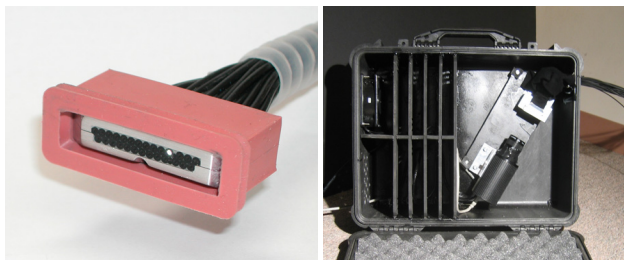
Figure 8: *Left: A picture of the sensor head with linear fiber array. The source fiber is lit. Right: The fiber array leads to a camera in a light-proof box. The box is cooled to minimize imaging sensor noise.*

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 9 shows the sensor head placed on a face. We found that pressure variations on the skin caused by the mechanical movement
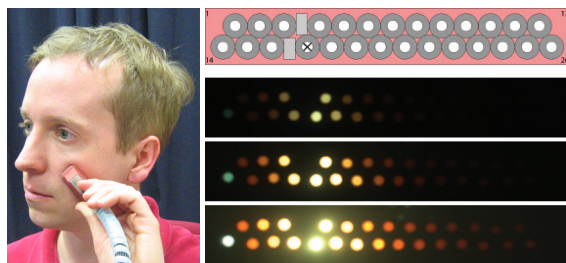
Figure 9: *Left: The sensor head placed on a face. Top: Sensor fiber layout. The source fiber is denoted by a cross. Bottom: An HDR image of the detector fibers displayed with three different exposure values.*

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. For more details on the subsurface device and calibration procedure see [Anonymous 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
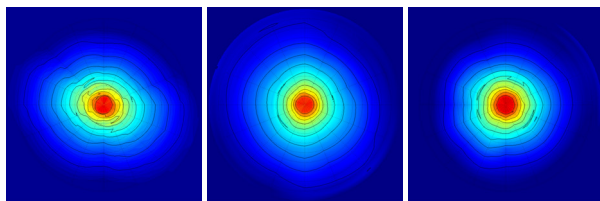
Figure 10: *Subsurface scattering curves for abdomen, cheek, and forehead measured along 16 1D profiles.*

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.

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. Note that these parameters also captures high-frequency albedo variations (redness, freckles etc.) However, it would be impractical to measure them densely across the face using our sensor head. Instead, we rely on the high-resolution photographs from the face-scanning dome to estimate diffuse albedo (see Section 6) and use the sensor head to measure slowly varying translucency. 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.

From the measured σ_a and σ_s' data we derive the effective transport coefficient:

$$\sigma_{tr} = \sqrt{3\sigma_a(\sigma_a + \sigma_s')} \approx 1/\ell_d \,, \tag{9}$$

with ℓ_d the diffuse mean free path. ℓ_d or $1/\sigma_{tr}$ provides a measure of skin translucency. We found that it shows little variation across a face. We interpolate $1/\sigma_{tr}$ from the three measurements to obtain a dense translucency map for the face.

Nbr. of Subjects	External Factors	Race	Gender	Age
Dataset A				
1	Normal / Cold Hot / Sweat Lotion / Makeup Powder	Caucasian	Male	31
Dataset B				
50	Normal	Caucasian	Male	13-59
5	Normal	Caucasian	Female	24-55
6	Lotion / Makeup	Caucasian	Female	23-56
20	Normal	Asian	Male	22-74
4	Normal	Asian	Female	22-65
2	Lotion / Makeup	Asian	Female	24-30

Table 1: *Relevant information about our collected data.*

9 Rendering

We implemented our reflectance model in a high-quality Monte Carlo ray tracer for offline rendering. For interactive rendering of our model we use the approach by Kautz and McCool [1999]. [3] The inputs for both systems are four textures with the coefficients of the NMF basis BRDFs, the four NMF basis BRDFs, the albedo map (R_d), and the translucency map ($1/\sigma_{tr}$).

