

A Layered, Heterogeneous Reflectance Model for Acquiring and Rendering Human Skin

Craig Donner*
*Columbia University

Tim Weyrich^{†‡}
†Princeton University

Eugene d'Eon[§]
‡University College London

Ravi Ramamoorthi*

Szymon Rusinkiewicz[†]
§NVIDIA Corporation




Figure 1: Our model approximates heterogeneous light transport in skin through the inter-scattering of light between layers. Shown are hand-drawn spatially-varying parameter maps of two scattering layers, and of an infinitesimally thin absorbing layer between them. The maps have been scaled in intensity (10x for melanin, 20x for hemoglobin) to show detail. Given this simple input, our method renders heterogeneous and volumetric effects that cannot be simulated using previous methods.

Abstract

We introduce a layered, heterogeneous spectral reflectance model for human skin. The model captures the inter-scattering of light among layers, each of which may have an independent set of spatially-varying absorption and scattering parameters. For greater physical accuracy and control, we introduce an infinitesimally thin absorbing layer between scattering layers. To obtain parameters for our model, we use a novel acquisition method that begins with multi-spectral photographs. By using an inverse rendering technique, along with known chromophore spectra, we optimize for the best set of parameters for each pixel of a patch. Our method finds close matches to a wide variety of inputs with low residual error.

We apply our model to faithfully reproduce the complex variations in skin pigmentation. This is in contrast to most previous work, which assumes that skin is homogeneous or composed of homogeneous layers. We demonstrate the accuracy and flexibility of our model by creating complex skin visual effects such as veins, tattoos, rashes, and freckles, which would be difficult to author using only albedo textures at the skin’s outer surface. Also, by varying the parameters to our model, we simulate effects from external forces, such as visible changes in blood flow within the skin due to external pressure.

CR Categories: I.3.7 [Computing Methodologies]: Computer Graphics—Three-Dimensional Graphics and Realism

Keywords: Skin reflectance, subsurface scattering, BSSRDF, reflection models, layered materials, light transport

1 Introduction

Human skin exhibits a striking range of appearance due to physiological and structural variations across its surface and within its layers. This heterogeneity may occur naturally, such as freckles, splotches, veins, or rashes, or be introduced artificially via external forces (e.g. pressure causing blood to increase or decrease) or pigmentation such as tattoos. When building shading models to describe, and designing methods to acquire, skin’s appearance, it is imperative to take these variations into account.

Understanding the appearance of skin is particularly challenging, as skin reflectance is dominated by subsurface scattering [Igarashi et al. 2005]. In addition, the spectral reflectance of skin is complex due to its many chemical constituents and structural variation [van Gemert et al. 1989]. Existing skin models, such as those based on the diffusion approximation [Jensen et al. 2001; Donner and Jensen 2005], generally rely on homogeneous approximations for both rendering and acquisition. To give the impression of heterogeneity, they rely on an approximate translucency modulated by a surface albedo texture [Jensen and Buhler 2002; Hery 2003]. Acquiring or generating albedo textures for use under complex and dynamic lighting conditions or during physiological changes in skin, however, remains a difficult problem.

To address these challenges, we introduce a layered, heterogeneous reflectance model for skin (Section 4). Our model is based on recent work on multi-layered translucent materials [Donner and Jensen 2005], but instead of computing simple profiles for materials, we use this layered light transport to model heterogeneous scattering of light. Though we assume local homogeneity around a particular point, we apply the spatially-varying properties of the layers as light scatters between them. This allows our model to capture spatial variation without applying albedo textures to the skin’s outer surface. Instead, we drive our model with intuitive parameter maps over the surface of the individual scattering layers, and allow thin absorbing layers between the scattering layers. We describe an efficient rendering method, based on applying multiple Gaussian blurs to the irradiance (Section 4.3).

We apply our model to capturing and reproducing skin appearance (Section 5). We use our physically-inspired skin model with multiple layers and heterogeneous distributions of chromophores

(pigments and dyes) among them. We describe a novel technique for recovering maps of these properties from multi-spectral images of real skin (Section 6). Our acquisition setup uses a digital camera, flash, and narrow-band color filters. Parameter maps are recovered using inverse rendering, assuming known chromophore absorption spectra. This method is versatile, and we obtain representative chromophore distributions over a wide range of skin types and body regions.

We evaluate our methods by synthesizing realistic images of skin, as illustrated in Figure 1. We also demonstrate how to simulate physiological changes in skin by modifying the maps that drive the model parameters. We show the effect of heterogeneity and evaluate our acquisition method, demonstrating skin renderings from acquired data.

2 Previous Work

The reflectance of translucent materials, and skin in particular, has received considerable attention in computer graphics. Here we focus on the work related to shading models and reflectance measurements, as they relate to our proposed methods. A thorough survey of techniques is presented by Igarashi et al. [2005].

2.1 Reflectance models

Reflectance models for skin range from simple BRDF approximations to the scattering of light [Hanrahan and Krueger 1993; Stam 2001], to comprehensive models that simulate small-scale anatomical and physiological detail [Krishnaswamy and Baranoski 2004]. The former rely on 1D approximations to light transport, and hence do not capture the characteristic subsurface scattering of translucent materials. Models that rely on an accurate volumetric light transport simulation are typically too complex and too slow for practical applications in graphics. Models based on the diffusion approximation offer a compromise between accuracy and complexity. Jensen et al. [2001] introduced the diffusion dipole approximation along with a simple technique for measuring optical parameters of materials, including two samples of skin. The method is limited to homogeneous, semi-infinite slabs with constant scattering and absorption properties.

Despite these limitations, methods using the dipole model to approximate the appearance of heterogeneous materials are often employed. Jensen and Buhler [2002] and Héry [2003] fix one parameter (translucency) of the model and vary another parameter to match a given albedo texture. Others have used the dipole model with fully spatially varying parameters [Tariq et al. 2006], but this use of the model in the presence of heterogeneity, particularly when there are large variations in parameters within small areas, is not well defined. All of these approaches assume that the transport of light between two points on the material depends only on a single set of homogeneous parameters. Weyrich et al. [2006] use a modulation texture on top of the homogeneous scattering process, which renders faithfully under uniform lighting. Such modulation textures, however, produce incorrect results under structured illumination, as we show in Section 4.

For rendering thin slabs and multi-layered translucent materials, Donner and Jensen [2005] introduced the diffusion multipole approximation. Although this model more accurately captures the reflectance of translucent layered materials, it is difficult to obtain optical parameters for specific materials. The multipole model was later applied to create a spectral shading model specifically for skin [Donner and Jensen 2006], which uses a small set of chromophore parameters to control overall skin appearance. Spatial variations across the surface, however, are still approximated with an albedo modulation texture. D’Eon et al. [2007] have recently implemented the multipole model to run on modern graphics hardware. Though this method allows interactive rendering of layered translucent materials, it still uses the albedo texture approximation.

More general heterogeneous BSSRDF models have been developed with the specific goal of obtaining model parameters from measurements [Goesele et al. 2004; Tong et al. 2005; Peers et al. 2006]. These models capture the relationship between incident and exitant light, but ignore the underlying physical and physiological structure of the material. For skin, it is particularly valuable to understand the relationship between structure, pigmentation, and reflectance, as this allows the derivation of physical models for predicting the appearance of arbitrary samples of skin. Wang et al. [2007] acquire a volumetric representation of heterogeneously scattering materials. They describe a numerical method for approximating heterogeneous diffusion on the GPU, but require a complex discretization of the geometry into a coupled polygrid.

2.2 Skin Analysis

Early work that acquired and analyzed skin reflectance is based on pure BRDF models. Dana et al. [1999] and Marschner et al. [1999] use a camera and a light source to acquire images of skin under different angles of incidence and reflection. Debevec et al. [2000] and Weyrich et al. [2006] use parallel hardware to rapidly acquire the reflectance field of a human face and subsequently fit a BRDF model to the acquired data. Purely BRDF-based approaches, however, cannot reproduce subsurface scattering.

