

A Hyperspectral Space of Skin Tones for Inverse Rendering of Biophysical Skin Properties

Carlos Aliaga¹ Menqi Xia^{1,2} Hao Xie¹ Adrian Jarabo¹ Gustav Braun¹ Christophe Hery¹

¹Meta Reality Labs Research ²EPFL




Figure 1: Our model estimates the biophysical skin properties from a single RGB diffuse reflectance with baked occlusion, allowing de-occlusion, and spectral upsampling in both the visible and near infrared (NIR) spectra. a) Input diffuse reflectance captured in Lightstage. b-d) Reconstruction (absolute error as inset) in the visible (b) and near-infrared range (d), as the product of the reconstructed skin reflectance (c) and recovered occlusion (c inset). These are computed from the estimated biophysical parameters, namely: e) melanin concentration, f) hemoglobin (blood) concentration, g) blood oxygenation, h) eumelanin/pheomelanin ratio, and i) epidermal thickness (in μm).

Abstract

We present a method for estimating the main properties of human skin, leveraging a hyperspectral dataset of skin tones synthetically generated through a biophysical layered skin model and Monte Carlo light transport simulations. Our approach learns the mapping between the skin parameters and diffuse skin reflectance in such space through an encoder-decoder network. We assess the performance of RGB and spectral reflectance up to $1\ \mu\text{m}$, allowing the model to retrieve visible and near-infrared. Instead of restricting the parameters to values in the ranges reported in medical literature, we allow the model to exceed such ranges to gain expressiveness to recover outliers like beard, eyebrows, rushes and other imperfections. The continuity of our albedo space allows to recover smooth textures of skin properties, enabling reflectance manipulations by meaningful edits of the skin properties. The space is robust under different illumination conditions, and presents high spectral similarity with the current largest datasets of spectral measurements of real human skin while expanding its gamut.

CCS Concepts

• *Computing methodologies* → *Reflectance modeling; Reconstruction;*

1. Introduction

The appearance of skin is the result of the interaction of light with the tissues forming the skin. Accurately modeling it requires accounting for its different layers, as well as the absorbers and scatterers inside the tissue. It is the combination of these parameters that gives skin its color.

Given the importance of virtual humans in computer graphics and vision, several models have been proposed for modeling the skin at various levels of realism. The most precise ones have been derived from the biophysical composition of the skin [DJ06; KB04; BK10; CBKM15; IAJG15], mapping between the biophysical parameters describing the composition and structure of the skin and

optical parameters governing light transport inside the skin, which is ultimately responsible of the diffuse reflectance of skin.

Unfortunately, authoring these biophysical models is challenging, and extracting them from non-invasive measurements is of great interest beyond computer graphics for gaining a deeper understanding on the structure and composition of skin samples. This has led to several inverse methods for reconstructing such parameters from photographs, leveraging biophysical constraints. However, these models either require hyperspectral captures [GAS*19; ZDP*19; GGD*20; GGR*21], or impose severe simplifications [AS17].

In this paper, we present a robust method for extracting the biophysical parameters describing the appearance of human skin from RGB images of the diffuse reflectance. Our method builds upon a hyperspectral space of human skin reflectance, obtained from Monte Carlo simulations on a biophysically-based model of human skin. We extensively validate our biophysical space on existing databases, presenting high spectral similarity with current existing databases.

Based on our space of human skin reflectance, we develop an inverse mapping between RGB measurements, the biophysical parameters describing the skin, and a spectrally-upsampled skin albedo. We derive this mapping using a supervised encoder-decoder neural network, where the encoder moves from RGB skin reflectance to biophysical parameters, and the decoder translates the latter to spectral reflectance. Our method is robust, accounts for the overdarkening due to occlusion which previous models infer as higher concentration of melanin, and produces biophysical concentration maps that roughly agree with medical literature.

We demonstrate our approach on a variety of skin tones reconstructing the skin reflectance with minimal error. While our goal is not to obtain in-vivo accurate measurements for medical imaging, we still provide the most complete biophysical inversion model up-to-date, which can be directly used for rendering photorealistic virtual humans (see Figure 11 and Supplementary), and enabling realistic biophysically-inspired editing of human skin such as tanning and blushing. In addition, we are confident that our methodology opens the door for future non-invasive, in-vivo measurements of skin melanin concentration or hemoglobin oxygenation, from low-cost off-the-shelf measurement devices.

In summary, our **contributions** are:

- An expressive space of spectral skin reflectance constrained by biophysical parameters of skin, that agrees with measurements reported in tissue optics and medical research.
- A learned inverse mapping from RGB skin reflectance to biophysical skin properties, that allows recovery of high-resolution spatially-varying maps of skin properties, that allow expressive biophysically-motivated edits on the skin reflectance.
- A learned forward mapping between biophysical skin parameters and hyperspectral skin reflectance in both the visible and the near-infrared (NIR) regimes. Together with the inverse RGB-to-parameters mapping, this allow us to perform spectral upsam-

pling of the reflectance of the skin leveraging the gamut space of human skin.

2. Related Work

Accurately capturing and modeling the appearance of human skin has been a very active area of research for decades. Here we focus on appearance modeling and reconstruction, and refer to a number of extensive surveys for a broader perspective on human modeling [INN07; NMM*19], face appearance capture [KRP*15], or capture of translucent materials in general [FJM*20].

Biophysical skin models Beyond the typical approach in VFX of using a diffuse albedo texture for modeling the skin reflectance, a number of models have been proposed to relate the appearance of human skin with its biophysical structure and composition. [THM99; TOS*03] proposed a simple model accounting for the distribution of melanin and hemoglobin. Later, more advanced models accounting for the multilayer nature of skin, as well as more comprehensive list of chromophores have been proposed [DJ06; CBKM15; IAJG15; KB04; BK10], as well as dynamic models that model the temporal change of the skin coloration from in-vivo measurements of melanin and hemoglobin concentrations [JSB*10]. Our work builds upon these models, and extend them to model a wider range of skins, to create our space for skin reflectance.