To achieve high-quality images we use the analytic BSSRDF approximation by Jensen et al. [2001]. We transform the R_d values of the albedo map to apparent albedo values α' by inverting Equation (5). We derive the model parameters σ'_s and σ_a from σ_{tr} using [Jensen and Buhler 2002] $\sigma'_s = \alpha'\sigma'_t$ and $\sigma_a = \sigma'_t - \sigma'_s$, with $\sigma'_t = \sigma_{tr}/\sqrt{3(1-\alpha')}$. Surface reflectance at each surface point is computed using a linear combination of the NMF basis BRDFs.

We show the comparisons between real and synthetic images for different faces and different viewpoints in Figure 11. The camera and light source calibrations of the dome were used to reproduce identical conditions. We observe that our model reproduces the photographs very well, including the shape of the specular highlights. Figure 12 shows another example, including renderings of the different components in our model.

10 Skin Reflectance Analysis

We now present an analysis of the three components of our face reflectance model for different external conditions and for a large population of people of different gender, race, and age. Table 1 shows the relevant information for the data we collected. To date, we captured 87 people: 70 male and 17 female. The data is heavily skewed towards male Caucasians and does not currently contain any African Americans. Our data collection effort will be ongoing for the foreseeable future. Because each capture session takes about 30 minutes we could not capture all individuals under all external conditions. Instead, we first analyze how various external conditions affect one person (dataset A), and then analyze the remaining subjects (dataset B). Some women in dataset B wore makeup or facial lotion. Figure 13 shows a few representative subjects in dataset B.

Our model treats each component independently, allowing us to render or edit each one without affecting the other two. Consequently, we analyze each component separately and do not consider the correlation between them. In the future we would like to perform a more in-depth statistical analysis that takes cross-correlations into account.

---

[3]See companion video.

Figure 13: *Representative subjects in dataset B.*

10.1 Analysis of Surface BRDFs

Surface BRDFs capture the monochrome light reflection on the oil-skin layer. They can be represented with a low-dimensional NMF basis of densely-sampled BRDFs. The effect of the surface BRDF on overall appearance is relatively pronounced, especially for different external conditions (dataset A) (see Figure 19).

To analyze the surface BRDFs of dataset A we randomly chose 5,000 points and fit the complete non-parametric BRDF basis as discussed in Section 7. We then computed the average of the coefficients of all 5,000 points. Figure 14 shows the average BRDF for each external condition applied to a sphere and lit with point light sources from two different directions. As expected, there are noticeable differences between these BRDFs, especially between lotion / hot and cold / powder.

To analyze the space of surface BRDFs of a larger population, we fit our non-parametric BRDF model to 5,000 (dataset A) or 2,000 (dataset B) randomly chosen points, respectively. Similar to the approach in Section 7 we applied NMF dimensionality reduction to obtain a low-dimensional manifold that characterizes the BRDF space for each dataset. Figure 15 shows relative reconstruction error as a function of the number of basis BRDFs for one person, dataset A, and dataset B, respectively. To show the three curves in
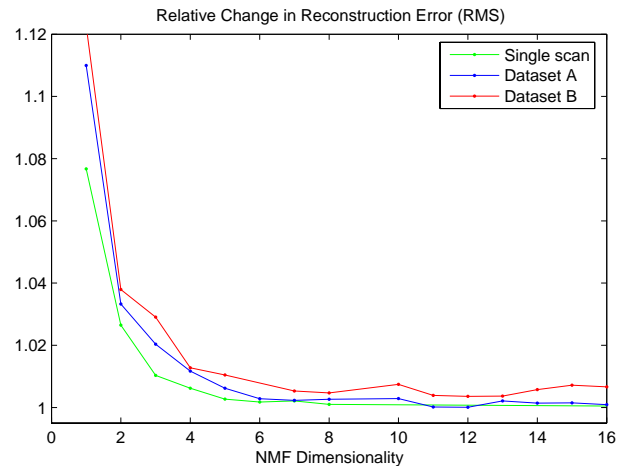

Figure 15: *Reconstruction error as a function of the number of NMF basis.*

the same plot we normalized the relative errors to a common scale. As expected, the plot suggests that four NMF basis BRDFs are sufficient for one person, whereas dataset A requires six and dataset B requires at least eight NMF basis BRDFs. Matusik et al. [2003] concluded that a 45-dimensional linear (PCA) basis is required to span the space of isotropic BRDFs. Our results suggest that BRDFs representing skin are a small subset of all isotropic BRDFs.

