A Heterogeneous Reflectance Model Applied to Human Skin: Rendering and Acquisition

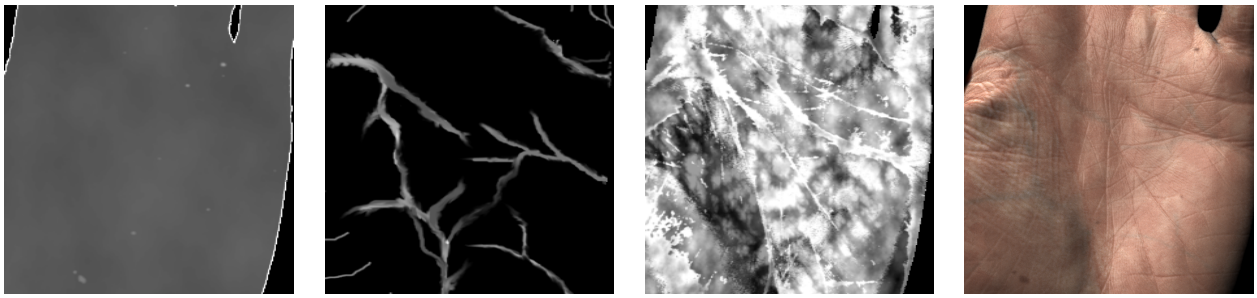

Melanin (top layer) | Inter-layer absorption | Hemoglobin (bottom layer) | Final image

Figure 1: *Our model approximates light transport in heterogeneous materials through the inter-scattering of light between layers. Here we show the spatially varying parameters of the top and bottom scattering layers, and the infinitesimally thin absorbing layer between them. Only these three parameter maps were used to generate the rendered final image on the right; there are no albedo textures.*

Abstract

We introduce a heterogeneous, spectral, layered reflectance model for translucent materials. Although the model is general, in this paper we apply it to faithfully reproducing the complex variations in skin pigmentation. Our reflectance model captures heterogeneity through the inter-scattering of light between layers. Each layer has an independent set of spatially-varying parameters, and we allow arbitrary variations in parameters over the surface of a layer. We assume that the local properties of each layer are slowly varying with respect to the mean free path of light. For greater physical accuracy and control, we introduce an infinitesimally thin absorbing layer between scattering layers. To obtain parameters for our model, we describe a novel acquisition method to infer heterogeneous model parameters from multi-spectral photographs of skin patches. We use an inverse rendering technique to globally optimize the best set of parameters for each pixel of the patch. We also show that our method finds close matches to a wide variety of inputs with low residual error. We demonstrate our heterogeneous reflectance model by rendering complicated skin visual effects such as veins, tattoos, rashes, and freckles, without the use of albedo textures. Also, by varying the parameters to our model, we achieve effects from external forces, such as external pressure changing blood flow within the skin.

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, realistic image synthesis

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 for acquiring, skin's appearance, it is imperative that we understand and 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 quite complex due to its many chemical constituents and structural variation [van Gemert et al. 1989]. Existing models used for skin, 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].

Our main contribution is a reflection model for rendering heterogeneous multi-layered translucent materials. Although our model is general, in this paper we apply it towards capturing and reproducing the heterogeneity of skin appearance. Our model is based on recent work on multi-layered translucent materials [Donner and Jensen 2005]. The key difference, however, is rather than 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, as light scatters between layers, we apply the spatially-varying properties of the layers. This allows our model to simulate complex, heterogeneous visual effects without albedo textures.

Though it can be driven by arbitrary parameter maps, to facilitate the use of our model we describe a novel technique for directly recovering model parameters from multi-spectral images of real skin patches using inverse-rendering. Our method is accurate and versatile enough to capture a number of physiological properties over a wide range of body regions.

We evaluate our methods by comparing our obtained parameters to acquired data, and by rendering realistic novel images of skin without albedo textures, such as in Figure 1. We also demonstrate how to simulate physiological changes in skin by modifying model parameters (see Figure 8).

2 Previous Work

The reflectance of translucent materials such as skin has received considerable attention in computer graphics. Here we focus on the work related to reflectance measurement and shading models, as they relate to our proposed methods. For a more thorough survey of techniques, see 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].

As they rely on 1D approximations to light transport, the former 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. As a successful compromise between accuracy and complexity, models based on the diffusion approximation have become popular. 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 have been developed to deploy the dipole model in rendering spatial surface variations. Jensen and Buhler [2002] and Héry [2003] fix one parameter (translucency) of the model and vary diffuse reflectance to match a given albedo texture. This approach, however, assumes that the transport of light between 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 seen in Section 4. Others have employed the dipole model with fully 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.

Donner and Jensen [2005] introduced the diffusion multipole approximation for thin slabs and multi-layered translucent materials. Although their model more accurately captures the reflectance of translucent layered materials, such as skin, it is difficult to obtain optical parameters for specific materials. The multipole model was later applied to create a spectral shading model specifically for rendering skin [Donner and Jensen 2006], that includes a small set of chromophore parameters to control overall skin appearance. Spatial and subsurface variations other than the differences between homogeneous layers, 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 their method allows interactive rendering of layered translucent materials, it does not improve on the model itself.

More general heterogeneous models that capture the BSSRDF of materials have been developed with the specific goal to obtain model parameters from measurements [Goesele et al. 2004; Tong et al. 2005; Peers et al. 2006]. These models, however, capture only the relationship between incident and exitant light, leaving the underlying physical and physiological structure of the material unknown. For skin, it is particularly valuable to understand the relationship between structure, pigmentation, and reflectance, as they allow 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] present systems 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 individual 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 at a single surface point. Weyrich et al. [2006] measure scattering profiles using an array of optical fibers, where one fiber is lit to confine illumination to a small spot. The other receiving fibers provide a low-resolution estimate of the profile. The use of their device is limited to sufficiently flat skin areas where the sensor can be placed. Tariq et al. [2006] avoid this shortcoming by using projected stripe patterns, shifted across a face, and observe the scattering response of a moving step function of illumination. 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 in contrast to the similarly simplifying assumption of fixing skin translucency σ_{tr}^{-1} [Jensen and Buhler 2002], as ours 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 *in vivo* measurements of skin [Goesele et al. 2004; Tong et al. 2005; Peers et al. 2006; Wang et al. 2007].