Skin capture Since early efforts to acquire the skin BRDF [MWL*99], several approaches have considered skin appearance capture as a spatially-varying diffuse albedo and specular reflectance reconstruction from carefully controlled setups [MHP*07; GHP*08; GFT*11; GTB*13], including dynamic [GRB*18] and single-shot captures [RGB*20]. We leverage the data from these works, and extract biophysical parameters from the diffuse reflectance. With the same goal as our work, several authors have focused on retrieving the biophysical structure and components of the skin from measurements of the diffuse reflectance of skin, either based on RGB [DWd*08] or hyperspectral images [GAS*19; GGR*21; ZDP*19]. Closest to our solution, [GGD*20] reconstructs spectral skin properties based on a simplified skin model [JSB*10; DJ06] through a novel measurement system, benefiting from narrow-band LEDs. Our work reconstructs the skin biophysical parameters from single-image RGB measurements of the diffuse reflectance of the skin, based on a carefully validated space generated from Monte Carlo simulations. In addition, thanks to our encoder-decoder reconstruction architecture, we are able to recover smooth spatially-varying maps of the skin parameters, that qualitatively match the observations from medical literature.

Spectral upsampling Since the development of the *gamut of valid reflectance* [Mac35b], there has been a number of works attempting to solve the ill-posed inverse problem of obtaining spectral reflectance data from RGB measurements, using pre-tabulated box-based values for obtaining a discrete approximation of the reflectance spectrum [Mac35a; Smi99], optimizing continuous parametric functions [MSHD15; JH19], using Fourier mo-

ments [TWF21], or via clustered PCA [OYH18]. Our model up-samples a continuous spectral reflectance function of the skin in both the visible and near-infrared spectrum, by leveraging the reduced gamut of skin reflectances.

3. A Biophysically-Based Space for Skin Reflectance

In order to develop our inverse and forward mappings from skin reflectance to biophysical parameters, we first build a space for such skin reflectance that is comprehensive enough to cover most skin conditions. Unfortunately, capturing the wide variety of skin tones is challenging given the large number of needed subjects, and, even with a complete space, we would still lack the biophysical structural parameters of each skin type.

Instead, we build our space from synthetic data, by using Monte Carlo simulations on a virtual replica of human skin, under a wide variety of relevant biophysical parameters. In the following, we describe the skin model used for building our space, and detail our simulations. Later in Section 4, we demonstrate that our space correlates with existing measured spaces, and analyze its properties.

3.1. Human skin model

Human skin is a multilayered tissue of roughly parallel layers, each of them with different concentrations of chromophores (absorbers that give the skin its color) as well as slightly different indices or refraction. Unfortunately, explicitly modeling such complexity, even stochastically using radiative transfer theory, is not only too expensive in terms of computation, but also results into an unmanageable parameter space.

We decided to restrict our model to two layers; epidermis and dermis, since a similar assumption has proven to be adequate in the past [MM02] and since it fits our purpose. The epidermis is composed of two parts: the *stratum corneum* and the living epidermis. The former is the outermost layer, and has minimal effect on the skin diffuse reflectance given its low absorption and thickness (5 and 20 μm depending on the body location [CSLM19]). Thus, we focus only on the later layer, with parametrizable thickness t .

The dermis is composed by two sub-layers: the reticular dermis and the papillary dermis. Previous works [DWD*08] shown that these can be safely simplified as a single layer, with average scattering and absorption parameters. We model it as a semi-infinite medium, omitting sub-dermal tissues (fat, cartilage, muscles), which are heavily dependent on the anatomical location of the subject. Also, we empirically found that including the dermal thickness has minimal effect on the diffuse reflectance.

We assume that both layers have an index of refraction (IOR) of 1.4, resulting from the weighted sum of the corresponding sub-layers. While this is not completely accurate, it allows to remove the interface between the two layers, and the effect of accurately considering the exact IOR is known to be marginal [LWC12]. Following previous well-established models from tissue optics [Jac13], we use the classic radiative transfer

equation, and treat each layer as an homogeneous medium defined by its spectral absorption (μ_a) and scattering (μ_s) coefficients. We do not account for spatial [VWH17; JAG18] or angular correlation [JAM*10] between scatterers, since there is no measured data of non-stationary particles distributions in human tissue.

Absorption The absorption of each layer μ_{a_i} is the result of the additive contribution of each chromophore absorption μ_{a_c} present in each layer:

$$\mu_{a_i} = \sum_{c \in C_i} \mu_{a_c} = \sum_{c \in C_i} \frac{V_c p_c \epsilon_c}{w_c} \quad (1)$$

where C_i is the set of chromophores contained in the layer i ; V_c is the volume fraction of the substance containing the chromophore c ; p_c is the concentration in g/L; ϵ_c is the molar extinction of the chromophore; and w_c is its molar weight (Table 3 in Appendix A).

Similar to previous approaches [KB04; IAJG15; CBKM15; DJ06], we include the effect of melanin, determined by the melanosomes volume fraction V_m , containing the two types of melanin (eumelanin $\mu_{a_{eu}}$ and pheomelanin $\mu_{a_{ph}}$) in the epidermis, governed by the melanin type ratio ϕ_m , which greatly varies through skin type. We also incorporate the hemoglobin present in blood V_b in the dermis: oxygenated hemoglobin $\mu_{a_{hbO_2}}$, responsible for the saturated reddish tint, and deoxygenated hemoglobin $\mu_{a_{hb}}$, responsible for a desaturated purple color; ϕ_h being the hemoglobin type ratio. Other chromophores included by the model are the beta-carotene $\mu_{a_{\beta-c}}$ in both epidermis and dermis, and bilirubin $\mu_{a_{bil}}$ in the dermis, contained in blood. Additionally, we also include a baseline of skin absorption $\mu_{a_{base}} = 7.84 \times 10^8 \lambda^{-3.255}$ [Sai92] already employed in previous models.