Figure 11: *Comparison of real photographs of five subjects (top) with fits to our model (bottom).*
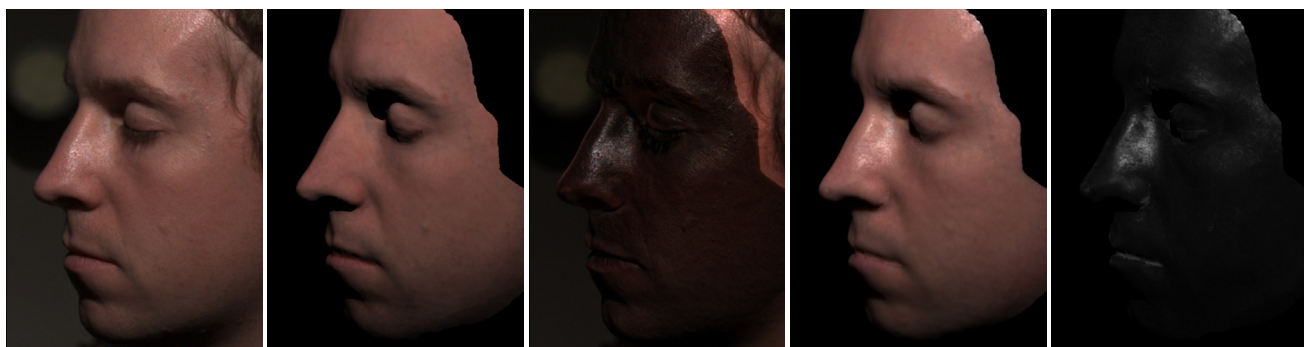
Figure 12: *Components of our model. (a) One of the input images with single light source illumination. (b) 3D surface scan shaded with the BRDF approximation of the diffuse subsurface term. The model is lit using the camera and light source calibration derived for view (a). (c) Subtracting the estimated diffuse reflection term (b) from (a) reveals surface reflection. (d) Reconstructed model. (f) Reconstructed surface reflection.*

| Normal | Lotion | Cold | Hot | Make-up | Sweaty | Powder |
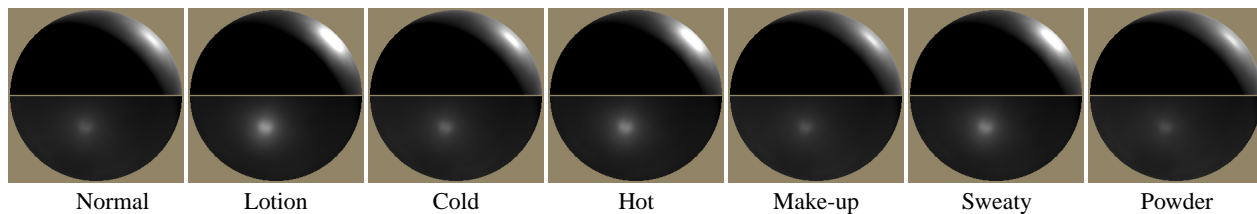
Figure 14: *Visualization of the average surface BRDF for dataset A. Top: Back lighting from grazing angle. Bottom: Front lighting from lower left.*

10.2 Analysis of Diffuse Albedo Maps

The diffuse albedo is the component of our model that captures most of the intra-personal and extra-personal variations. Small-scale intra-personal albedo variations are due to skin imperfections, markings, scars, etc. and typically show high spatial frequency across the face. Overall extra-personal albedo variations are mainly due to race, gender, tanning, or other external factors.

We first transform all albedo maps in each dataset into a common and decorrelated color space using the method discussed in [Heeger and Bergen 1995][Section 3.5]. In the following analysis we process each transformed color channel independently.