More recent work on acquiring parameters of subsurface scattering concentrates on radial scattering profiles, which requires structured illumination to observe the profiles. Jensen et al. [2001] use a tightly-focused beam of light to produce well-defined, structured illumination, which exposes the scattering profile around a single surface point. Weyrich et al. [2006] measure similar scattering profiles using point illumination and a sensor array of optical fibers. The use of a contact device, however, limits applicability to flat skin areas where the probe can be placed. Tariq et al. [2006] avoid this shortcoming by projecting stripe patterns on a face, shifting them through time, and observing the resulting scattering response of a moving step function of illumination. Ghosh et al. [2008] have concurrently with this work used structured illumination to capture and reproduce the layered reflectance of faces. In our setup, we control incident illumination by attaching a black tape occluder to the skin. In addition, we avoid having to measure individual scattering profiles at each point of the skin by fixing the reduced scattering coefficient $\sigma'_s(\lambda)$. This choice is more physically accurate than the similarly simplifying assumption of fixing the effective transport length σ_{tr}^{-1} [Jensen and Buhler 2002], as it holds for the vast majority of structures found in skin [Jacques 1998; Donner and Jensen 2006]. Other acquisition methods to capture general, heterogeneously scattering materials require extensive acquisition times and do not lend themselves to measurements of skin *in vivo* [Goesele et al. 2004; Tong et al. 2005; Peers et al. 2006; Wang et al. 2007].

Rather than modeling reflectance from first principles, a common approach is also to tabulate a particular skin sample (such as an actor’s face) under different lighting conditions [Georghiadis et al. 1999; Debevec et al. 2000; Borshukov and Lewis 2003; Cula et al. 2004; Hawkins et al. 2004; Wenger et al. 2005]. While these methods faithfully reproduce the measured skin’s reflectance, they do not provide insight on how skin interacts with light. This insight is valuable when developing digital characters, or when simulating poses or conditions (such as changes in pigmentation) significantly different from the captured data.

More specialized image-based methods derive chromophore concentrations from single images. Tsumura et al. [2003] qualitatively estimate melanin and hemoglobin distributions in a face using independent component analysis (ICA). This method is based on a purely absorptive reflectance model that ignores the influence of scattering and depth on color. Cotton et al. [1999] derive melanin distributions from photographs by analyzing color deviations from a model based on Kubelka-Munk theory. Though this technique has

gained popularity in the medical physics community, the underlying model is still effectively a BRDF, and cannot capture the lateral subsurface transport of light. It is this subsurface transport that is critical to reproducing the appearance of skin.

Often, reflectance models and acquisition methods have been developed in parallel. Continuing in this tradition, we design our measurement system around a physically-based model, and use our measurements to derive guidelines for parameter choice.

3 Background: Layered Translucent Materials

Given absorption and reduced scattering coefficients σ_a and σ'_s , and the thicknesses of the layers, the diffusion multipole model [Donner and Jensen 2005] uses sums of contributions from point sources arranged in mirrored configurations about each layer to compute reflectance and transmittance profiles. These profiles are radially symmetric, and assume the radiant exitance from the homogeneous layer is diffuse. Note that these profiles define the 2D radiant exitance over the surface given an input ray of light. We will later extend this formulation to generate 2D heterogeneous profiles.

Given a two-layer translucent material, we denote its reflectance and transmittance profiles as R_{12}^+ , T_{12}^+ , R_{12}^- and T_{12}^- . The + and - indicate whether light is propagating deeper into (+) or out of (-) the material before it scatters. The total radiant exitance is the convolution of the incident flux Φ at the surface of the material with its forward diffuse reflectance profile:

$$M(x, y) = \iint \Phi(x', y', \vec{\omega}) R_{12}^+(x-x', y-y') dx' dy' = \Phi * R_{12}^+. \quad (1)$$

The forward reflectance profile is itself composed of multiple terms that capture the degree of inter-scattering between layers. Thus, R_{12}^+ includes the sum of the profile accounting for light that scatters only in the top layer R_1^+ , and the convolution of the profiles accounting for light that scatters from the top layer into the lower layer, back into the top layer, and out of the top layer at the surface:

$$\begin{aligned} R_{12}^+(x, y) = & R_1^+(x, y) + \iiint T_1^+(x', y') R_2^+(x'-x'', y'-y'') \\ & \cdot T_1^-(x-x'', y-y'') dx' dy' dx'' dy'' \\ & + \dots = R_1^+ + T_1^+ * R_2^+ * T_1^- + \dots \end{aligned} \quad (2)$$

where T_1^- is the backward transmission profile of the top layer, and R_2^+ is the forward reflectance profile of the lower layer. This equation assumes that light emitted from the material is either scattered directly from the top layer, or scattered down into the bottom layer before returning and exiting. Thus, it represents a single inter-scattering of light between the layers.

Accounting for the full inter-scattering of light requires summing a series of convolutions:

$$\begin{aligned} \Phi * R_{12}^+ = & \Phi * (R_1^+ + T_1^+ * R_2^+ * T_1^- + T_1^+ * R_2^+ * R_1^- * R_2^+ * T_1^- \\ & + T_1^+ * R_2^+ * R_1^- * R_2^+ * R_1^- * R_2^+ * T_1^- + \dots). \end{aligned} \quad (3)$$

Due to the associativity of convolution, it is more efficient to perform the convolution with irradiance after the sum of layer convolutions has been computed. Note that when the profiles in (3) are transformed into the Fourier domain, the resulting sum of pointwise products is a geometric series, and reduces to a simple form that is inexpensive to evaluate when the profiles are radially symmetric [Donner and Jensen 2005].

D'Eon et al. [2007] observed that a set of Gaussian functions provides a useful basis for accurately representing radially symmetric diffusion profiles. By expressing each profile as a linear combination of k Gaussians of variance v_i :

$$R_1^+ = \sum_{i=1}^k w_i G(v_i), \quad (4)$$




Figure 2: Cross-sections of skin with varying amounts of melanin. The samples have been stained to highlight the melanin distribution. From left to right, the samples represent light skin, moderately pigmented skin, and heavily pigmented skin. Note that melanin is distributed fairly evenly throughout the epidermis, with a concentration at the junction between the epidermis and dermis. Images used with permission from [Matts et al. 2007].

convolution by a non-separable profile R_1^+ is efficiently computed as a sum of separable convolutions.

4 A Heterogeneous Reflectance Model for Skin

One of the most striking visual components of skin is its color, caused by subsurface scattering. Light scatters among and is absorbed by skin's structural and chemical constituents [Igarashi et al. 2005]. The absorption of light in skin is largely due to chromophores, chemicals that selectively absorb different wavelengths of light. Scattering occurs from small-scale cellular structures, collagen, chromophores, as well as high-frequency changes in index of refraction. This is generally modeled at the macroscopic level as scattering from particles [Jacques 1996]. Both scattering and absorption vary significantly through the visible spectrum and across the surface.

From the standpoint of pigmentation, skin is roughly divided into two layers: the epidermis at the surface, with the dermis below. Some pigmentation is focused in a narrow region between the epidermis and dermis [Matts et al. 2007]. Melanin, for example, is produced by cells at the base of the epidermis, and thus has a significant concentration there (see Figure 2). Artificial pigmentation, such as in a tattoo, is usually injected into the upper dermis to prevent the pigment from dissipating as epidermal cells regenerate [Bernstein 2006].

We model skin as a two-layer heterogeneous material, and account for inhomogeneous pigmentation in skin in two ways: by varying optical properties over the surfaces of the skin layers, and by introducing thin absorbing layers between the scattering layers (see Figure 3). Our formulation of heterogeneity models the physical characteristics skin, where the epidermis is thin and highly scattering, and has the highest gradient of optical properties of the skin layers.

In general, representing the heterogeneous transport of light would require computing the pairwise transport between each set of points on the surface [Peers et al. 2006]. A common approximation is to assume local homogeneity and fit to a simple model, such that the residual becomes an albedo texture [Jensen and Buhler 2002; Weyrich et al. 2006; Tariq et al. 2006]. The dipole, however, is a poor choice for representing materials with significant heterogeneity, as it approximates the complete transport between points with a single radially-symmetric profile. This produces artifacts in the resulting renderings, such as the overly-smooth shadow boundaries shown in Figure 9, bottom right.