Rather than attempting to model reflectance directly, a common approach is also to tabulate a particular skin sample (such as an actor's face) under different lighting conditions [Georghiades 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 exist that derive chromophore concentrations from single pictures. Tsumura et al. [2003] qualitatively estimate melanin and hemoglobin distributions in a face using independent component analysis (ICA). The method is based on a purely absorptive reflectance model that ignores scattering and is used to vary the effect of melanin and hemoglobin in the input image. Cotton et al. [1999] derive melanin distributions from photographs by analyzing their 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.

In practice, analysis methods have driven the development of new reflection models, even as physically-based models help determine the acquisition and analysis to be performed. Continuing in this tradition, we design our measurement system around a physical model, and use our measurements to derive guidelines for obtaining physical model parameters from calibrated RGB images.

3 Background: Light Transport in Layered Translucent Materials

Light transport in layered translucent materials is well-approximated by the diffusion multipole model [Donner and Jensen 2005]. Given absorption and reduced scattering coefficients σ_a and σ_s', and the thicknesses of each of the layers, the multipole model 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 emittance over the surface given an input ray of light.

Given a two-layer translucent material, we denote its forward directed diffuse reflectance and transmittance profiles as R_{12}^+ and T_{12}^+. We define the forward direction as the incident direction of light on the surface of the material. The backward directed profiles are R_{12}^- and T_{12}^-. The total reflectance is the convolution of the incident flux Φ at the surface of the material with its forward diffuse reflectance profile

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

where $r = \sqrt{x^2+y^2}$, and $r'' = \sqrt{(x-x')^2+(y-y')^2}$.

The forward reflectance profile itself is composed of multiple terms that capture the degree of inter-scattering between layers. Thus, R_{12}^+ is the sum of the profile accounting for light that scatters only in the top layer R_1^+, and the profile 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

$$\begin{aligned} R_{12}^+(r) &= R_1^+(r) \\ &\quad + \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} T_1^+(r')R_2^+(r'')T_1^-(r''')dx'dy'dx''dy'' \quad (2) \\ &= R_1^+ + T_1^+ * R_2^+, * T_1^-, \end{aligned}$$

where $r' = \sqrt{x'^2+y'^2}$, and $r''' = \sqrt{(x-x'')^2+(y-y'')^2}$. We denote this convolution between profiles with $*$. 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 requires summing a series of similar convolutions

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

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 Equation 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 [Donner and Jensen 2005].

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

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

convolution by a non-separable profile R_1^+ is efficently computed as a sum of separable convolutions, where each can also be computed hierachically starting with the result of the previous, narrower convolution. The convolution of two sum-of-Gaussian profiles reduces to a polynomial multiplication when the appropriate Gaussian basis is selected [2007] allowing very efficient evaluation of equation 3.

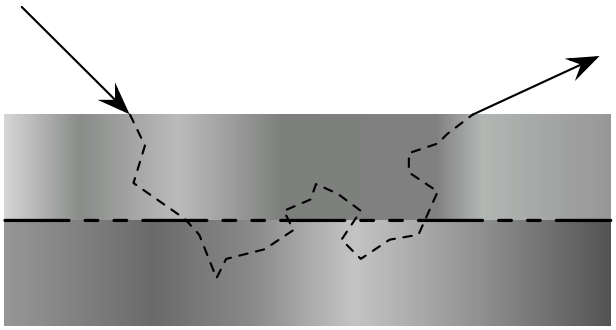

Figure 2: *Heterogeneous light transport in our model. Multiple layers each interact with light in different ways, and their optical properties vary spatially. Light is also absorbed as it passes through an infinitesimally thin absorbing layer between scattering layers.*

Both the Fourier domain and Gaussian polynomial methods extend trivially to more than two layers. This is done by computing the overall profiles of the first two layers, and using those as the responses of a single top layer.

4 A Heterogeneous Reflectance Model

Many real translucent materials are inhomogeneous compositions of elements with different properties, such as leaves, marble, and skin. We account for varying parameters in two ways, first by varying optical properties over the surfaces of layers, and second by introducing thin absorbing layers between scattering layers (see Figure 2).

In general, representing the heterogeneous transport of light would require computing the pair-wise 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, where the residual becomes an albedo texture [Jensen and Buhler 2002; Weyrich et al. 2006]. The dipole, however, is a poor choice for representing materials with significant heterogeneity, as it approximates the transport between points with a single radially-symmetric profile. This is clearly evident across shadow boundaries, as shown in Figure 3.

4.1 Heterogeneity Through Inter-scattering

The multipole model assumes that materials are composed of homogeneous layers, that is, the optical properties of a layer are fixed over its surface. Note that the layer profiles give the homogeneous light transport between points on the surface.

If the properties of the layer vary slowly relative to the mean free path of light in the layer, then we can make the assumption that the local properties of the material are homogeneous. The exitant radiance of the layer is approximately

$$L(\vec{x}_o,\vec{\omega}_o) \approx \int_0^{\infty}\int_{2\pi} f_i(\vec{x}_i,\vec{\omega}_i) R_{\vec{x}_o}(r) L(x_i,\vec{\omega}_i) f_o(\vec{x}_o,\vec{\omega}_o) d\vec{\omega}_i d\vec{x}_i \quad (5)$$

where $r = ||\vec{x}_o - \vec{x}_i||$, and where f_i and f_o indicate the modulation of light at the surface of the layer due to Fresnel effects or a BRDF. Note that we use the profile at the point of exitant radiance $R_{\vec{x}_o}$ to predict the contribution of nearby points. This homogeneous assumption allows the use of simple analytical models such as the dipole and multipole to predict local light transport.