As a result, the total spectral absorption for epidermis and dermis are defined respectively as

$$\begin{aligned} \mu_{a_e} &= V_m (\phi_m \mu_{a_{eu}} + (1 - \phi_m) \mu_{a_{ph}}) + (1 - V_m) (\mu_{a_{\beta-c}} + \mu_{a_{base}}), \\ \mu_{a_d} &= V_b (\phi_h \mu_{a_{hb}} + (1 - \phi_h) \mu_{a_{hbO_2}} + \mu_{a_{bil}} + \mu_{a_{\beta-c}}) + (1 - V_b) \mu_{a_{base}}. \end{aligned}$$

Scattering For modeling the scattering in both the epidermis and dermis, we use the reduced wavelength-dependent scattering coefficient μ'_s computed according to the fit introduced by Jacques [Jac13], which is generic for a wide range of human tissues:

$$\mu'_s(\lambda) = a \left(f_R \left(\frac{\lambda}{\lambda_r} \right)^{-4} + (1 - f_R) \left(\frac{\lambda}{\lambda_r} \right)^{-b_M} \right), \quad (2)$$

where the wavelength λ is normalized by a reference wavelength $\lambda_r = 500$ nm, the scaling factor $a = \mu'_s(\lambda_r)$, f_R is the relative contribution of Rayleigh scattering, and b_M characterizes the wavelength dependence of Mie scattering. We set $a = 36.4$, $f_R = 0.48$, and $b_M = 0.22$ following the coefficients reported in the optics literature for human skin [BGT11]. We model the scattering phase function using the standard Henyey-Greenstein phase function, parameterized by a spectrally resolved anisotropy factor $g(\lambda)$ as [VJSS89]

$$g(\lambda) = 0.62 + \lambda \cdot 0.29 \cdot 10^{-3}. \quad (3)$$

Both epidermis and dermis share the same anisotropy.

Parameter	Description	Epidermis	Dermis
V_m	Melanin Volume Fraction	[0.001, 1]	-
V_b	Blood Volume Fraction	-	[0.001, 1]
t	Thickness [μm]	[10, 350]	2100
ϕ_m	Ratio of melanin types	[0.001, 1]	-
ϕ_h	Ratio of hemoglobin types	-	[0.001, 1]

Table 1: Parameters of our five-dimensional space for skin reflectance, including their ranges. For the full list of biophysical parameters and constants used in our model, we refer to Table 3 in the Appendix.

3.2. Parametrizing the space

We create our reflectance space by varying the skin properties in the ranges listed in Table 1, leading to a five-dimensional space defined by the four main chromophores in epidermis and dermis (eu- and pheomelanin in the epidermis, and oxy- and deoxygenated hemoglobin in the dermis), and the epidermal thickness, which proved to be critical for parameter estimation [ZDP*19]. Other parameters like bilirubin and β -carotene concentrations remain fixed to common values of human skin measurements found in literature (see Table 3 in Appendix A).

Accounting for the epidermal thickness over the face, along with varying values in melanin concentrations, help us achieve local dark zones, such as moles, and generalize over any skin type. We also allow melanin and hemoglobin to go beyond the usual values for human adults measured in the literature [MM02], in order to automatically handle outliers found in the face, such as the lips, which exhibit very thin epidermis and higher blood concentration, or other cases like underlying veins and capillary veins, or areas with abnormal melanin concentration like freckles or spots. This range expansion is also reasonable for the oxygenation level, since it can vary a lot depending on the physical state of the person, and for the melanin type ratio, where there is little agreement in the available measured data.

The final 5D parameter space is sampled as follows: melanin and hemoglobin are selected cubically ($\sqrt[3]{V_m}$) and quartically ($\sqrt[4]{V_b}$) respectively, to better adjust to their non-linear effect on reflectance, while epidermal thickness, melanin type ratio and hemoglobin type ratio all are uniformly sampled. Example slices of the resulting space are shown in Figure 2.

3.3. Computing the diffuse reflectance

We compute the diffuse reflectance of a skin patch by using brute-force Monte Carlo random walks inside the skin. We also considered using the Kubelka-Munk model [AS17], but found it inaccurate (see Section 3 of the Supplemental). For each skin type, we run a monochromatic simulations for wavelengths between 380 and 1000 nm, at steps of 2 nm, an increment we found sufficient for avoiding spectral aliasing. We sample one million random walks per wavelength and skin tone, with initial directions randomly chosen from a cosine-weighted distribution centered at the inverse normal at the boundary of the skin, in agreement with what we later use at rendering time (Section 6).




Figure 2: Matrices of resulting skin reflectance as a function of melanin and blood volume fraction (x and y axis), with $V_m \in [0.1, 43]\%$ and $V_b \in [0.1, 30]\%$, for different epidermal thickness, melanin type ratio and blood oxygenation; note that these parameters significantly affect the appearance of the skin.

4. Analysis of our Reflectance Space

To evaluate the spectral accuracy and colorimetric coverage of our hyperspectral representations, we compare our reconstructed skin reflectances against the Leeds skin spectral dataset [XYZ*17]. The Leeds dataset is the most comprehensive skin data, which was collected from over 1000 subjects across four different ethnicities (Caucasian, Chinese, Kurdish, and Thai), and spectrophotometric measurements were done on different body locations not including faces (forehead, cheek, inner arm, and back of hand). Note that in the Leeds dataset, the measurement device, spectrophotometer, was used with the specular included mode, which may have small inconsistencies with our diffuse reflectances. Also ideally, we would like to compare our reconstructed skin textures against the hyperspectral captures of the same subjects using image color/spectral difference [LU14], which we consider as future work.

4.1. Skin diffuse reflectance gamut

In terms of skin color gamut, Figure 3 shows the CIELAB (under CIE D65 and 1931 2-degree standard observer) distributions of our skin spectra versus the Leeds skin spectra. The bottom row in Figure 3 shows the projection plane of $a^* - b^*$ representing the chromaticity in CIELAB space, and the top row shows the plane of $L^* - b^*$ including the lightness dimension. From the overlapped contour comparisons in the right column, our spectra clearly have a more comprehensive coverage in those dimensions, meaning our hyperspectral manifold has a better diversity synthesizing more skin conditions, which may be limited by the subject population and the point-based measurements in the Leeds dataset.