4.1 Heterogeneity Through Inter-scattering

We begin by considering the multipole model, which assumes that materials are composed of *homogeneous* thin layers. A layer profile predicted by this model describes the light transport between points on the surface. We now assume that the properties of the layer are spatially-varying, though we constrain the variation to be slow relative to the mean free path of light. Because the properties are




Figure 3: *Heterogeneous light transport in our model. Light is scattered among multiple layers, each with spatially-varying parameters. It is also absorbed as it passes through infinitesimally-thin absorbing layers between the scattering layers.*

locally homogeneous, we approximate the exitant radiance (here shown for the case of reflectance from a two-layer material) as:

$$L(\vec{x}_o, \vec{\omega}_o) \approx \iint f_i(\vec{x}_i, \vec{\omega}_i) R_{12, \vec{x}_i}^+(\vec{x}_i, \vec{x}_o) L(x_i, \vec{\omega}_i) f_o(\vec{x}_o, \vec{\omega}_o) d\vec{\omega}_i d\vec{x}_i, \quad (5)$$

where f_i and f_o indicate the modulation of light at the surface of the layer due to Fresnel effects or any other BRDF. Note that R_{12, \vec{x}_o}^+ , the radially symmetric reflectance profile at the point of exitant radiance, is used to predict the contribution of nearby lit points.

In the homogeneous case, convolving two layers together to find their total transmittance is a radially-symmetric process and yields a homogeneous response. In our heterogeneous formulation, however, convolution of layer responses depends on the *local* position on the interface *between* the layers:

$$T_{12, \vec{x}_o}^+(x, y) = \iint T_{1, \vec{x}_i}^+(x', y') T_{2, \vec{x}_o}^+(x - x', y - y') dx' dy'. \quad (6)$$

Here, the forward transmittance profile of two heterogeneous layers at the point of exitance \vec{x}_o is T_{12, \vec{x}_o}^+ . It depends on the convolution of the second layer’s profile at \vec{x}_o , T_{2, \vec{x}_o}^+ , with the transmitted responses of the first layer (T_{1, \vec{x}_i}^+) over all points \vec{x}_i at the interface between the layers. Since there is a separate profile at each \vec{x}_i , the heterogeneities of each layer affect the transport of light *between* layers. This convolution of many different profiles produces a heterogeneous transport.

The incident flux on the lower layer takes the form of a set of points from different profiles, each profile having different shape. The resulting overall profile has a complex, piecewise structure (see Figure 4). Computing the overall response requires performing a full 2D convolution, a considerably more expensive process, as the efficient methods described in the previous section no longer apply.

Just as in the case of homogeneous layers, light that scatters from the top layer into the bottom layer may scatter back into the top layer. Accounting for these heterogeneous inter-scattering events requires evaluating (3) in 2D. For efficiency, we use the sum-of-




Figure 4: *The false-color images at left show homogeneous (top) and heterogeneous (bottom) profiles. The leftmost images show the result of convolving a constant 16×16 pixel irradiance with the top layer profiles, while the middle images show the result after convolving with both layers’ profiles. The plots at right show the center horizontal scanline of each image. Note that there are significant differences, even under this simple illumination.*

weighted-Gaussian representation of (4), as we will discuss in Section 4.3.

Although the above formulation does not capture heterogeneous light transport within a single layer, when a layer is thin, light does not spread very far before propagating to adjacent layers. Also, when the scattering of the layer is high, the mean free path of light is shorter within the layer. The gradients of optical properties within this layer are then small relative to the short mean free path. Thus, this approach of modeling layers is most accurate when layers with high parameter gradients are thin and highly scattering; those with lower-frequency variation may propagate light further with less error. Also note that, unlike with the dipole model, the transport between points on the surface now takes many paths between the layers, and thus involves using potentially different profiles.

4.2 Inter-scattering Absorption

As mentioned above, there may be localized areas of high absorption at the junction between skin layers, due to either natural or artificial pigmentation. We account for this by introducing infinitesimally thin absorbing layers between scattering layers. These simply attenuate light that scatters out of a scattering layer before it enters the next scattering layer. This complicates computing profiles, since the result of each convolution in (3) (except the last of each term) must be *multiplied* by a 2D attenuation map A . The resulting terms include both convolutions and products, and hence are no longer associative: $(A * B) \cdot C \neq A * (B \cdot C)$. Thus, (4) becomes:

$$\Phi * R_{12}^+ = \Phi * R_1^+ + \Phi * T_1^+ \cdot A * R_2^+ \cdot A * T_1^- + \Phi * T_1^+ \cdot A * R_2^+ \cdot A * R_1^- \cdot A * R_2^+ \cdot A * T_1^- + \dots, \quad (7)$$

where $*$ and \cdot have equal precedence and are evaluated left to right. The attenuation A is controlled by an absorbance parameter map P :

$$A(x, y) = e^{-P(x, y)}, \quad (8)$$

where greater absorbance allows less light to transmit through the layer.

Note that even without these absorbing layers, light transport in our model is heterogeneous due to inter-layer scattering, as we described in the previous subsection. Also, these absorbing layers can be thought of as similar to albedo textures at the surface of the skin: they allow precise control over the inter-scattering of light.

4.3 Efficient 2D Heterogeneous Convolution

Accounting for heterogeneities in our model requires performing the set of 2D convolutions and point-wise multiplications in (7). For arbitrary skin layers with complex changes in parameters, pre-computing and storing a profile per point over the entire surface would be prohibitively expensive. So, we perform the evaluations as needed, using graphics hardware to accelerate the computation.

Ideally, evaluating (7) requires performing a convolution of irradiance over an arbitrary surface. However, instead of implementing convolution over non-planar meshes, we perform all calculations on a 2D planar grid representing a local region of the surface. Specifically, to shade some point \vec{x}_o , we first determine a set of nearby points \vec{x}_i by taking a regular grid in a parameter space and mapping back to the surface according to the local UV parameterization. We collect the irradiance at the \vec{x}_i , multiplying by an area distortion metric so that each element represents the *total* irradiance contained within its area mapped on the surface.

We then perform the convolutions and multiplications indicated in (7). Following d’Eon et al. [2007], we represent diffusion profiles by a set of weights for a fixed set of Gaussians. We constrain the variances of the Gaussian set to be integral powers of an initial variance. This allows the blur of the next wider Gaussian to be computed from the results of the previous one, as successive blurs sum variances: $G(2v) = G(v) * G(v)$.




Figure 5: We blur irradiance using sums of weighted Gaussians that represent the heterogeneous profiles of the material. In this case, we illustrate the data flow of reflectance computation.

Given a profile R , represented as in (4), convolving it with the irradiance Φ is equivalent to blurring the irradiance by each Gaussian and accumulating the results:

$$\Phi * R = \sum_{i=1}^k w_i (\Phi * G(v_i)). \quad (9)$$

Because a Gaussian blur is separable, the cost of convolution is linear in the number of samples. Convoluting two profiles and accumulating for inter-layer scattering is performed similarly, through continued weighted blurring of the irradiance (see Figure 5).

Note that the discretization of the computation onto a grid requires a good parameterization of the surface being rendered: the UV coordinates of samples determine where irradiance is computed, how the parameter and attenuation maps are stored, and where final radiance is computed. Ideally, points that are optically close to each other in world space should be near each other in UV space, and vice versa.

4.4 Shading Thin Geometry

In the presence of thin geometry, the UV-distance between shading and lit points may have significant error. This is because points on opposite sides of a thin surface are likely to be far from each other in UV space. To account for this discrepancy, we project the shading point to the lit surface along a vector towards the light.

We use this projected point’s UV coordinates along with those of the irradiance samples to determine the local frame of the irradiance grid: the UV distance from irradiance samples to the projected point determines their location in the grid. We then perform convolutions to calculate the transmittance through the material. Because the geometry is thin, we construct a three layer material on-the-fly, consisting of a thin dermis layer bounded by two epidermis layers. We use the original and projected shade points to find the local material properties.