Convolving two homogeneous layers together to find their total transmittance is a radially symmetric process and produces a homogeneous response. With the above formulation of spatially-varying profiles, however, convolution of layer responses depends on the

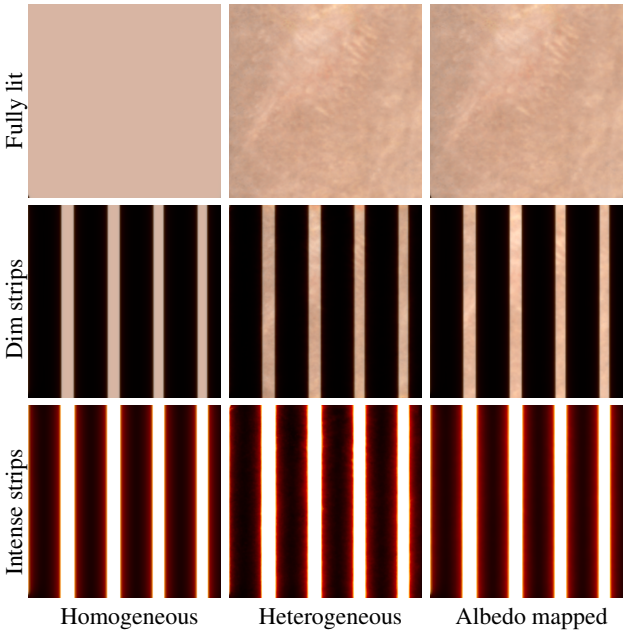

Figure 3: *Renderings of a heterogeneous material (skin) with average homogeneous parameters (left), heterogeneous parameters using our method (middle), and homogeneous parameters modulated by an albedo texture (right). The top row shows fully lit images, while the mideel and bottom row are lit by strips of light of different intensity.). Note that although the middle images seem similar, clear heterogeneities are visible in the brightly lit image, and are not captured by the homogeneous model.*

local position across the interface between the layers. For example, the forward transmittance profile of two heterogeneous layers is

$$T_{12}^{+}(x,y) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} T_{1,x',y'}^{+}(x-x',y-y')T_{2,x',y'}^{+}(x',y')dx'dy'. \quad (6)$$

This is because the incident flux on the lower layer takes the form of a set of points from different profiles, each with different shape. The resulting profile has a complex, sometimes discontinuous structure (see Figure 4). Compared to the radially symmetric convolution of homogeneous layers in Equation 2, the full 2D convolution is considerably more expensive to compute.

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 Equation 3 in 2D. Using the sum-of-weighted-Gaussian representation in Equation **??**, however, keeps evaluation of these convolutions efficient, as we discuss in Section 4.3.

Note that when a layer is highly scattering, light that propagates to its adjoining layers does not spread very far, and conversely. Thus, this approach of modeling layers works best when layers with high frequency changes in properties are highly scattering, while those with lower frequency variation may propagate light more with less error. This is significantly more accurate than using just the dipole model with an albedo texture, as the transport between points on the surface now takes many paths.

Figure 3 demonstrates the differences in scattering between homogeneous and heterogeneous profiles. The first column and second column shows the rendering of a flat patch of skin using average homogeneous and heterogeneous parameters. The third column uses an albedo texture derived from the first two images. In the second row, the patches are lit with strips of light. Though the appearance

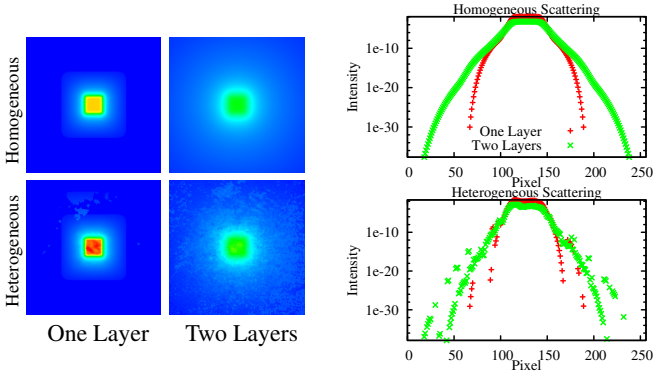

Figure 4: *The false-color images to the left show homogeneous (top) and heterogeneous (bottom) profiles. The left image shows the result of convolving a 16x16 pixel input irradiance with one profile, and the middle image shows the result after convolving with a second profile. The plot to the right shows the center horizontal scanline of each image. Note that there are significant differences, even under the source illumination, between the homogeneous and heterogeneous convolutions.*

seems similar, when lit with brighter strips (bottom row), clear heterogeneities are visible using spatially-varying parameters, while the albedo-mapped image retains its homogeneous characteristics.

Our formulation models the physical characteristics of real materials. In skin, for example, the epidermis is thin and highly scattering, and has the highest frequency variation of the skin layers. The same is true of plant leaves, where the back side of the leaf has a thin, highly scattering layer.

4.2 Inter-scattering Absorption

Real heterogeneous materials have additional types of heterogeneity, such as localized areas of high absorption. We account for this form of absorption by introducing an infinitesimally thin absorbing layer between scattering layers. This absorbing layer modulates light that scatters out of a scattering layer before it enters the next scattering layer. This, however, presents a complication when convolving layer profiles, as the result of each convolution in Equation 3 (except the last of each term) is multiplied by the absorbing layer. The resulting sum of 2D convolutions and products is no longer associative, and Equation **??** becomes

$$\begin{aligned}\Phi * \mathfrak{R}_{12}^{+} = \Phi * \mathfrak{R}_{1}^{+} + \big(\big(\big(\big(\Phi * \mathfrak{T}_{1}^{+}\big)\mathcal{A}\big) * \mathfrak{R}_{2}^{+}\big)\mathcal{A}\big) * \mathfrak{T}_{1}^{-} \\ + \big(\big(\big(\big(\big(\big(\big(\Phi * \mathfrak{T}_{1}^{+}\big)\mathcal{A}\big) * \mathfrak{R}_{2}^{+}\big)\mathcal{A}\big) \\ * \mathfrak{R}_{1}^{-}\big)\mathcal{A}\big) * \mathfrak{R}_{2}^{+}\big)\mathcal{A}\big) * \mathfrak{R}_{1}^{-} + \cdots,\end{aligned} \quad (7)$$

where multiplication by the 2D absorption map $\mathcal{A}$ performs the absorption of the thin layer. As the convolution of a product is not the product of convolutions, previous approaches for efficiently computing Equation 7 cannot be applied; each individual convolution must be performed.