An albedo map bears many similarities to a stochastic image texture. Consequently, we apply the texture analysis method of Heeger and Bergen [1995]. We 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 [Simoncelli and Freeman 1995]. 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 histograms of all filter responses including a low-pass and a high-pass (30 total) and the histogram of the original albedo map are concatenated into a $256 \times 31 \times 3 = 23,808$ element vector H. This vector can be viewed as a generative model of a texture for a given person. For example, we can use this vector for albedo transfer between two subjects using histogram matching [Heeger and Bergen 1995] (see Section 11).

Figure 16 shows color histograms of the original albedo map averaged over different groups of people in dataset B. For clarity we used the original RGB space to compute these histograms. As expected, the plots show a different color distributions depending on race and gender.

To analyze the extra-personal variations, we resample the albedo maps of dataset B into a common (u, v) parameter space using point-correspondences (20 feature points per face) and radial basis function interpolation [Pighin et al. 1998]. To obtain reliable statistics we cluster points in corresponding face regions. Currently, we only specify beard versus no-beard regions, but one could use a finer granularity and distinguish between chin, cheek, forehead, nose, eyelid, and beard areas. For each face i and each region r we compute a histogram vector H_i^r as described above. All vectors H_i^k for a specific region k are stored as column vectors in a matrix M_h. For example, M_h for the beard regions in dataset B has dimensions $87 \times 23,808$ (there are 87 subjects in dataset B). We can now run PCA on matrix M_h to compute a region-specific basis for albedo map histograms. Each point in this reduced space corresponds to the albedo of a particular person. We will use this PCA basis to synthesize new albedo maps in Section 11.

Cula and Dana [Cula and Dana 2002] use a very similar method to analyze bi-directional texture functions (BTFs) of a collection of skin patches. However, they do not use their model for image synthesis.

10.3 Analysis of Translucency

The translucency component accounts for non-local subsurface scattering in the epidermis and dermis. It is a slowly varying effect that is responsible for much of the red color and soft shadows we see in human faces. It is important to note that translucency cannot be estimated directly from images, which is why we additionally use subsurface measurements.

Table 2 shows the mean and variance of σ_{tr} for dataset B. The measurement points on 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
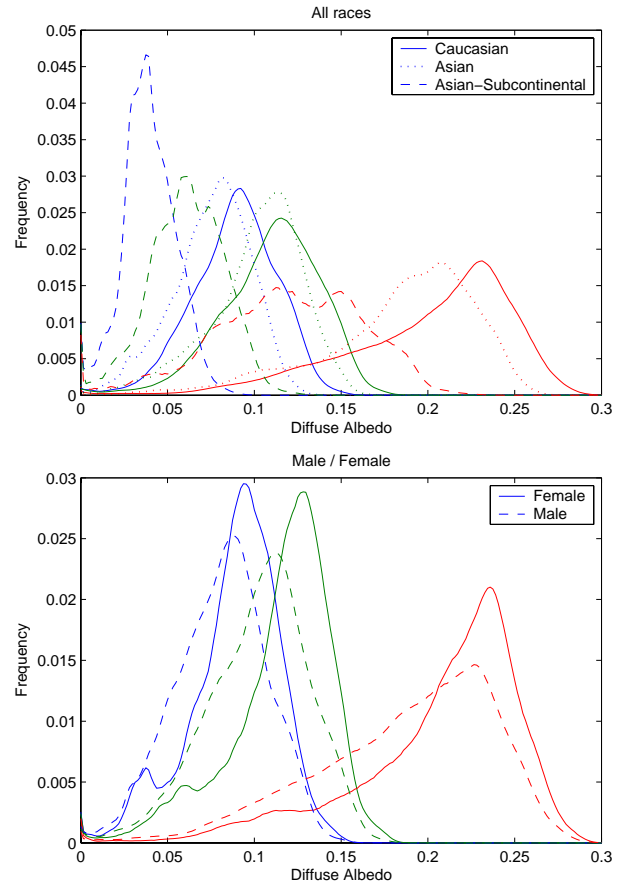

Figure 16: *Average color histograms for different race and gender groups in dataset B. Subsequent plots show, from right to left, data for red, blue, and green channels (also indicated by the line colors).*

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 2: *Mean and variance of σ_{tr} for dataset B.*

to place there. Overall, translucency values do not vary much between measurement points and between individuals. In practice, one could approximate it using a single value for the whole face.