Finally, we weight the total contribution of reflectance and transmittance based on the normal at the point being shaded \vec{n}_s , and the average normal at the lit points \vec{n}_l [Donner and Jensen 2005]:

$$L = \frac{1}{2}(1 + \vec{n}_s \cdot \vec{n}_l) R + \frac{1}{2}(1 - \vec{n}_s \cdot \vec{n}_l) T. \quad (10)$$

4.5 Implementation Details

We have implemented our heterogeneous shading model in two forms. The first is a GPU-only version that is limited to simple geometry, shadows, and moderate texture resolution, but generates an image in under a second. For more general scenes we have also integrated our algorithm into a CPU-based Monte Carlo ray tracer. While the GPU version calculates the convolution of the shading points under all image pixels in parallel (assuming uniform directional irradiance), the CPU version samples the irradiance in a small area around the shading point, as in [Jensen et al. 2001].

Since separable Gaussian blurs are highly parallelizable, both renderers employ the GPU as a computational engine to facilitate efficient computation. We construct irradiance grids as described in Section 4.3, and upload them as 2D textures to the GPU. Recall that

each irradiance sample maps into a texel of the irradiance texture. As the CPU renderer becomes bound by the speed the GPU can convolve profiles, we group nearby shade points and process them together for efficiency. These points then share irradiance samples, which reduces the number of irradiance samples required per point. When rendering the surface using the CPU renderer, we use irradiance and convolution texture sizes of 128×128 . Higher resolutions gave no appreciable increase in quality.

We represent each profile as a set of 18 Gaussians, starting from a minimum standard deviation of 0.01 mm (the mean free path of light in the epidermis), with increasing powers of this value as discussed in Section 4.3. To efficiently determine the Gaussian weights for the optical properties at a particular point, we precompute lookup tables (LUTs) indexed by σ_a , as scattering is fixed in our skin model (see Section 5.1). Each LUT is 18×256 pixels, with each column giving the 18 Gaussian weights for a particular value of σ_a . For a given σ_a , we perform a non-linear least squares optimization to find the best set of Gaussian coefficients that approximate the resulting multipole profile. Each layer then has its own LUTs. We have found 256 logarithmically-spaced samples of σ_a to be sufficient for rendering skin.

5 Heterogeneous Spectral Shading of Skin

The most prominent chromophores in the epidermis are melanin and carotene. Melanin acts as optical protection from harmful UV radiation—it is typically increased by tanning—and dominates light absorption in the epidermis. As melanin is not a pure substance, we approximate it with blends of “eumelanin” and “pheomelanin”. Small amounts of the yellow/orange carotene [Sayre and Black 1992] are also found in the epidermis, often dependent on diet. In contrast, in the dermis the primary chromophore is hemoglobin, the bright-red chemical that carries oxygen in the blood. Figure 7 shows (normalized) absorption spectra of these chromophores—see [Jacques et al. 2001] for tabulated versions of these spectra.

Although in our model the epidermis and dermis are planar layers, in real skin the interface between them is strongly corrugated. In approximating the thickness of the epidermis, we effectively incorporate part of the dermis. For this reason, we include a small amount of hemoglobin in the epidermis layer. This inclusion also helps model the effects of increased blood flow, or erythema, which gives skin a reddish appearance. Since we include the effects of melanin in the absorbing layer between the epidermis and dermis, we have not found it necessary to incorporate melanin in the dermis.

Summarizing, we model skin as a two-layer translucent material, similarly to Donner and Jensen [2006]. However, our skin model incorporates six *spatially-varying* physiological parameters:

Parameter	Description	Typical range
C_m	Melanin fraction	0 – 0.5
β_m	Melanin type blend	0 – 1
C_{he}	Hemoglobin fraction (epi)	0 – 0.05
C_{bc}	Carotene fraction	0 – 0.05
C_{hd}	Hemoglobin fraction (dermis)	0 – 0.1
ρ_s	Oiliness	0 – 1

Table 1: Physiological parameters describing skin reflectance in our model. We use 2D parameter maps to define the parameters above over a surface. Using (11) and (12), these directly determine the optical properties (σ'_s and σ_a) of the skin layers.

5.1 Spectral Skin Properties

In order to use the above parameter maps to compute (spectral) reflectance and transmittance profiles, we must evaluate the total




Figure 6: Measurement setup. Left: Flash inside its case (back and top removed), facing the positioner ring with the gray card. Filter wheel visible at top of case. Right: View of camera and polarizers.

absorption in each layer. In the epidermis, the absorption is:

$$\begin{aligned} \sigma_a^{\text{epi}}(\lambda) = & C_m (\beta_m \sigma_a^{\text{em}}(\lambda) + (1 - \beta_m) \sigma_a^{\text{pm}}(\lambda)) \\ & + C_{\text{he}} (\gamma_e \sigma_a^{\text{oxy}}(\lambda) + (1 - \gamma_e) \sigma_a^{\text{deoxy}}(\lambda)) \\ & + C_b \sigma_a^{\text{bc}}(\lambda) + (1 - C_m - C_{\text{he}} - C_b) \sigma_a^{\text{baseline}}, \end{aligned} \quad (11)$$

where λ is the wavelength of light in nanometers, C_m is the total volume fraction of melanin, σ_a^{em} and σ_a^{pm} are the absorption coefficients for the two types of melanin (eumelanin and pheomelanin), and β_m blends between them. C_{he} is the total volume fraction of hemoglobin, γ_e controls the relative amounts of absorption from oxygenated σ_a^{oxy} and deoxygenated σ_a^{deoxy} hemoglobin. Finally, σ_a^{bc} is the absorption of carotene, and $\sigma_a^{\text{baseline}}$ is a baseline absorption of the remaining skin tissue.

The total spectral dermis absorption is :

$$\begin{aligned} \sigma_a^{\text{derm}}(\lambda) = & C_{\text{hd}} (\gamma_d \sigma_a^{\text{oxy}}(\lambda) + (1 - \gamma_d) \sigma_a^{\text{deoxy}}(\lambda)) \\ & + (1 - C_{\text{hd}}) \sigma_a^{\text{baseline}}, \end{aligned} \quad (12)$$

where γ_d controls the oxygenation of the dermal hemoglobin. Note that we fix γ_e and γ_d to be 0.7, a typical value.

We fix the reduced scattering coefficient σ_s' of light in skin in our model: in the epidermis, we describe it by the following power law [Jacques 1998]:

$$\sigma_s'(\lambda)^{\text{epi}} = 14.74\lambda^{-0.22} + 2.2 \times 10^{11} \times \lambda^{-4}, \quad (13)$$

while in the dermis we reduce this by half. We take the thickness of the epidermis to be constant at 0.25 mm, while the dermis is considered to be semi-infinite (for thick surfaces — our approach for thin structures such as ears is described below). Also, we fix the indices of refraction of the two layers to be 1.4 and 1.38.

We make two final additions to the two-layer model. First, to account for surface reflectance, we use the Torrance-Sparrow BRDF [1967] and scale its contribution by an “oiliness” factor ρ_s , as described by Donner and Jensen [2006]. Second, as mentioned above, we assume that some fraction of the melanin is concentrated between the epidermis and dermis. Based on our empirical observations, we set this additional absorption to be 17.5% of the total absorption of melanin in the epidermis. We also account for other localized absorption (e.g. from veins) with this absorbing layer, though it must be manually added (such as in Figure 1).

This two-layer model, with spatially-varying absorption, may now be used in the rendering procedure described in Section 4.3. To render thin materials, such as ears, we determine the actual thickness of the dermis as the average distance from the shade point \vec{x}_o to contributing lit points, as described in Section 4.4. We then construct a *three-layer* material, containing a top epidermis, a thin inner dermis, and a bottom epidermis. We compute transmittance profiles using this three-layer configuration. Note that this involves multiplication by both the front- and back-side absorption layers.

6 Reflectance Measurements

Our heterogeneous model has five parameters that control subsurface scattering (ρ_s controls surface reflection). One option for obtaining parameters for the model would be to turn to the medical literature: the optical properties of skin have previously been measured and tabulated [van Gemert et al. 1989; Tuchin 2000]. These tabulations, however, were not acquired with graphics applications in mind, but rather were intended to define parameter ranges that characterize healthy and pathological tissue. Choosing good values to achieve a desired look therefore remains a non-trivial task.