Since the parameter map $\mathcal{P}$ for the absorbing layer controls the absorbance of the layer, the actual absorption as light passes through the layer is

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

4.3 Efficient Implementation

In previous work, since the profiles were homogeneous and radially symmetric, Equations 3 and **??** could be efficiently precomputed and stored as a single radially symmetric profile. Equation 7, however, requires each 2D convolution to be performed in turn. Precomputation of several 2D convolutions over an entire surface would be prohibitively expensive to compute and store.

Instead, we compute profiles on the fly when shading points on the surface. We represent each profile by a set of weighted Gaussians. Convolving the irradiance with a profile amounts to blurring the irradiance by each Gaussian, and accumulating the weighted result. Since we must evaluate a large number of convolutions to evaluate Equation 7, efficiency and performance are of utmost importance. As a Gaussian blur is separable, performing these blurs is a task ideally suited for the parallelism of modern graphics hardware (GPUs) [d'Eon et al. 2007].

To evaluate Equation 7, we represent irradiance incident at a point $\vec{x}_o$ as a 2D texture, where each pixel represents light incident at a point $\vec{x}_i$. The location of $\vec{x}_i$ in the texture depends on the relative position of the two points, and the local UV parametrization of the mesh. Thus, the relative UV coordinates of $\vec{x}_o$ and all its contributing $\vec{x}_i$ determine the world-space size of the texture.

Given a profile represented by the set of k Gaussians $\mathfrak{G} = \sum_{i=1}^{k} G(v_i)$, the convolution of the irradiance Φ with this profile is a weighted sum of blurs of different widths

$$\Phi * \mathfrak{G} = \sum_{i=1}^{k} w_i G(v_i) * \Phi \tag{9}$$

Since a Gaussian blur is separable, however, this formulation leads to an efficient method of convolution as a two-pass process. For each Gaussian in the set, we blur the irradiance texture and accumulate the weighted result.

Convolving two profiles is done in a similar fashion. We perform each convolution in turn, starting with the irradiance and the first forward reflectance profile. The results of the previous convolution are then used as the input to the next. If there is an absorbing layer, then the results of a convolution are multiplied by a discretized absorption texture before becoming the input to the next set of blurs.

Note that although a convolution is a point-wise product in the Fourier domain, the cost of a Fourier transform is $O(n \log n)$, where n is the number of profiles samples. A separable blur, however, is only $O(n)$. Thus actually performing the convolutions brute-force is more efficient.

We constrain the variances of the set of Gaussians to be integral powers of an initial variance. With this formulation, we start with the thinnest Gaussian, and the next wider blur in the set is easily computed from the results of the previous one, as successive blurs sum variances.

Note that the discretization of the irradiance requires a good UV parametrization of the surface being rendered. We use the UV coordinates of irradiance samples to determine the properties of the layer at that point. Ideally, points that are close to each other in world space should be near each other in UV space, and *v*ice versa. This parametrization is also necessary to determine the correct absorption between layers in the local frame.

When computing transmittance profiles, however, the UV parametrization between shading and lit points will 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 far surface, and use its UV coordinates along with those of the irradiance to determine the blur frame. This approach is more general for complex geometry, as opposed to always evaluating a translated reflectance profile, as done in previous work.

This formulation of convolutions followed by multiplications gives the flexibility needed to approximate the spatially varying properties of materials, and account for the transmittance through thin absorbing layers. These features make this model ideal for simulating the reflectance of skin, which we discuss in the next Section.

5 Spectral Shading of Heterogeneous Skin

One of the most striking visual components of skin is its color, which is 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 absorb light. Scattering occurs from small-scale cellular structures, collagen, and also from chromophores, as well as high-frequency changes in index of refraction. This is generally modeled at a high-level as scattering from particles [Jacques 1996]. Both scattering and absorption vary significantly over the visible spectrum, as well as across the surface. This makes skin a good candidate for rendering using the heterogeneous reflectance model described in the previous section.

5.1 Skin Pigmentation

From the standpoint of pigmentation, skin is divided into roughly two layers, the epidermis at the surface, with the dermis below. The human epidermis contains several chromophores, but the most prominent are melanin and carotene. Melanin acts as optical protection from harmful UV radiation, while hemoglobin is the primary carrier of oxygen in blood. Small amounts of carotene [Sayre and Black 1992] are also found in the epidermis, often dependant on diet, but light absorption is dominated by melanin. In the dermis, hemoglobin is the primary chromophore. Figure 7 shows the absorption spectra of these chromophores.

Some pigmentation is focused between the epidermis and dermis. Melanin, for example, is produced by cells at the base of the epidermis, and thus has a significant concentration there [Matts et al. 2007] (see Figure 5). Artificial pigmentation, such as a tattoo, is usually injected into the upper dermis to prevent the pigment from dissipating as epidermal cells regenerate [Bernstein 2006]. These forms of pigmentation can be handled as absorbing layers in our heterogeneous model.

5.2 Spectral Skin Properties

Similarly to Donner and Jensen [2006], we model skin as a two-layer translucent material, and compute the reflectance and transmittance profiles based on the concentrations of their respective chromophores. Note that these profiles are spectral.