Figure 17 shows closeups of the subjects with minimum (0.3558, 0.7932, 1.5173) and maximum (0.9171, 1.5682, 1.6569) values for σ_{tr} in dataset B. Note that we define translucency as $1/\sigma_{tr}$. There
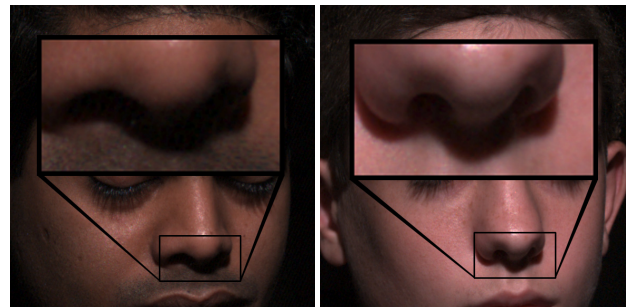
Figure 17: *Photographs of subjects with minimum (left) and maximum (right) translucency values in dataset B. The differences at shadow boundaries are subtle.*

are subtle differences visible at shadow boundaries. Figure 18 shows closeups computed with our model using the same minimum and maximum translucency values. Note that the model is capable

Figure 18: *Synthetic images with minimum (left) and maximum (right) translucency values.*

of reproducing the subtle differences of Figure 17.

11 Face Editing

Similar to previous work [Pellacini et al. 2000; Matusik et al. 2003] we define meaningful parameters for face editing. The user assigns arbitrary traits to each face using binary classification (trait present or not). We use normal, cold, hot, sweat, lotion, makeup, powder for dataset A, and Asian, Asian-Subcontinental, Caucasian, male, female for dataset B. The user can choose a face and change its reflectance properties according to any of the defined traits, e.g., making a face look more tanned. We apply this general idea to surface BRDFs and albedo maps. Translucency maps could be handled in a similar way, if desired.

Similar to [Blanz and Vetter 1999; Matusik et al. 2003] we use mean differences to navigate the low-dimensional spaces of surface BRDFs (using their NMF basis) and albedo histograms (using their PCA basis). We compute the average of the basis vectors in each complementary pair of clusters associated with a trait (i.e., those faces with, and those without). The differences between the complement averages provide *trait vectors* that can be used for navigation and interpolation between traits. I.e., we use convex combinations of trait vectors and apply them to the data of a source face.

Specifically, we compute trait vectors for the NMF basis of surface BRDFs in dataset B (see Section 10.1). To compute a new (target) BRDF, we apply a linear combination of the trait vectors to the BRDF coefficients of a source face.

Figure 19 shows progressions of adding different traits to source faces. The last row in the figure shows an example where we change the surface BRDF using the "lotion" trait. The albedo map and translucency remain constant in this case.

For albedo map changes we compute trait vectors using the PCA basis of albedo histograms (see Section 10.2). A linear combination of trait vectors is applied to a basis histogram vector H of a source face, resulting in a target histogram vector H'. We then apply the histogram-matching technique of Heeger and Bergen [1995] to match H to H'. We either use the basis histograms of the whole albedo map or of facial regions (beard versus no-beard area).

Note that Heeger and Bergen start their texture synthesis with a noise image as the source. We could do the same (with satisfactory results). However, for most applications it makes more sense to start from an original 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.

Rows one through three in Figure 19 show examples of changing albedo maps *and* surface BRDFs using various trait vectors. Translucency remains constant in all cases.