To avoid manually picking model parameters, previous work has inferred model parameters from RGB textures of skin [Jensen and Buhler 2002; Donner and Jensen 2006], with additional measurements to constrain the fit [Weyrich et al. 2006]. As our model comes with an increased number of parameters, these methods would be difficult to apply. Instead, we describe a new skin measurement protocol to derive representative parameters for our model.

6.1 Design

Our design is inspired by the acquisition device used by Jensen et al. [2001]: a digital SLR camera (Canon EOS 20D) observes a skin sample under orthogonal incident illumination; the camera view is 30° off-orthogonal to avoid retro-reflection. A photographic xenon flash (Canon 580 EX II) serves as a light source. In addition, crossed polarizers in front of flash and camera largely eliminate surface reflection. See Figure 6 for a schematic.

With five degrees of freedom at each surface point, it is not possible to restrict measurements to three (RGB) channels alone, as has been done by all previous skin studies in graphics. Instead, to allow for discrimination of the model constituents, we measure at multiple, narrow frequency bands. To this end, we place narrow band-pass filters (10 nm-wide band, Cheshire Optical and Newport® Optics) in front of the flash. At a 165 mm (6.5”) distance to the skin sample, the flash’s output modulated by the filters’ spectra still provides enough radiant energy for imaging, while the interference filters’ frequency shift of non-orthogonal rays stays below 0.1 nm within a volume of interest of 29 mm (1.14”) diameter. Flash and filters, mounted in a filter wheel, are contained in a fixed enclosure that eliminates stray light. Skin samples are positioned at a fixed location in front of the setup, with a blackened metal ring around this location aiding proper positioning.

We selected the filter bands based on simulations using our skin model; we performed an exhaustive search for filter combinations that provided an optimal signal-to-noise ratio with respect to the intrinsic model parameters. Note that due to wavelength-dependent scattering in skin, it is not sufficient to solely consider the constituents’ absorption spectra; the complete reflectance model must be evaluated. Our choice of filters also avoids the ragged region around 420 nm in the spectrum of the xenon flash, even though this region might otherwise have been helpful in discrimination. Figure 7 shows the filter spectra used, overlaid with the spectra




Figure 7: Filters that optimally separate the model constituents, overlaid with the absorption of chromophores used in the model.

of the model chromophores. For maximum irradiance control we structure the incident illumination by affixing opaque black plastic tape onto the skin, with a 7.5 mm square window cut out of the tape. The camera observes the skin through this window, and the boundary of the cut-out region serves as an occluder modulating the incident illumination by a box function¹.

For each filter, we acquire a reference image of a Spectralon™ reflectance target placed at the same location as the skin samples. Relating subsequent measurements to these reference images calibrates for incident flux Φ , for camera vignetting, and for the camera’s spectral sensitivity. To track changes in the flash intensity, a 30% gray card (x-rite) is included in the setup, such that it is simultaneously lit by the light source and observed by the camera throughout all measurements.

6.2 Measurement Procedure

We position the taped skin sample in front of the measurement apparatus, taking a series of images with the filter wheel advanced after each shot (f/10, 1/100 sec, ISO 200–800 depending on wavelength). We also acquire and subtract off a black image, using the same setup but with an opaque inset in the filter wheel.

The cross-polarization filters out glossy surface reflection and single-scattering events. Surface reflectance, however, also affects the amount of light entering and exiting the skin and therefore has an effect on the observed diffuse reflectance [Donner and Jensen 2005]. Jensen et al. [2001] assume a smooth surface in their measurements, so that surface transmittance can be described by a transmissive Fresnel term $F_t(\vec{\omega}, \eta) = 1 - F_r(\vec{\omega}, \eta)$. For skin, with its varying oiliness ρ_s , transmittance can differ significantly from F_t : for our measurement geometry, typical dry skin (Torrance-Sparrow roughness $\rho_r = 0.35$) would have 28% higher transmittance and 78% lower reflectance than a smooth surface. We eliminate this measurement uncertainty by applying a thin film of ultrasound gel (Aquasonic® Clear) to the skin. This is a clear gel, with an index of refraction close to water, that is designed to adhere closely to skin. The gel remains glossy throughout the measurement procedure and removes the influence of the skin’s surface BRDF by inducing the transmittance described by a smooth Fresnel term.

Before further processing, we identify the cut-out region in each image separately, and resample the captured images within it. This way, we implicitly compensate for potential subject motion between images. A final image alignment step further improves correspondence between the different wavelength components.

6.3 Model Fit

The stack of images acquired from a skin patch represents a single, multi-spectral image of that patch. We estimate model parameters for each multi-spectral pixel using inverse rendering [Yu et al. 1999]. Starting from average skin parameters, a gradient-descent optimization repeatedly renders the skin patch and alters the model parameters to minimize the differences between the rendering and the acquired image. The optimization uses a multi-spectral GPU implementation of our layered, heterogeneous skin model, as described in Section 4.3. At each step of the optimization, the model includes an absorbing layer between the epidermis and dermis which contains 17.5% of the predicted melanin concentration. Additional absorption not directly predicted by the model, such as from veins, becomes part of the residual of the fit. Accounting for this would require manually editing the parameter and absorption maps, but we have chosen not to do so in our fits.

The inverse renderer simulates the incident illumination, modulated by the black-tape occluder, and computes radiant exitance

at each pixel for each measured wavelength. After each gradient descent iteration, we copy nearby skin parameters to occluded regions. We simulate a one-centimeter-wide boundary of the occluded region to accurately account for inter-scattering between visible and occluded regions. The gradient descent simultaneously optimizes the parameter vector at each pixel, estimating the objective function’s gradient using forward-differences. By simultaneously evaluating all forward differences (by rendering the full patch), the scattering cross-talk between surface points implicitly couples the independent optimizations, leading to a consistent solution of heterogeneous model parameters. The strong coupling between pixels due to the scattering profiles’ large support leads to rapid convergence, typically within 400 iterations.

6.4 Error Analysis

Great care has been taken to radiometrically calibrate all components of the setup. We found that special attention had to be paid to effects due to imperfect depolarization. From systematic measurements we derive that Spectralon’s reflectance varies by $\pm 5.27\%$ between s- and p-polarization of the camera’s polarizer, which must be taken into account when radiometrically calibrating the system. In addition, we find that scattering in skin depolarizes even more imperfectly than Spectralon: on a gel-treated skin sample oriented obliquely to the camera (to avoid glossy reflection), we observe a polarization-dependent intensity variation of $\pm 20.0\%$ around the average diffuse reflectance. We include these biases in our calibration, unlike previous work in graphics and vision that assumes that cross-polarization transmits exactly 50% of diffuse reflectance.

The radiometric accuracy has direct influence on the convergence of the inverse rendering; errors in scale slow down convergence and increase the residual. We achieve optimal convergence by scaling our data by 1.02, a small correction factor that indicates a fairly good radiometric calibration. On 256×256 pixel patches, the optimization typically converges to a residual below 3% RMS reflectance. We believe that this residual is due to camera noise and to effects in the tissue that are not simulated by our model.

7 Results

Given parameter maps that define the model parameters in Table 1, Equations (11)–(13) directly define the scattering and absorption properties of the skin layers. These then let us define the Gaussian coefficients which represent each layer’s reflectance and transmittance profiles. We then use the multipole model to convolve irradiance with profiles as described in Section 4.5. Note that unless stated otherwise all images rendered in this paper use a surface roughness of 0.35.

The renderings in this section were performed on a Dual-Core Intel Core 2 Duo 3.0 GHz machine with one NVIDIA GeForce 8800 GTX and one GeForce 8800 GT (G92) graphics card. The GPU renderer computes a multi-spectral image on a single GPU in about a second, while the CPU version takes under 15 minutes per image.

We render images using 8 spectral samples, at wavelengths 400, 435, 470, 505, 540, 575, 610, and 645 nm. We have found this set to give results within 1% error in final RGB color, relative to a reference using 151 samples spaced every 2 nm. Note that this is a different set of wavelengths than used for acquisition, as the latter is optimized to expose the differences between model constituents. To obtain RGB color values after rendering with the spectral model, we convert from spectral to XYZ using the CIE 2° Standard Observer color matching functions \bar{x} , \bar{y} , \bar{z} [CIE 1931].