The total spectral absorption with the above epidermal chromophores is

$$\begin{aligned}\sigma_{\mathrm{a}}^{\mathrm{epi}}(\lambda) = \; & C_{\mathrm{m}}\left(\beta_{\mathrm{m}}\sigma_{\mathrm{a}}^{em}(\lambda) + (1-\beta_{\mathrm{m}})\sigma_{\mathrm{a}}^{\mathrm{pm}}(\lambda)\right) \\ & C_{\mathrm{he}}\left(\gamma_e\sigma_{\mathrm{a}}^{\mathrm{oxy}}(\lambda) + (1-\gamma_e)\sigma_{\mathrm{a}}^{\mathrm{deoxy}}(\lambda)\right) \\ & + C_b\sigma_{\mathrm{a}}^{\mathrm{bc}}(\lambda) + (1 - C_{\mathrm{m}} - C_{\mathrm{he}} - C_b)\,\sigma_{\mathrm{a}}^{\mathrm{baseline}}\end{aligned} \tag{10}$$

where λ is the wavelength of light in nanometers, C_{m} is the total volume fraction of melanin, and β_{m} blends between the two different types of melanin, eumelanin $\sigma_{\mathrm{a}}^{\mathrm{em}}$ and pheomelanin $\sigma_{\mathrm{a}}^{\mathrm{pm}}$. C_{he} is the total volume fraction of hemoglobin, γ_e controls the relative amounts of absorption from oxygenated $\sigma_{\mathrm{a}}^{\mathrm{oxy}}$ and deoxygenated $\sigma_{\mathrm{a}}^{\mathrm{deoxy}}$ hemoglobin, and $\sigma_{\mathrm{a}}^{\mathrm{baseline}}$ is a baseline absorption of the remaining skin tissue. $\sigma_{\mathrm{a}}^{\mathrm{bc}}$ is the absorption of carotene. We include hemoglobin in the epidermis of our model to simulate the effects of increased blood flow, or erythema, which gives skin a reddish appearance. See [Jacques et al. 2001] for tabulated versions of the absorption spectra of the chromophores.

The total spectral dermis absorption is

$$\begin{aligned}\sigma_{\mathrm{a}}^{\mathrm{derm}}(\lambda) = \; & C_{\mathrm{hd}}\left(\gamma_d\,\sigma_{\mathrm{a}}^{\mathrm{oxy}}(\lambda) + (1-\gamma_d)\,\sigma_{\mathrm{a}}^{\mathrm{deoxy}}(\lambda)\right) \\ & + (1 - C_{\mathrm{hd}})\sigma_{\mathrm{a}}^{\mathrm{baseline}}\end{aligned} \tag{11}$$

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 Equations 10 and 11, these directly determine the optical properties (σ_s' and σ_a) of the skin layers.*

where γ_d controls the oxygenation of the dermal hemoglobin.

We fix the reduced scattering coefficient σ_{s}' of light in skin in our model, describing it by the following power law

$$\sigma_{\mathrm{s}}'(\lambda)^{\mathrm{epi}} = 14.74\lambda^{-0.22} + 2.2\times 10^{11}\times\lambda^{-4}. \qquad (12)$$

where the scattering in the dermis is that of the epidermis reduced by half.

For computing reflectance, we fix the thickness of the epidermis to be 0.25mm, while the dermis is semi-infinite for reflectance. Also, we fix the indices of refraction of the two layers to be 1.4 and 1.38.

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]. Table 1 summarizes the parameters to the skin model.

5.3 Rendering Heterogeneous Skin

To render images of skin using our heterogeneous model, we use 2D parameter maps over a surface to control the values of the model parameters in Table 1, as shown in Figure 1. Using these parameters in Equations 10 and 11 gives the absorption coefficient σ_{a} (recall that σ_{s}' is fixed per wavelength) of the skin layers at each point on their surfaces. We then use these optical properties to construct Gaussian representations of each profile. We convolve these profiles on the fly to compute the final reflectance using Equation 7 as described in Section 4.3.

In addition to melanin in the epidermis, based on the observations of [Matts et al. 2007], we assume that some fraction of melanin is concentrated within an absorbing layer between the epidermis and dermis. When computing Equation 7, we assume that an additional 1.75% absorbance is concentrated at each point in the absorbing layer. We base this amount on our empirical observations.

To render thin materials, such as ears, we determine the actual thickness of the dermis as discussed in Section 7. We then construct a three layer material with a top epidermis, an thin inner dermis, and a bottom epidermis, and compute transmittance profiles using the multipole model of the three layer configuration. Note that this potentially involves multiplying by both the front- and back-side absorbance layers.

6 Reflectance Measurements

Our heterogeneous model has five parameters that control subsurface scattering, only ρ_s is separate. One option for obtaining parameters for the model is 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 intended to define physiological parameter ranges of healthy and pathological tissue. Choosing good values to achieve a certain desired visual look is still a non-trivial task.

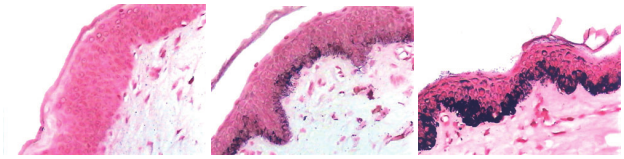

Figure 5: *Cross sections of skin with varying amounts of melanin. The samples have been stained to highlight the melanin distribution. From left to right, the samples represent light skin, moderately pigmented skin, and heavily pigmented skin. Note that melanin is distributed fairly evenly throughout the epidermis, with a concentration at the junction between the epidermis and dermis. Images used with permission from [Matts et al. 2007].*

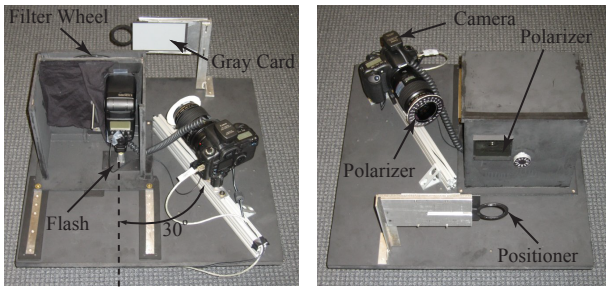

Figure 6: *Our measurement setup.* Left: *Flash inside its case (back and top removed), facing the positioner ring with the gray card. Filter wheel visible at top of case.* Right: *View onto camera and flash polarizers.*

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], partially performing measurements to constrain the respective fit [Weyrich et al. 2006]. As our model comes with an increased number of parameters, these methods do not apply. Instead, we perform our own skin measurements to derive representative parameters for our model. In Section 7.2 we show how these measurements can be used to find reasonable constraints that allow our model to be fit to calibrated RGB data.