4.2. Spectral similarity validation

Colorimetric comparisons ensure that our spectral manifold is representative and diverse under the given viewing condition (D65 & 2-deg observer). Here we directly compare the spectral similarity and the color mismatch when the spectra are viewed across different lighting conditions, i.e., illumination metamerism [Ber19]. Illumination metamerism refers to the phenomenon where one pair of spectra match in color under a given illumination (e.g., D65) but have color mismatches under another illumination (e.g., incandescent light). In Figure 4, the black lines in each subfigure correspond to 16 representative skin spectra from the Leeds dataset, from dark to light skin (left to right, up to down). For each Leeds skin




Figure 3: CIELAB distributions of our skin spectra (middle) versus the Leeds dataset [XYZ*17] spectra (left). The CIE D65 and 1931 2-degree standard observer were used for calculations. Bottom row shows the projection plane of $a^* - b^*$ representing the chromaticity, and top row shows the plane of $L^* - b^*$ including the lightness dimension. Our spectra clearly have a more comprehensive coverage in those colorimetric dimensions from the contour comparison in the right column.




Figure 4: Metamer spectra under CIE D65 and 2-degree standard observer from ours (colored dash lines) versus the Leeds dataset (black solid lines). The average RMSE over all the metamers is included in each panel, which generally is less than 0.05.

spectrum, our spectral manifold is screened to find the metamers ($< 1\Delta E_{ab}^*$ in CIELAB) as matched color under D65. It can be observed that those metameric spectra have reasonably high similarity compared with the Leeds spectra, with an average root mean square error less than 0.05 [IRB02]. In a few cases, there are more differences in the longer wavelength region, which is probably caused by the specular component included in the reflectance measures in the Leeds dataset.

4.3. Illumination metamerism evaluation

The spectral difference between the Leeds data versus their metamers, without considering the weights of illumination spectra and human observers, may not necessarily translate to visible differences under different illuminations. To simulate the potential color mismatches caused by the change of illumination, the color differences between the 16 Leeds spectra and their metamers from our spectral manifold were calculated under other illuminations in addition to D65. Eight representative illuminations com-

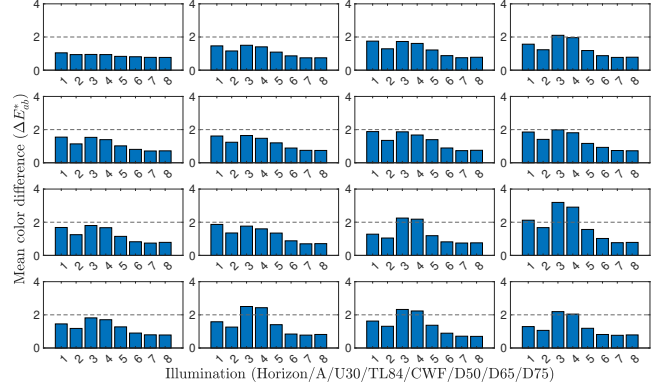


Figure 5: Illumination metamerism between 40 Leeds spectra in Figure 4 versus their corresponding metamers from our spectral manifold. The y-axis is the average color mismatch over all the metamers for the eight representative illuminations, indexed on the x-axis, color temperature from low (Horizon) to high (D75).

monly used in the color-critical applications are considered, including different color temperatures and different spectral characteristics, such as smooth incandescent light and CIE daylight, as well as spiky fluorescent lights. Figure 5 shows the average color difference across all metamers, and the maximum color difference across all the illuminations can be used as a common metric for quantifying illumination metamerism [Ber19] (similarly for observer/device metamerism [XFM20]). Note that under D65, the maximum color difference came from the metamer selection criterion, i.e., $< 1\Delta E_{ab}^*$ in CIELAB. Compared with this baseline, the average color differences are mostly less than $2\Delta E_{ab}^*$ in CIELAB. There are higher color differences in low color temperatures, as well as fluorescent lights, which differ most to the D65 used to create the metamers. The maximum color difference we used outlines the potential worst mismatches only under very different illuminations. Usually $2 \sim 3\Delta E_{ab}^*$ is acceptable for most colorimetric reproduction applications [Ber19], for which our metamer spectra have very promising accuracy, therefore are robust to illumination changes.

From both colorimetric and spectral comparisons above, our hyperspectral manifold presents a more comprehensive and diverse colorimetric coverage than the spectrophotometric measurements in the Leeds dataset. The realistic spectral characteristics in ours provide not only colorimetric accuracy, but also spectral accuracy that leads to less color mismatches across different illumination conditions, as well as potential across different observers/cameras. More analysis can be found in Supplementary Section 2.

5. Mapping Reflectance to Skin Properties and Back

Our biophysical model defines the forward mapping from the skin parameters to the skin reflectance, by means of a Monte Carlo simulation or alternative, a look-up table (LUT). For the inverse process, we need to characterize the mapping from the skin reflectance back to biophysical skin parameters. The non-bijective nature of such mapping, where many combinations of skin properties can lead to the same reflectance, makes this task challenging. A simple

approach like a LUT would lead to quantized results and eventually non-smooth recovered parameters (see Supplemental Material, Section 4). Instead, we opted to learning both the inverse and forward mappings, moving from RGB skin reflectance ρ_{RGB} to our 5D biophysical parameter space γ , and from γ to a hyperspectral skin reflectance ρ .

Overview To define both the inverse and forward mappings, we train an encoder-decoder network in a fully-supervised fashion. This network takes a single RGB value of human skin reflectance ρ_{RGB} , and returns both the biophysical parameters and the spectrally-upsampled skin reflectance. In addition, it returns an estimated occlusion value ϵ , which helps obtain more uniform distributions of melanin volume fraction, melanin ratio and epidermis thickness.

In particular, our encoder f_E maps from $\rho_{\text{RGB}} \in \mathbb{R}^3$ to both the parameter set $\gamma \in \mathbb{R}^5$ and the occlusion value $\epsilon \in \mathbb{R}$, while the decoder f_D takes γ and obtains the hyper-spectral reflectance (310 bands) $\rho \in \mathbb{R}^{310}$, as

$$f_E: \mathbb{R}^3 \rightarrow \mathbb{R}^5 \times \mathbb{R} \quad \text{and} \quad f_D: \mathbb{R}^5 \rightarrow \mathbb{R}^{310} \quad (4)$$

$$\rho_{\text{RGB}} \mapsto (\gamma, \epsilon) \quad \gamma \mapsto \rho$$

Note that we can therefore obtain the initial RGB value ρ_{RGB} as the product

$$\rho_{\text{RGB}} = \text{toRGB}(\epsilon f_D(\gamma)). \quad (5)$$

Network architecture Both the encoder and the decoder are multilayer perceptrons (MLPs) of four fully-connected layers, with two hidden layers of 70 neurons for the encoder, and two hidden layers of 512 neurons for the decoder. We use the tanh activation function for all neurons in the hidden layers, and a sigmoid for the output of both MLPs.