In our implementation, we compute the first 5 terms of (7) for blue wavelengths, and 6 terms at green and red wavelengths. This is sufficient to compute the inter-scattering of light between layers while preserving over 99% of the profiles’ spectral energy. To limit the computational cost when rendering thin

¹ Note that this allows us to observe part of a line convolution of the scattering profile along the tape edges, where skin areas closer to the occluding tape observe less light through scattering from nearby regions.




Figure 8: Changing skin parameters with mechanical deformation. The example shows a reduction of blood flow (hemoglobin concentration) after clenching and releasing the hand. The parameter maps (scaled 20x) for hemoglobin are shown in the insets.

three-layered geometry, we compute only terms with at most two $R^+ * R^-$ factors for each adjacent pair of layers, and only terms with forward transmittance profiles. These terms take the form: $T_1^+ * (R_2^+ * R_1^-)^{*0-2} * T_2^+ * (R_3^+ * R_2^-)^{*0-2} * T_3^+$. The exponent denotes the number of inter-scattering convolution factors (from 0 to 2).

7.1 Using the Model

To demonstrate the practicality of our model, we manually created chromophore maps and an absorbing layer for a model of the palm of a hand [Ma et al. 2007], as shown in Figure 1. Though the model is generated from high-resolution normal data, and this contributes to the realism of the rendered image, the color, luminance, and contrast come from the underlying shading model.

Modifying parameter maps: The parameter maps (scaled in intensity to show detail) shown in Figure 1 correspond to epidermal melanin, inter-layer absorption due to blood vessels at the junction between the skin layers, and dermal hemoglobin (no epidermal hemoglobin is used). The contribution of melanin to the absorbing layer is not shown. Note the simplicity of the maps; the complex interaction of light is captured by the heterogeneous layered model. These effects would be difficult to achieve using albedo maps. Direct access to the physical parameters allows realistic rendering of changes of the skin condition. Figure 8 shows the same model before and after undergoing a possible mechanical deformation (e.g. clenching the hand into a fist and then releasing), which forces blood to flow away from the deformed areas. We simulate this effect by changing the concentration of hemoglobin in our skin model. The hemoglobin parameter maps are inset into the rendered images. The values of the parameter maps have been scaled to better detail their differences.

Heterogeneous vs. homogeneous: Figure 9 demonstrates the differences in scattering between homogeneous and heterogeneous profiles. We compare renderings of a flat patch of skin, using Donner and Jensen’s [2006] homogeneous model with average homogeneous parameters (first row), and our model with heterogeneous (third row) parameters. The second row adds an albedo texture to the homogeneous model derived by dividing the heterogeneous image by the homogeneous one. Note that creating the albedo texture by hand would be difficult even for a skilled artist, and in this case required first rendering the heterogeneous image. In the second column, the patches are lit with non-uniform, structured illumination using strips of light. Though the appearance seems similar, when lit with brighter strips (third and fourth columns) clear heterogeneities are visible using spatially-varying parameters, while the albedo-mapped image retains its homogeneous characteristics.

Thin material transmission: Figure 10 shows an image of a back-lit ear transmitting light, with parameter and absorption maps adding the effect of internal veins and freckles. Here we sample

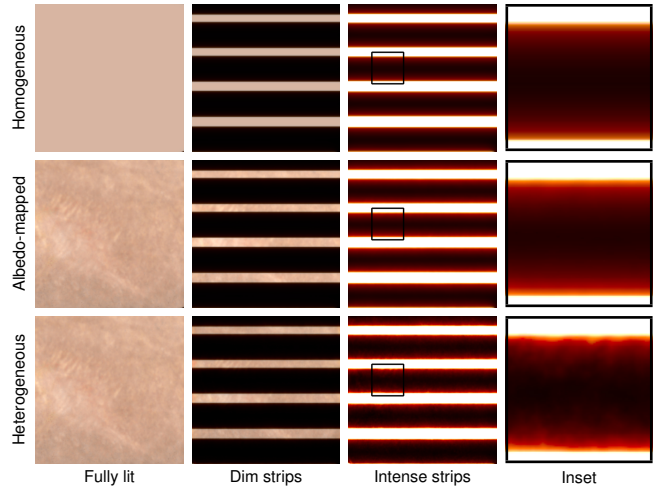


Figure 9: Renderings of skin with homogeneous parameters (top), homogeneous parameters modulated by an albedo texture (middle), and heterogeneous parameters using our method (bottom). The left column shows fully lit images, while the remaining columns are lit by strips of light of different intensity.




Figure 10: A back-lit ear rendered with absorbing layers. Transmittance is calculated using the multipole model, not an approximate transmittance using a reflectance profile. Veins and freckles are controlled by parameter maps over the model geometry.

the illumination normally, but project the shade points to the lit side of the ear to obtain the convolved irradiance as described in Section 4.4.

7.2 Skin Measurements

To evaluate our model using actual chromophore concentrations of real skin samples, we measured 39 skin patches across different




Figure 11: Measured skin patch (top) compared to its model fit (bottom) at each of the nine measured wavelengths, and in RGB (right).

Skin Type	n	Interior Arm				
		C_{he}	C_{hd}	C_m	β_m	C_{bc}
I	1	0.000184	0.004133	0.000126	0.0238095	0.000240
II	1	0.002971	0.007399	0.00571	0.833975	0.000010
III	2	0.001441	0.004674	0.005364	0.151566	0.000081
IV	2	0.002359	0.004503	0.013851	0.11501	0.001105

Skin Type	n	Exterior Arm				
		C_{he}	C_{hd}	C_m	β_m	C_{bc}
I	1	0.001933	0.005001	0.002346	0.0767263	0.001543
II	1	0.000342	0.004226	0.002718	0.894776	0.000000
III	1	0.001766	0.004072	0.009446	0.143447	0.000113
IV	1	0.003252	0.006598	0.019121	0.137964	0.002384
V	1	0.006558	0.012924	0.038484	0.0603627	0.006445

Table 2: Chromophore distributions for different skin types.

skin types, conditions, and regions of the body. A measurement takes roughly 20 seconds; the limiting factor is the manual advancing of the filter wheel. The skin patches cover roughly 500×500 pixels in the camera image. After bilinearly resampling them to 512×512 , we finally downsample (box-filter) the patch to 256×256 for inverse rendering, which provides moderate noise reduction.

Reconstructions: Figure 11 shows a patch of skin that has been reconstructed with our acquisition procedure. Shown are the nine acquired single-wavelength images and renderings of the reconstructed skin model. The rendering features a smooth surface, mimicking the effect of gel application. Even though we re-stage the conditions at acquisition time, it is not possible to directly compare the reconstruction to an RGB image taken under the same conditions: correctly displaying such an image in sRGB would require a full spectral calibration of the flash and the camera response, which we avoid in our setting. Figure 12 (a) shows the corresponding eumelanin, pheomelanin, and hemoglobin distributions in this patch.

Note that the chromophore distributions in Figure 12 show distinct features such as scars and freckles separated into different layers. Also, although Figure 12 (b) shows some hemoglobin within the freckle, it has a distinctly different structure from the melanin distributions. Figure 13 shows additional reconstructions, demonstrating the variation within single subjects and across skin types, with selected coefficient maps in Figure 12 (a–d).

Chromophore variation: For further analysis, we eliminate outlier pixels within each patch reconstruction by discarding pixels where one of the coefficients falls outside the 5- and 95-percentiles; all numbers refer to the full set of 39 scans. First, we concentrate on skin type. Of subjects that cover the Fitzpatrick skin types I through V [Fitzpatrick 1988], we measured patches on the lower arm, both on the inside and on the more strongly pigmented outside. This provides us with the skin-type dependent average concentrations shown in Table 2. Unfortunately, most of the statistics are with respect to a single subject each (column n). Differences between skin types are quite noticeable and, as to be expected, mainly established in the melanin concentration.