6.1 Design

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

With five degrees of freedom at each surface point, however, it is not possible to restrict measurements to three (RGB) channels alone, as has been done by all previous skin studies in graphics. To allow for discrimination of the model constituents, we measure at multiple, narrow frequency bands. To this end, we subsequently place narrow band-pass filters (Chesire Optical, Newport® Optics) in front of the camera flash. At 165 mm (6.5″) distance to the skin sample, the flash's output modulated by the filters' narrow spectrum still provides enough radiant energy for imaging, while the 10 nm-wide 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

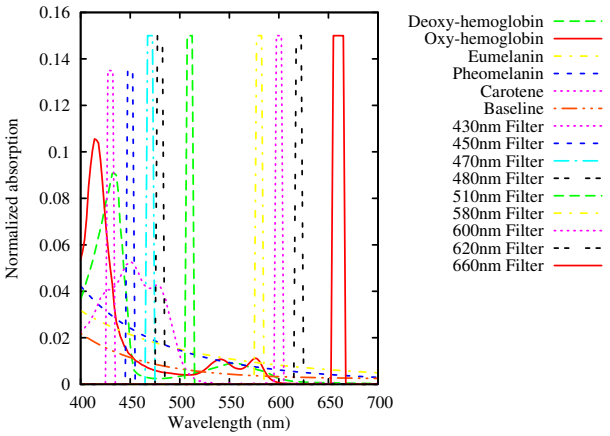

Figure 7: *Filters used for an optimal discrimination of the model constituents overlayed on the absorption of chromophores used in the model.*

that provide 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 filter choice also takes care to avoid the ragged region around 420 nm in the spectrum of the xenon flash.

For maximum irradiance control we structure the incident illumination by affixing black electrical 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 function[1].

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 statically mounted in the setup that is simultaneously lit by the light source and observed by the camera throughout all measurements.

Figure 7 shows the filter spectra used, overlaid with the spectra of the model constituents.

6.2 Measurement Procedure

We position the taped skin sample in front of the measurement apparatus, taking subsequent images, advancing the filter wheel after each shot (F/10, exp 1/100, ISO 200–800, depending on wavelength). The filter wheel also contains an opaque inset used for black-image acquisition.

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 has therefore 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 from F_t by up to ***XX***. 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 closely adhere to skin. The gel remains glossy throughout the measurement procedure and removes the influence of skin's surface BRDF by inducing a transmittance described by a smooth Fresnel term.

Before further processing, we resample the captured images within the cut-out region. By determining the cut-out's location for each image separately, we implicitly compensate for potential subject motion between subsequent 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 [Marschner 1998]. Starting from average skin parameters, a gradient-descent optimization subsequently renders the skin patch and alters the model parameters to minimize the differences between the rendered patch and the acquired image. We use a multi-spectral GPU implementation of our skin model that extends the rendering method of d'Eon et al. [2007] as described in Section 7.

The inverse renderer simulates the incident illumination modulated by the black-tape occluder and computes radiant exitance at each pixel for each measured wavelength. The gradient descent simultaneously optimizes the parameter vector in each pixel, constantly estimating the objective function's gradient from 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 profile's large support leads to a comparatively fast convergence after typically 400 iterations.

6.4 Error Analysis

Great care has been taken to radiometrically calibrate all components. Besides random camera noise, a few systematic error sources due to the polarization potentially introduce a bias. The cross-polarization, intended to eliminate surface reflection, may be imperfect, potentially letting light leak through the crossed polarizers. We found, however, that in our setting even mirror reflection is attenuated down to noise level. The polarization itself adds further effects. 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 less than Spectralon, which is an important observation, as previous work in graphics often relies on cross-polarization to transmit exactly 50% of the diffuse reflectance. On a (gel treated) skin sample that is reflecting away from the camera, to avoid glossy reflection under parallel polarization we observe a polarization-dependent intensity variation of $\pm 19.96\%$ around the average diffuse reflectance. These biases are included in our calibration.

The radiometric accuracy has direct influence on the convergence of the inverse rendering. With the proper calibration, inverse rendering typically converges to a residual of below RMS 4% reflectance within 400 iterations. Further iterations hardly reduce this error, such that the remaining mismatch is due to noise and to potential effects in the tissue that are not simulated by our model.

7 Results

We have implemented the heterogeneous shading model described in the paper in two forms. The first is a GPU-only version that is limited to simple geometry, shadows, and texture resolution. We have also implemented a CPU version in a Monte Carlo ray tracer, but it is

[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; in our data-fitting process, however, we do not directly exploit this fact.

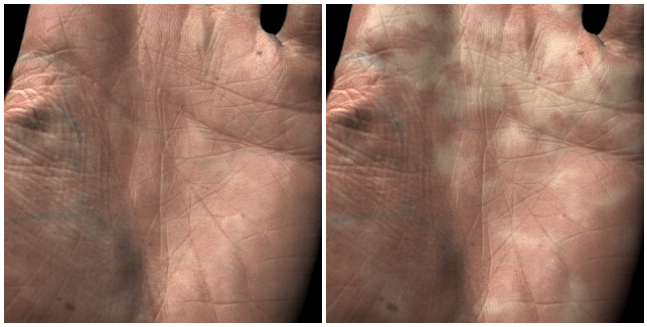

Figure 8: *Changing skin parameters with mechanical deformation. The example shows a reduction of blood flow (hemoglobin concentration) at after clenching and releasing the hand.*

limited by the convolution speed of the GPU. The GPU version calculates the convolution of the shading points under all image pixels in parallel, while the CPU version samples the irradiance in a small area around the shading point.

The renderings in this section were performed on a Quad-Core Intel Xeon 2.0GHz machine with two NVIDIA GeForce 8800GTX and one GeForce 8800GT (G92) graphics cards. The GPU renderer computes one image using a single GPU in about a second, while the CPU version took about 30 minutes per image. We discuss the limitations on efficiency of the CPU method in Section 7.3.