Dataset We generate a dataset of 600k pairs of 5D skin parameter vector γ and the corresponding normalized spectral reflectances ρ . Then, we augment the dataset with 14 scaling intensities ϵ per parameter tuple, uniformly sampled from 0.01 to 1, leading to a dataset of 8.4M RGB values computed using Equation (5), with the classic RGB conversion using D65 as the illuminant CIE 1931; we use ρ_{RGB} as input for the network. We split the dataset in 80% and 20% for training and validation respectively. For validation, we sample the skin parameters according to the uniform distribution. For training, we rely on Quasi-Monte Carlo for a better coverage of the skin parameters, using a low-discrepancy sequence (Halton [Hal64]). Both are subsequently non-linearly remapped following the warping procedures described in Section 3.2.

Training loss We train the encoder-decoder simultaneously in a fully-supervised fashion, by minimizing the error of the biophysical parameters, occlusion value and final reflectance, estimated by the network ($\hat{\gamma}$, $\hat{\epsilon}$ and $\hat{\rho}$, respectively). This error is defined using the following loss function:

$$\mathcal{L} = \mathcal{L}_E + \mathcal{L}_D + \mathcal{L}_{Full}. \quad (6)$$

Here \mathcal{L}_E explicitly minimizes the error of the encoder function, following

$$\mathcal{L}_E = \text{MSE}(\hat{\gamma}, \gamma) + \text{MSE}(\hat{\epsilon}, \epsilon). \quad (7)$$

The second summand \mathcal{L}_D attempts to minimize the error of the decoder as

$$\mathcal{L}_D = \text{SAM}(\hat{\rho}, \rho), \quad (8)$$

where SAM models the spectral angle mapper function [KLB*93], which measures the difference between two spectra by treating them as n-dimensional normalized vectors and computing their angle, following

$$\text{SAM}(\rho_1, \rho_2) = \cos^{-1} \left(\frac{\rho_1}{|\rho_1|} \cdot \frac{\rho_2}{|\rho_2|} \right), \quad (9)$$

with $|\cdot|$ the vector norm, and \cdot the dot product operator.

Finally, \mathcal{L}_{Full} measures the error in the reconstruction as the distance in value and chromaticity, following

$$\mathcal{L}_{Full} = \mathcal{L}_1(\text{Lum}(\text{toRGB}(\hat{\rho}')), \text{Lum}(\rho_{\text{RGB}})) + \text{SAM}(\hat{\rho}', \rho) \quad (10)$$

with $\hat{\rho}' = \hat{\epsilon} \hat{\rho}$. The first term measures the distance in luminance, using the standard luminance in XYZ from D65 RGB i.e., $\text{Lum}(R, G, B) = 0.2126R + 0.7152G + 0.0722B$. The distance in chromaticity for the estimated scaled spectral reflectance $\hat{\rho}'$ and the ground truth spectral reflectance ρ is measured using SAM again. Note that SAM is agnostic to the amplitude of the signal, so it is suitable for comparing the ground truth, normalized spectral reflectances, against the estimated ones (scaled by the occlusion).

Implementation details We implemented our model in Pytorch, using the Adam optimizer [KB14], with a learning rate of 10^{-4} and a batch size of 256. A grid search was performed to find the optimal hyper parameters listed, including the number of hidden neurons that we set to 70 and 512 for encoder and decoder, respectively.

6. Results

We demonstrate our method on five subjects of different ethnicities and genders captured using a Lightstage [DHT*00]. The diffuse RGB reflectance is extracted using the standard Lightstage pipeline, including high-quality diffuse-specular separation, as a 4K HDR texture. Our network takes less than 2 seconds to proceed each 2kx2k map on a Nvidia RTX A6000.

For rendering, we implement a custom skin model in Blender Cycles, using a double GGX specular lobe and a diffuse component computed using Monte Carlo random walks on a homogeneous medium. The medium parameters are obtained using numerical albedo inversion [WVH17] from the reconstructed RGB diffuse reflectance of the skin, as well as its mean free path and its anisotropy parameter. Note that using the spectral reflectance in a renderer with spectral support should be possible, by using our decoder in render time, but it is left as future work. Details and more rendering results can be found in Supplemental Sections 5 and 6.