12 Conclusions and Future Work

In this paper we have proposed a simple and practical skin model that is powerful enough to represent most aspects of facial skin appearance. We combine an analytic model for subsurface scattering with a low-parameter non-parametric BRDF model. An important feature of our model is that all its parameters can be robustly estimated from measurements. Renderings using our model are capable of reproducing photographs of real human faces taken under arbitrary illumination and pose. We fit our model to data of a large and diverse group of people. We have also addressed the problem of editing facial reflectance. To accomplish this goal we have developed intuitive user parameters such as age, race, gender, tan-level, etc.

Face appearance is of course not only determined by face reflectance. A lot of realism comes from facial hair, which is currently not represented at all in our model. We would like to develop novel representations for eyebrows, eyelashes, mustaches, and beards. Fine facial hair leads to the important "velvet" look of skin near grazing angles [Koenderink and Pont 2003]. We would like to extend our reflectance model to account for this effect. We believe our methods could also be extended to model reflectance of other important parts of the face, such as the eyes, lips, and teeth.

Head motion or slight changes in facial expressions lead to noise in our reflectance measurements and to visible blur in renderings from our model (e.g., near the lips). We are developing better procedures to detect registration errors and to improve the quality of our measurements.

At the moment we do not consider global-illumination effects such as inter-reflections or self-shadowing due to fine-scale geometry. We would like to improve the estimation of reflectance components by adding these effects by using methods similar to Yu et al. [1999].

References

ANONYMOUS. 2005. Measurement devices for face scanning. Tech. rep., Anonymous. Will be published before SIGGRAPH.

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

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.

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

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

DESBRUN, M., MEYER, M., AND ALLIEZ, P. 2002. Meshes and parametrization: Intrinsic parameterizations of surface meshes. *Computer Graphics Forum 21* (Sept.).

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,

Figure 19: *Face editing by adding trait vectors to different source faces. From top to bottom: beard, Caucasian, tan, and lotion trait.*

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.

GOLDMAN, D., CURLESS, B., HERTZMANN, A., AND SEITZ, S. 2003. Shape and spatially-varying brdfs from photometric stereo. Tech. Rep. 04-05-03, University of Washington.

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

HARO, A., GUENTER, B., AND ESSA, I. 2001. Real-time, photo-realistic, physically based rendering of fine scale human skin structure. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, 53–62.

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.

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.

KAUTZ, J., AND MCCOOL, M. 1999. Interactive rendering with arbitrary BRDFs using separable approximations. In *Rendering Techniques '99 (Proceedings of the Tenth Eurographics Workshop on Rendering)*, Springer Wien, New York, NY, 281–292.

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.

LEE, D., AND SEUNG, H. 1999. Learning the parts of objects by non-negative matrix factorization. *Nature 401*, 788–791.

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.

MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics (SIGGRAPH 2003) 22*, 3 (July), 759–770.

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

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.

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

PELLACINI, F., FERWERDA, J., AND GREENBERG, D. 2000. Toward a psychophysically-based light reflection model for image synthesis. *Computer Graphics 34*, Annual Conference Series, 55–64.

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.

RUTGERS. Rutgers skin texture database. Web page. http://www.rutgers.edu/rutgers_texture.

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

SIMONCELLI, E. P., AND FREEMAN, W. T. 1995. The steerable pyramid: A flexible architecture for multi-scale derivative computation. In *ICIP '95: Proceedings of the 1995 International Conference on Image Processing*, IEEE Computer Society, vol. 3, 3444–3452.

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.

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.

WILLIAMS, L. 1978. Casting curved shadows on curved surfaces. In *Computer Graphics (Proceedings of SIGGRAPH 78)*, vol. 12, 270–274.

YU, Y., DEBEVEC, P., MALIK, J., AND HAWKINS, T. 1999. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *Computer Graphics*, SIGGRAPH 99 Proceedings, 215–224.

ZWICKER, M., PFISTER., H., BAAR, J. V., AND GROSS, M. 2001. Surface splatting. In *Computer Graphics*, SIGGRAPH 2001 Proceedings, 371–378.

ZWICKER, M., PFISTER, H., BAAR, J. V., AND GROSS, M. 2002. Ewa splatting. *IEEE Transactions on Visualization and Computer Graphics 8*, 3, 223–238.