In both implementations, we represent each profile as a set of 18 Gaussians. The smallest variance in the set is 0.0001mm, and the other Gaussians have increasing powers of this value, as discussed in Section 4.3. We base the smallest size on the mean free path length of light in the epidermis, which is about 0.01mm. Each of the Gaussian weights are stored in a lookup table (LUT) indexed by σ_a, as scattering is fixed in our skin model. Each LUT is 18x256 pixels, where each column indexes the weights for a Gaussian of particular variance, and the rows for that Gaussian's value of σ_a. We have found that 256 samples of σ_a spaced logarithmically is sufficient for rendering skin. Note that for all images rendered in this paper, the surface oiliness parameter ρ_r was fixed at 0.35.

We render images using 8 spectral samples, at wavelengths 400nm, 435nm, 470nm, 505nm, 540nm, 575nm, 610nm, and 645nm. We have found this set to give results within 1% pixel error of 151 samples spaced every 2nm. Note that this is a different set of wavelengths than we use for acquisition, as that set is optimized to show the differences between model parameters. The GPU version of the renderer, however, uses the measurement wavelengths when performing inverse rendering. We perform four iterations of inter-scattering convolutions for blue light, and five for green and red light, after analyzing the total spectral error based on the number of inter-scattering iterations performed (e.g. the number of terms of the sum in Equation 7). Figure 9 summarizes the error for skin with minimal absorption, where error is highest. When rendering the surface using the CPU renderer, we use irradiance and convolution texture sizes of 128x128, and warp irradiance samples into it as described in Section 4.3. Higher resolutions gave no appreciable increase in quality.

7.1 Using the Model

Figure 8 shows images of the palm of the left hand before and after undergoing a possible mechanical deformation (e.g. clenching the and 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 parameter maps used to generate the left image are shown in Figure 1. Two

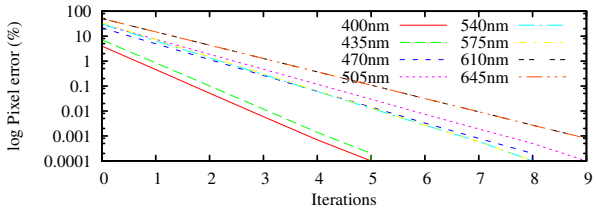

Figure 9: *Spectral error depends on the number of inter-scattering convolutions performed. We have found that using a small number of iterations to compute the inter-scattering of light between layers is sufficient to preserve over 99% of the spectral energy of the profiles.*

Figure 10: *Images of a hand with different pigmentation. Left: Texture synthesis of real captured data simulates freckles over the arm and wrist. The original data was taken from a person with red hair, and the freckles contain high amounts of pheomelanin. Middle: Using absorbing textures to create a tattoo on the back of the hand. Note that the tattoo appears to be beneath the surface (where it is). The colors of the tattoo are faded due to the scattering and absorption of light in the epidermis. This is why tattoos often appear blue. Right: Rendering using a more heavily pigmented dataset.*

maps correspond to Note the simplicity of the maps; the complex interaction of light is captured by the heterogeneous layered model.

Figure 10 demonstrates the use of our reflectance model to simulate different types of pigmentation. The left image uses texture synthesis to generate freckles over the hand and wrist. The original dataset was measured from a person with red hair; the freckles are concentrations of pheomelanin. The middle image demonstrates our model's thin absorbing layers between scattering layers. The pigment of the tattoo is represented using eight maps, one for each spectral sample, to control the amount of absorption between the skin layers. Note that the tattoo appears faded. This is because the light passing through it is attenuated from scattering and absorption in the epidermis. This is also why veins appear blue. The right image (scar?)

Figure 11 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 the illumination normally, but project the shade points to the lit side of the ear to obtain the convolved irradiance as described in Section4.3.

7.2 Skin Measurements

To evaluate our model using actual chromophore concentrations of real skin samples, we measured a number of skin patches across different skin types. A measurement takes roughly 20 second; the limiting factor is the manual advancing of the filter wheel.

The skin patches roughly cover 500x500 pixels in the camera image. After bilinearly resampling them to 512x512, we finally downsample (box-filter) the patch to 256x256 for inverse rendering, which provides a certain degree of noise-reduction; no other filtering is applied.

Figure 12 shows a patch of skin that has been reconstructed with our acquisition procedure. Figure 13 shows the eumelanin, pheomelanin

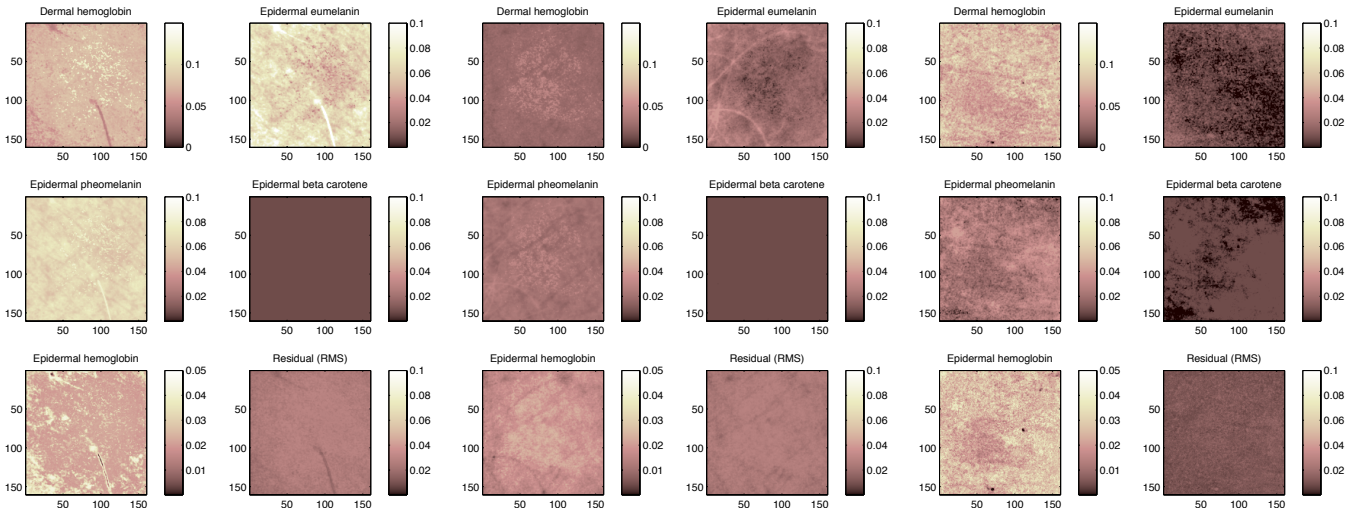

Figure 13: *Chromophore concentrations as derived from our measurements.* (a) *Asian-subcontinental, skin type VI, measured at the outside of the lower arm.* (b) *Asian-subcontinental, skin type V, measured at the outside of the lower arm.* (c) *Caucasian, skin type II, vessel on the back of the arm.*

and hemoglobin distribution in this patch.