Figure 6 shows the reconstruction results in RGB for subjects A-

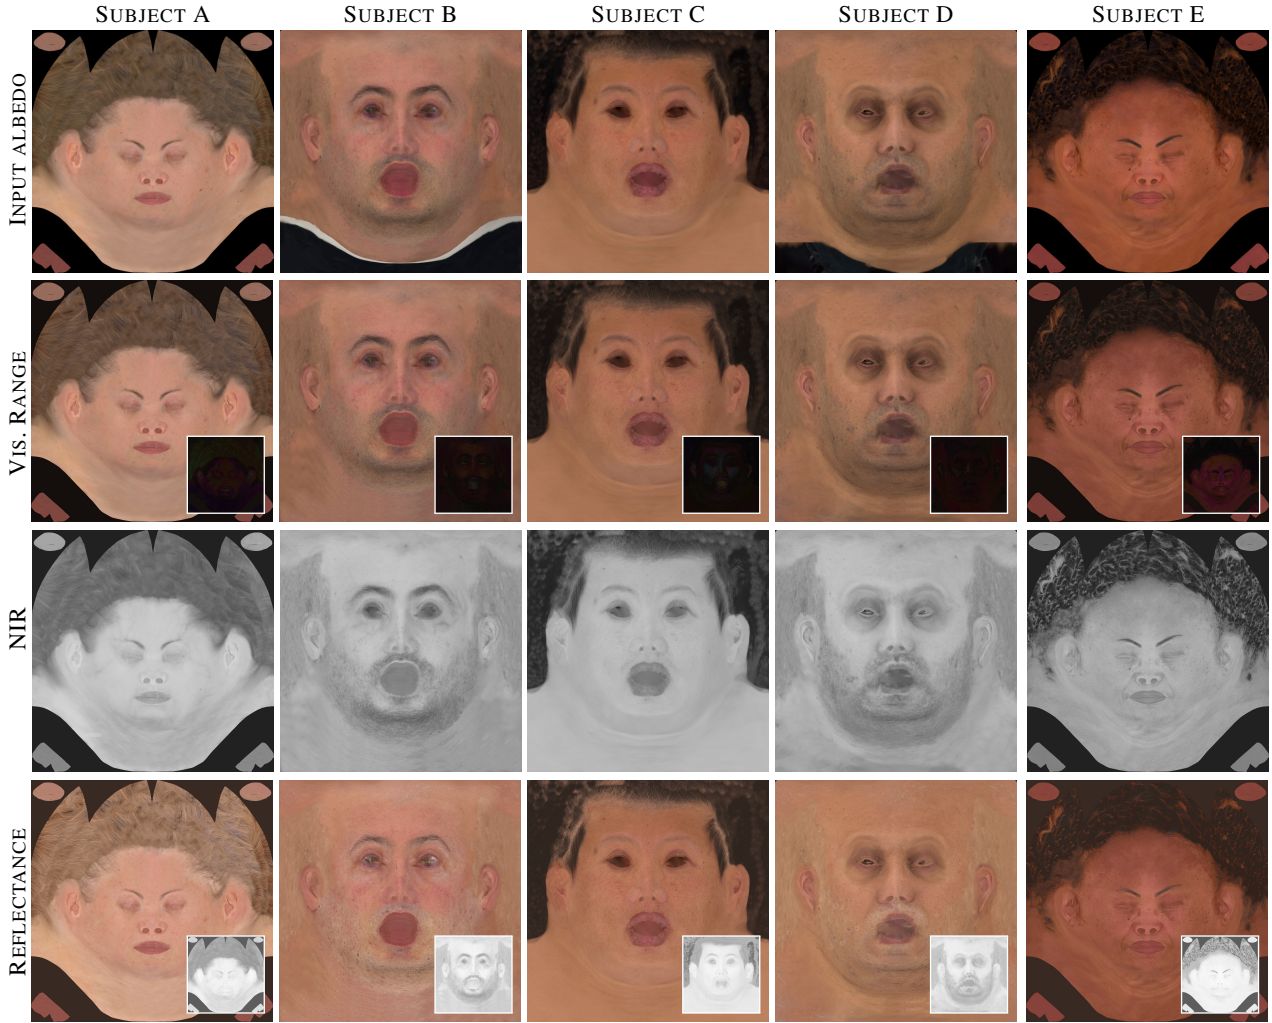


Figure 6: Estimated maps for five subjects of different skin types. Reconstruction mean squared error for the visible range (second row): $7.37 \cdot 10^{-5}$, $5.39 \cdot 10^{-5}$, $3.057 \cdot 10^{-5}$, $4.5 \cdot 10^{-5}$ and $7.61 \cdot 10^{-4}$. The recovered reflectance maps in near infrared (NIR) are computed as the integral of the spectral reflectance from 780 nm up to 1 micron. They look similar for different skin types, since light travels deeper in that region of the spectra where the absorption by the chromophores diminishes drastically. The last row shows the estimated normalized reflectance maps, with the predicted occlusion maps as insets: The model removes baked shadows and ambient occlusion, preventing the network to interpret darker areas by means of higher concentration of chromophores.

E, including the scaled reflectance (reflectance times the occlusion ϵ) in both the visible (toRGB($\rho\epsilon$)) and near infrared (NIR) ranges, as well as the reflectance toRGB(ρ) (the occlusion is shown in the insets). Our results show minimal error in the reconstruction.

Ablation study We evaluate the importance of our design choices by means of an ablation study, where we compare the performance of an RGB-to-RGB network (RGB), an RGB-to-Spectral network (SPECTRAL), and an RGB-to-Spectral network accounting for the de-occlusion parameter when estimating the reflectance (SPECTRAL+). Figure 7 shows that spectral upsampling produces less error in the reconstruction, while allowing to obtain NIR spectral values or using our model for spectral rendering. The addition of the occlusion in our network (SPECTRAL+) does not produce significant changes in the reconstruction, but is important when

generating the biophysical parameter maps, as shown in Figure 12: there, we can see more uniform distributions of the melanin ratio and thickness, with variations only around the eyes and lips, as described in the medical literature. Thus, while the reconstructed spectral diffuse albedo is not significantly affected, the biophysical parameters obtained with our final network are better aligned with real parameters, which might be relevant for e.g. editing tasks.

Robustness under different illuminants We evaluate our model trained for D65 illuminant over a set of skin tones lit by illuminants ranging from very warm to very cool color temperatures: 2500 to 10000 K, applying the von Kries chromatic adaptation model [vKri05] and the CAT02 transform matrix [Fai13]. A subset of 484 skin tones was created by following the sampling scheme in Section 3.2 and limiting the ranges to plausible values for melanin




Figure 7: Ablation study: reconstruction error. While the three versions of the model recover the original reflectance maps with minimal error, enhancing the decoder with dense spectral reflectance data (SPECTRAL) improves overall the RGB approach. The SPECTRAL model even outperforms the final version SPECTRAL+, at the cost of much higher sensitivity to occlusions, which translates into lower quality in the reconstructed maps, as can be seen in Figures 12 for subjects A and C. Such effect is found more pronounced in the RGB approach.