We also scanned multiple skin patches of single subjects to analyze intra-subject variations of skin appearance. Our initial expect-

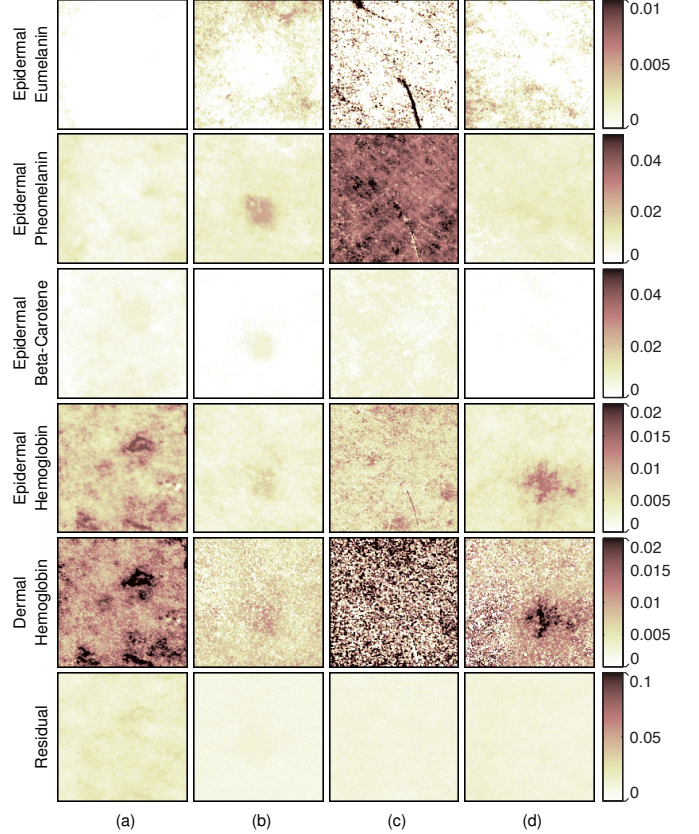


Figure 12: Chromophore concentrations as derived from our measurements, corresponding to Figure 13 (a–d).

tation to reveal correlations within a subject’s samples, potentially allowing for a subject-specific model with fewer parameters, was not met. It turns out that intra-subject variations can be larger than inter-subject variations measured at corresponding points. Even within a single skin patch, PCA analysis across coefficients still requires a 4-dimensional basis to reach an average 98.7% accuracy.

Model validation: To validate the model and its fit, we analyze correlations between the different model chromophores across all pixels in all patches. In accordance with the general gradient of hemoglobin in skin, there is a strong correlation between dermal and epidermal hemoglobin (correlation coefficient 0.87). However, the ratio between the per-patch average of C_{hd} and C_{he} ranges from 1.7 and 2.8 in most scans. Thus, we argue that despite the high correlation, this variability does not allow us to reduce the model by tying both C_{hd} and C_{he} to a single model parameter. Other general correlations among skin chromophores are not known. Indeed, we find only negligible correlation between melanins (0.034 between all pixels; 0.12 between mean concentration of patches). There is, however, a slight correlation among all reddish chromophores (hemoglobins, pheomelanin, beta-carotene), indicating that their




Figure 13: Example skin patch reconstructions (a,b,d,e,h,i) Caucasian, skin type III: bridge of the nose (a), freckle on arm (b) slight acne on arm (d), forehead (e), stronger acne on arm (h), scar on edge of the hand (i). (c,g) African, skin type V: exterior lower arm (c) (artifact due to hair), wrinkles at posterior side of wrist (g). (f) Caucasian, skin type II, exterior lower arm. (j) Asian-subcontinental, skin type IV, scar on back of hand. (k,l) Caucasian, red-haired, skin type I: exterior (k) and interior (l) lower arm.

separation is imperfect. In most cases, we observe unusually high pheomelanin-to-eumelanin ratios, which has to be attributed to this effect. We believe that a better separation would be possible by measuring far blue wavelengths, which unfortunately would be difficult to achieve with off-the-shelf light sources and cameras.

Chromophore editing and synthesis: Figure 14 shows examples of the direct access to skin conditions possible with our model. We synthesize new skin appearance by using the hemoglobin coefficients from one patch to replace hemoglobin in another patch. This shows how coefficient maps can independently be edited to achieve a desired effect. We further demonstrate the clear separation of chromophores by scaling hemoglobin and melanin concentrations independently. Note how our acquisition method was able to pick up the delicate structure of the faint melanin distribution in 13 (i).

The images in Figure 15 use texture synthesis to directly synthesize skin appearance from the acquired chromophore distributions. We use image quilting [Efros and Freeman 2001], to synthesize parameters over the surfaces. One of the original datasets used is shown in Figure 13 (d). In the left image we model a tattoo using an absorbing layer between the epidermis and dermis: the pigments of a tattoo are represented using eight absorption maps, one for each spectral sample. Scattering and absorption in the epidermis give depth to the tattoo and the veins in the right image. This also causes them to acquire a bluish tint.

8 Discussion

Spectral vs. RGB model: Our shading model renders images by simulating light transport in 8 spectral bands. We handle this with two separate shading passes on the GPU, where each pass processes 4 spectral bands in parallel. In many cases these spectral computations can be limited to the skin shader only, by tagging irradiance samples per light, without significantly modifying the entire rendering pipeline. Unfortunately, as shown by Donner [2006][Chapter 8.5.1], the non-linearity of the multipole model prevents a simple transform from a spectral to an RGB model, and RGB profiles introduce significant errors due to compression of the color space.

Design decisions: In designing our measurement system, we made several decisions to simplify the parameter optimization. In




Figure 14: Left: Assigning the hemoglobin distribution of the patch from Figure 13 (h) to 13 (f). Right: Increasing blood flow (left) and tan (right) of a sample with scar tissue. Note how capillaries do not reach scar tissue, while melanin manages to diffuse into it.

particular, we chose to fix the scattering coefficient σ_s' . In contrast to other potential constraints, fixing σ_s' does not affect the mean free path of light; absorption heterogeneities lead to realistic variations in how far light "bleeds" into shadows (see Figure 9). Physiologically, our choice of chromophores is fairly common in both graphics and biology. We do, however, implicitly assume that the chosen spectra (taken from measured data of real tissues) and an isotropic scattering model of two plane-parallel layers can accurately describe light transport in skin.

Fit convergence: All of our acquired datasets converged to a sufficiently low residual after 200 to 500 iterations. Convergence appears to be independent from the choice of initial values, but starting closer to the desired skin type speeds up convergence. For the final data fits we report, we arbitrarily chose the following initial values: epidermal eumelanin 0.73%, pheomelanin 0.3%, hemoglobin 0.3%, beta-carotene 0% and dermal hemoglobin 0.6%, which corresponds to very pale Caucasian skin.

9 Conclusion and Future Work

In this paper we introduced a spectral, layered shading model for heterogeneous skin. Our model captures these heterogeneities via inter-layer scattering, and accounts for spatially-varying changes in optical properties, as well as inter-layer absorption. We have demonstrated our model's ability to capture the heterogeneous appearance of skin, and to reproduce a wide range of skin appearance over various skin types and conditions. In particular, the use of simple and intuitive parameter maps allows our model to simulate complex visual phenomena that would be difficult to achieve using only albedo maps.



Figure 15: Left and right: Texture synthesis from scanned data to propagate an appearance across a surface. Left: An absorptive tattoo texture (inset) between epidermis and dermis reveals "depth".

We presented a measurement procedure that is capable of retrieving model parameters from living skin, thereby providing realistic, spatially-varying chromophore concentrations. Analyzing the measurements, we conclude that our skin model cannot be significantly simplified while maintaining physical quantities relevant for skin's visual appearance. Because wavelength restrictions have led to a slight correlation between reddish chromophores, the recovered concentrations may not be suitable for medical diagnosis; nonetheless, we feel that they are accurate enough for applications in computer graphics. The procedure itself is general enough to be applied to layered materials other than skin.

In summary, we believe that our rendering model represents an attractive tradeoff between variability, visual accuracy, and practicality. It provides access to its static and dynamic appearance parameters, and faithfully renders material heterogeneities, while being more efficient than a full simulation of scattering in heterogeneous materials.

So far, we have only demonstrated our acquisition method on small patches of skin. Acquiring multi-spectral measurements of whole body regions is conceivable and would require a larger light source and a geometric reconstruction for radiometric calibration. The main challenge would be to handle surface roughness: the currently used gel application does not scale practically. Future work may also evaluate the performance of our model on other layered materials. We also believe that the model can be used as an approximation for general heterogeneous materials by discretizing them into layers.