comparison of spectral sample images with the spectral photos

7.3 Discussion

Rendering time is currently dominated by performing Gaussian blurs on the GPU for both implementations. Our method is memory-bandwidth-limited due to texture accesses in the Gaussian blur kernels. We have found that multiple GPUs in one machine improves the speed of the convolutions nearly linearly. We do not use SLI hardware, rather, we address each GPU independently in parallel. As there is little communication between the CPU and GPU other than control information (Gaussian width, scaling factors, etc.), our implementation does not saturate the PCI-Express bus even with three graphics cards in one machine. Thus, we expect our approach to scale well as computer with multiple PCI-Express slots are becoming more commonly available.

Our CPU implementation currently samples irradiance and performs convolutions for every shade point rendered. A more optimal solution would be to share convolution results with nearby shade points, or cache them for later use.

We have performed measurements of several patches of skin from different ethic groups, covering a range of

other things to put here? measurements? what we have learned? we've measured things and found them to match what they should be?

8 Discussion and Conclusions

Our model represents a tradeoff between variability and accuracy.

Model can be applied to other materials than skin, fairly general for homogeneous transport in layered materials.

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.

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 of *Computer Graphics Proceedings*, 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*, Computer Graphics Proceedings, 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.

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.

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*, Computer Graphics Proceedings, 164–174.

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

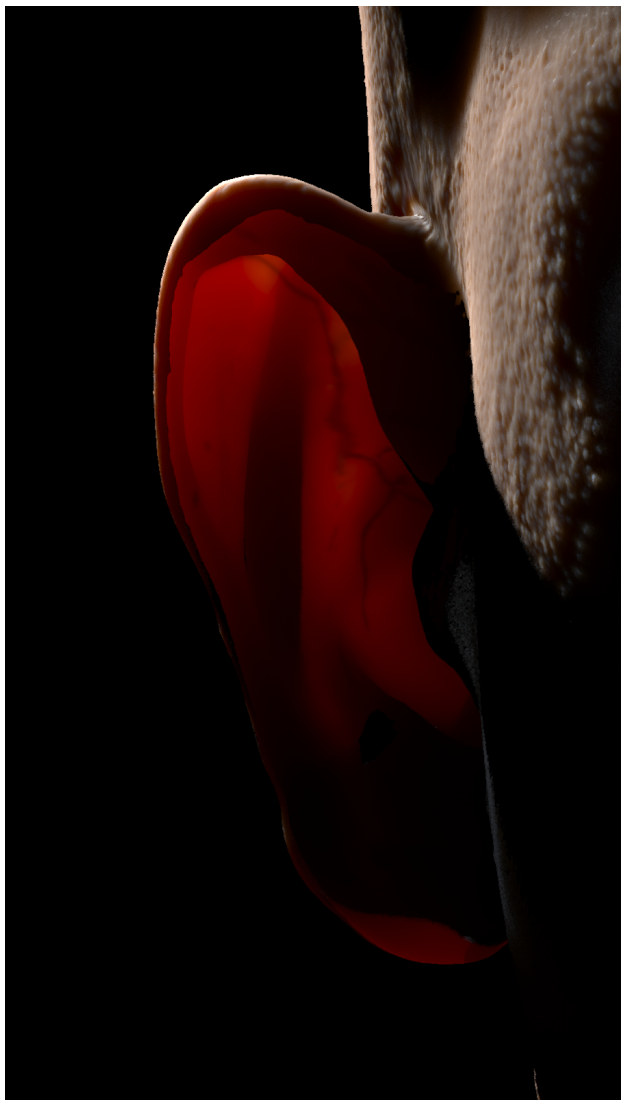

Figure 11: *A back-lit ear rendered with absorbing layers. The transmittance is calculated using the multipole model, not an approximate transmittance using a reflectance profile. The veins and freckles are controlled by parameter maps over the surface of the head and ear.*

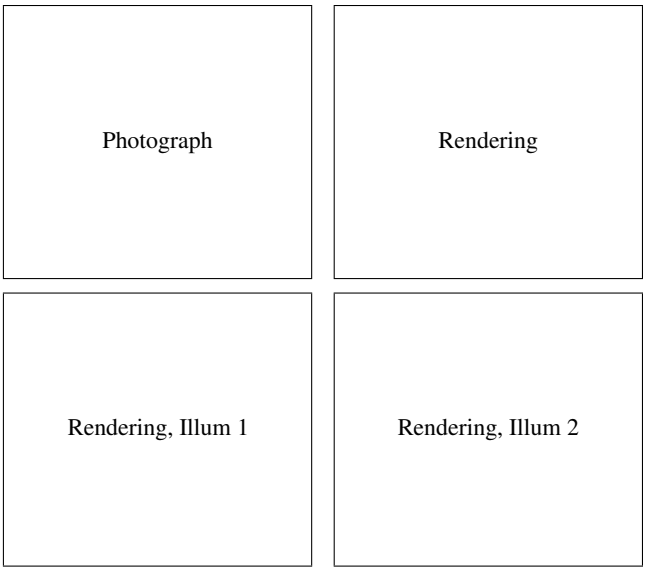

Figure 12: *Measured patch of skin with reconstruction.* Top-Left: *Skin patch as captured in the acquisition device.* Top-Right: *Rendering using the model parameters estimated from this patch.* Bottom: *Renderings under (spatially) varying illumination.*

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

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.

MARSCHNER, S. R. 1998. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University, Ithaca, NY.

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., RAMAMOORTHI, 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 Transactions on Graphics 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.