	2500 K	4500 K	8500 K	1000 K
Rec	$2.57 \cdot 10^{-2}$	$3.78 \cdot 10^{-2}$	$5.81 \cdot 10^{-2}$	$6.21 \cdot 10^{-2}$
V_m	$8.5 \cdot 10^{-3}$	$3.06 \cdot 10^{-2}$	$1.54 \cdot 10^{-3}$	$2.2 \cdot 10^{-3}$
V_b	$4.7 \cdot 10^{-2}$	$1.64 \cdot 10^{-2}$	$7.7 \cdot 10^{-3}$	$1.07 \cdot 10^{-2}$
t	$3.6 \cdot 10^{-3}$	$1.2 \cdot 10^{-3}$	$5.4 \cdot 10^{-4}$	$7.7 \cdot 10^{-4}$
ϕ_m	$4.8 \cdot 10^{-2}$	$2.01 \cdot 10^{-2}$	$1.16 \cdot 10^{-2}$	$1.16 \cdot 10^{-2}$
ϕ_h	$1.57 \cdot 10^{-1}$	$4.5 \cdot 10^{-2}$	$1.8 \cdot 10^{-2}$	$2.4 \cdot 10^{-2}$

Table 2: For the same subset of skin tones lit by 2500 K to 10000 K illuminants: RGB reconstruction error (MAE) in the visible spectrum of our model trained for D65 illuminant, and mean absolute differences between each of the corresponding estimated biophysical parameters and their D65 counterparts.

$V_m \in [0.001, 0.43]$ and hemoglobin $V_b \in [0.001, 0.2]$. Results are shown in Figure 8. The reconstruction error remains low for all the illuminants including the coolest and warmest color temperatures, which significantly differ from D65. The differences in the estimated properties between D65 and the rest of the illuminants in MAE are shown in Table 2. In terms of parameters, the distributions of the biophysical properties exhibit a high degree of similarity, with only subtle variations at the most extreme color temperatures: mainly a decrease in thickness and an increase in oxygenation level to accommodate warmer illuminants, or the opposite for cooler color temperatures. To show the effect over faces, we conducted a similar assessment in Figure 9, using the captures from

Figure 1. The reconstruction error remains minimal. However, it is worth noting that the model demonstrates greater accuracy in reproducing warmer tones by leveraging certain biophysical skin properties (e.g. hemoglobin and blood oxygenation), whereas reproducing extreme cool color temperatures proves to be a bit more challenging for the model. These results show that our model, even when trained for a particular illuminant, is still robust and consistent in other illumination conditions. Nevertheless, the network should be retrained for illuminants with emission spectra significantly different to the one used for training.

Comparisons to previous work We conduct a series of comparisons with the related and recent work [GGD*20]. Not having access to the training data, we decided to reproduce the authors' results from the examples in the paper with our approach. Figure 10 is a comparison on more true albedos from [GGD*20], which were obtained via the Antera device (under D65 illuminant). The general observation is that our reconstruction error is much lower, and we tested the model across more skin types. See the Supplemental Section 1 for additional tests.

Editing We show how we can manipulate directly in this space of inferred skin properties, scaling some of them up or down, in an intuitive and predictable manner. We run the neural decoder on these modified quantities to reconstruct biophysical albedos, and finally render them on 3D faces. Figure 11 shows edits over the

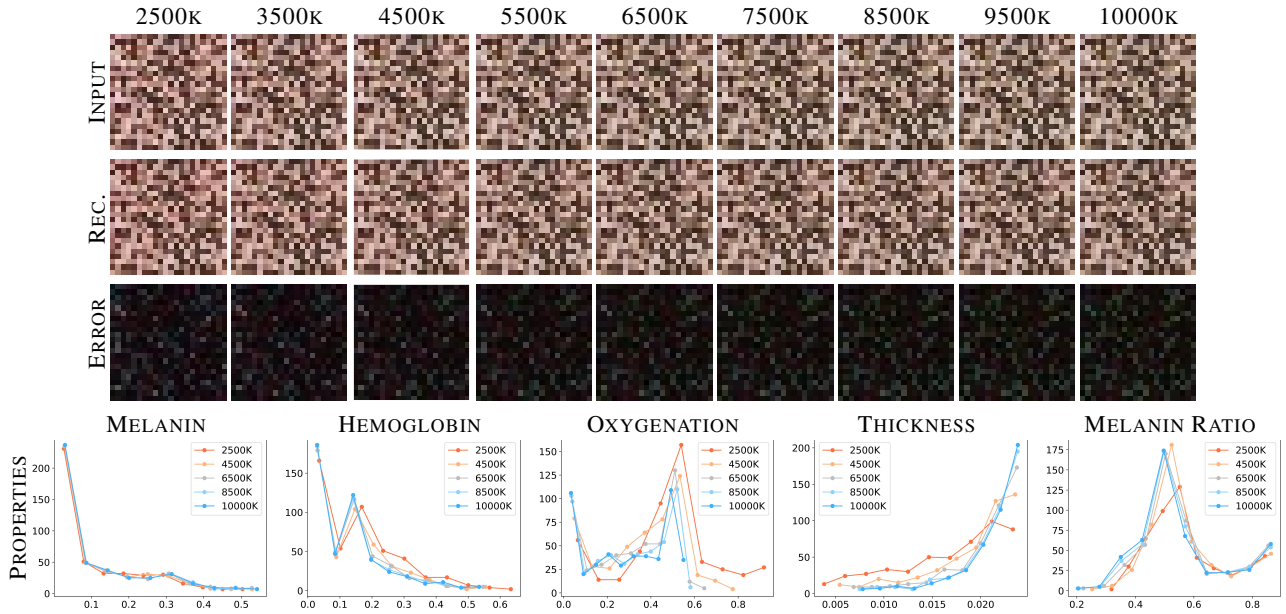


Figure 8: Robustness analysis under different color temperatures. We tested a subset of skin tones for a range of color temperatures from warmest to coolest (2500 to 10000 K), chromatically adapted. The mean absolute error (MAE) of the RGB reconstruction in the visible range remains low, as shown in the third row. The histograms of the estimated skin properties (bottom) for the range of color temperatures illustrate how melanin remains very similar for all illuminants. The remaining properties exhibit only small variations. From coolest to warmest, a) there is a progressive increase in the reconstructed hemoglobin concentration, oxygenation level, and pheomelanin ratio, the three chromophores mostly responsible for redder skin tones; and b) the reconstructed epidermis is thinner for warmer illuminants, which reduces the masking effect of melanin concentration and enhances the effect of the aforementioned chromophores.

rest of the parameters of the model. The level of blood oxygenation generates paler or more saturated skin colors. The thinning of the epidermis, which typically occurs with aging, translates into a more translucent look, revealing the heterogeneities of the underlying layers (e.g. capillary and veins), while a thicker epidermis results in a more opaque and rough appearance. Last, we vary several components to simulate tanning and flushing. While it is hard to fully validate the correctness of the recovered skin properties, we find the parametrization adequate to produce plausible human skin albedos.