10 Acknowledgements

This work was supported in part by the NSF (grants IIS 03-25867, CCF 03-47427, CCF 04-44916, CCF 05-41259, CCF 07-01775, and CCF 07-02580), a Sloan Research Fellowship, and an ONR Young Investigator award N00014-07-1-0900. We also acknowledge generous equipment donations by NVIDIA. Many thanks are due to Ryan Overbeck for help in creating parameter maps, and to Xiaojuan Ma for assistance in scanning skin patches. Thank you also to Paul Matts for invigorating discussions and allowing the use of skin cross-section images. We acknowledge the helpful and constructive comments from the Tiggraph and SIGGRAPH reviewers. The ear is a section of a high-resolution scanned model by XYZRGB, the palm model is provided generously to the community by Ma et al. [2007], and the hand model is a part of the AIM@Shape Repository.

References

BERNSTEIN, E. F. 2006. Laser treatment of tattoos. *Dermatol. Clin.* 24, 43–55.

BORSHUKOV, G., AND LEWIS, J. P. 2003. Realistic human face rendering for “The Matrix Reloaded”. In *ACM SIGGRAPH Sketches and Applications*, ACM, 1.

CIE. 1931. In *Commission Internationale de l’Eclairage Proceedings*.

COTTON, S. D., CLARIDGE, E., AND HALL, P. N. 1999. A skin imaging method based on a colour formation model and its application to the diagnosis of pigmented skin lesions. In *Proceedings of Medical Image Understanding and Analysis '99*, 49–52.

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.

DANA, K. J., VAN GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and texture of real-world surfaces. In *Proceedings of ACM SIGGRAPH 1999*, vol. 18, 1–34.

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

D’EON, E., LUEBKE, D., AND ENDERTON, E. 2007. Efficient rendering of human skin. In *Rendering Techniques*, 147–157.

DONNER, C., AND JENSEN, H. W. 2005. Light diffusion in multi-layered translucent materials. *ACM Trans. Graphic.* 24, 3, 1032–1039.

DONNER, C., AND JENSEN, H. W. 2006. A spectral BSSRDF for shading human skin. In *Rendering Techniques*, 409–417.

DONNER, C. 2006. *Towards Realistic Image Synthesis of Scattering Materials*. PhD thesis, University of California, San Diego.

EFROS, A. A., AND FREEMAN, W. T. 2001. Image quilting for texture synthesis and transfer. In *Proceedings of ACM SIGGRAPH 2001*, 341–346.

FITZPATRICK, T. 1988. The validity and practicality of sun-reactive skin times I through VI. *Arch. Dermatol.* 124, 6, 869–871.

GEORGHIADES, A. S., BELHUMEUR, P. N., AND KRIEGMAN, D. J. 1999. Illumination-based image synthesis: Creating novel images of human faces under differing pose and lighting. In *Proceedings of the IEEE Workshop on Multi-View Modeling and Analysis of Visual Scenes (MVIEW '99)*, 47–54.

GHOSH, A., HAWKINS, T., PEERS, P., FREDERIKSEN, S., AND DEBEVEC, P. 2008. Practical modeling and acquisition of layered facial reflectance. *To appear in ACM Trans. Graphic.* 27.

GOESELE, M., LENSCH, H. P. A., LANG, J., FUCHS, C., AND PETER SIEDEL, H. 2004. DISCO: Acquisition of translucent objects. *ACM Trans. Graphic.* 23, 3, 835–844.

HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. In *Proceedings of ACM SIGGRAPH 1993*, 164–174.

HAWKINS, T., WENGER, A., TCHOU, C., GARDNER, A., GÖRANSSON, F., AND DEBEVEC, P. 2004. Animatable facial reflectance fields. In *Rendering Techniques*, 309–320.

HERY, C. 2003. Implementing a skin BSSRDF. *ACM SIGGRAPH 2003 Course 9*, 73–88.

IGARASHI, T., NISHINO, K., AND NAYAR, S. K. 2005. The appearance of human skin. Technical Report CUCS-024-05, Columbia University.

JACQUES, S. L., PRAHL, S., AND LINDSEY, J. 2001. Optical properties spectra. *Oregon Medical Laser Center*, <http://omlc.ogi.edu/spectra>.

JACQUES, S. L. 1996. Origins of tissue optical properties in the UVA, Visible, and NIR regions. In *OSA TOPS on Advances in Optical Imaging and Photon Migration*, vol. 2, 364–371.

JACQUES, S. L. 1998. Skin optics. *Oregon Medical Laser Center News*, <http://omlc.ogi.edu/news/jan98/skinoptics.html>.

JENSEN, H. W., AND BUHLER, J. 2002. A rapid hierarchical rendering technique for translucent materials. *ACM Trans. Graphic.* 21, 576–581.

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

KRISHNASWAMY, A., AND BARANOSKI, G. V. G. 2004. A biophysically-based spectral model of light interaction with human skin. In *Proceedings of EUROGRAPHICS 2004*, vol. 23.

MA, W.-C., HAWKINS, T., PEERS, P., CHABERT, C.-F., WEISS, M., AND DEBEVEC, P. 2007. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Rendering Techniques*, 183–194.

- MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., TORRANCE, K. E., AND GREENBERG, D. P. 1999. Image-based BRDF measurement including human skin. In *Rendering Techniques*, 119–130.
- MATTS, P., DYKES, P., AND MARKS, R. 2007. The distribution of melanin in skin determined in vivo. *British Journal of Dermatology* 156, 4, 620–628.
- PEERS, P., VOM BERGE, K., MATUSIK, W., RAMAMOORTHY, R., LAWRENCE, J., RUSINKIEWICZ, S., AND DUTRÉ, P. 2006. A compact factored representation of heterogeneous subsurface scattering. *ACM Trans. Graphic.* 25, 3, 746–753.
- SAYRE, R. M., AND BLACK, H. S. 1992. Beta-carotene does not act as an optical filter in skin. *J. Photochemem. Photobiol. B: Biol.* 12, 83–90.
- STAM, J. 2001. An illumination model for a skin layer bounded by rough surfaces. In *Rendering Techniques*, 39–52.
- TARIQ, S., GARDNER, A., LLAMAS, I., JONES, A., DEBEVEC, P., AND TURK, G. 2006. Efficient estimation of spatially varying subsurface scattering parameters. In *Vision, Modeling, and Visualization*.
- TONG, X., WANG, J., LIN, S., GUO, B., AND YEUNG SHUM, H. 2005. Modeling and rendering of quasi-homogeneous materials. *ACM Trans. Graphic.* 24, 3, 1054–1061.
- TORRANCE, K., AND SPARROW, E. 1967. Theory for off-specular reflection from roughened surfaces. *J. Opt. Soc. Am.* 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 Trans. Graphic.* 22, 3, 770–779.
- TUCHIN, V. 2000. *Tissue Optics: Light Scattering Methods and Instruments for Medical Diagnosis*. SPIE Press.
- VAN GEMERT, M. J. C., JACQUES, S. L., STERENBORG, H. J. C. M., AND STAR, W. M. 1989. Skin optics. *IEEE Trans. Biomed. Eng.* 36, 12, 1146–1154.
- WANG, J., ZHAO, S., TONG, X., STEPHEN LIN, Z. L., DONG, Y., GUO, B., AND SHUM, H. 2007. Modeling and rendering of heterogeneous translucent materials using the diffusion equation. Technical Report MSR-TR-2007-59, Microsoft Research.
- WENGER, A., GARDNER, A., TCHOU, C., UNGER, J., HAWKINS, T., AND DEBEVEC, P. 2005. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Trans. Graphic.* 24, 3, 756–764.
- WEYRICH, T., MATUSIK, W., PFISTER, H., BICKEL, B., DONNER, C., TU, C., MCANDLESS, J., LEE, J., NGAN, A., JENSEN, H. W., AND GROSS, M. 2006. Analysis of human faces using a measurement-based skin reflectance model. *ACM Trans. Graphic.* 25, 1013–1024.
- YU, Y., DEBEVEC, P., MALIK, J., AND HAWKINS, T. 1999. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *Proceedings of ACM SIGGRAPH 1999*, 215–224.