7. Discussion and Future Work

We have presented a method for inverse estimation of skin's biophysical parameters, from RGB images of the diffuse reflectance. Our method is based on a carefully designed and validated space for skin reflectance, parametrized by the main biophysical parameters of human skin. As a side effect, our method allows to do a well-principled spectral upsampling of skin reflectances.

While our approach outperforms previous research and works well on a variety of skin tones, it is still not free from limitations. First of all, although we consider all the chromophores in our current model, we aim to widen the albedo space to handle uncommon skin shades resulting from rare chromophore concentrations, such as the excess of bilirubin, beta-carotene, or pathologies related to methemoglobin. Incorporating the effect of hair roots on the light

transport inside the skin would be interesting to better handle short beards or shaved skulls.

Currently we infer the biophysical properties locally, already obtaining coherent spatially-varying parameters. Nonetheless, incorporating the spatial domain in the inference, by introducing spatial priors and correlation between the different parameters of our model, could potentially improve results.

Additionally, our method is robust to different illumination conditions, specially if images are white-balanced. However, for illuminants with very different spectral signature, our neural network trained for D65 illuminants might result in errors. In these cases, re-training the model for new illuminants might be a better choice. Inferring the illuminant at the same time as the skin parameters and occlusion would make our method even more robust, but likely will require some extra information for inference due to metamerism.

Finally, we have not validated our recovered biophysical parameters against measurements in skin samples. Thus, while our recovered parameters roughly agree with measurements we do not claim that we recover them such that they can be used in medical inspection of skin. However, non-invasive in-vivo recovery of these parameters is an important research field, and current solutions present many limitations even in commercial products used in medical applications. For instance, although often relying on large spectral bands where oxy- and deoxyhemoglobin spectra mostly differ (660 and 940 nm) and often installed in regions less likely to be masked by melanin (inside surface of finger tips),

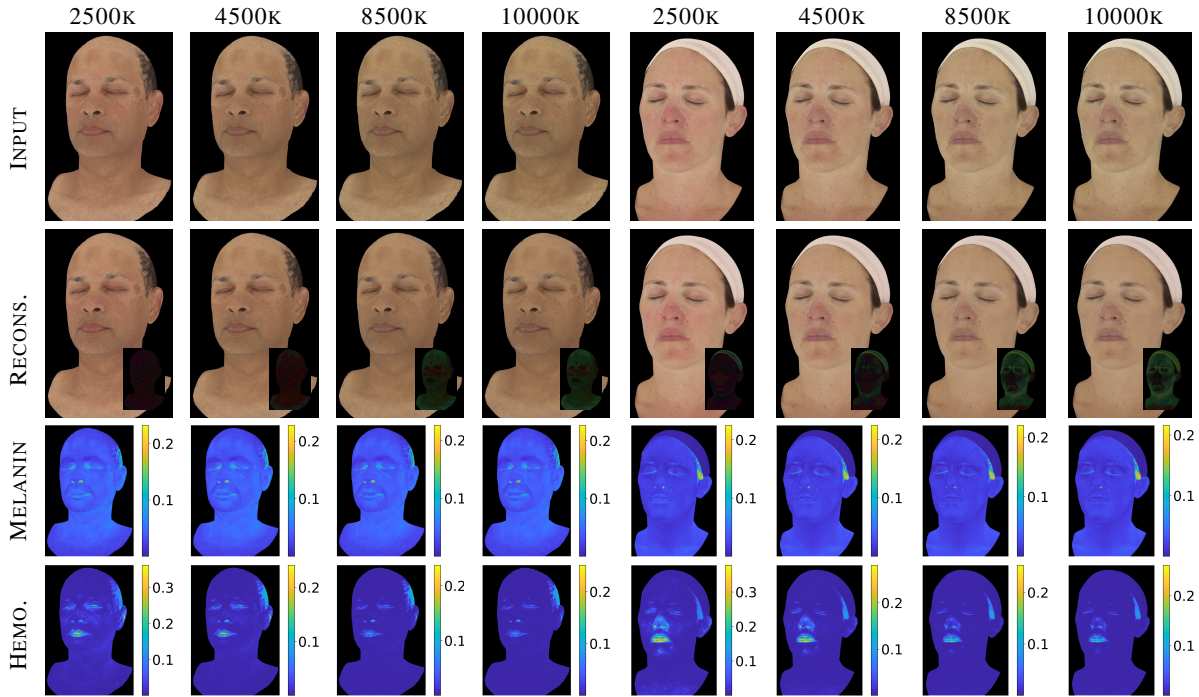


Figure 9: Faces under different color temperatures. For two faces of different skin type and four color temperatures (first row), our model trained for D65 is able to reconstruct the skin reflectance in the visible range (second row) with low error (MAE, shown as inset). The model adapts to warmer and cooler illuminants by slightly adjusting some of the estimated properties. Melanin and hemoglobin shown here, please see Supplemental Figure 6 for the remaining properties. Similarly, the behavior of oxygen saturation aligns with that of hemoglobin concentration. Conversely, epidermal thickness and eumelanin ratio, much like melanin concentration, exhibit minimal variation and remain relatively constant.

pulse oxymeters have shown inaccuracies for different pigmentation levels [SGD*22; FWW*22]. Validating and correcting our model against measurements, potentially incorporating spatial priors as mentioned above, is an exciting avenue of future work, with numerous applications beyond graphics.

Acknowledgments

We thank Olivier Maury and Akira Orikasa for general lookdev support, Christoph Lassner for conversations about the neural model, and Danielle Belko, Arkell Rasiyah, Tony Tung, Ronald Mallet and Yaser Sheikh for their face captures. Scans and actors for Subjects A and E were provided by Eisko.

References

- [AS17] ALOTAIBI, SARAH and SMITH, WILLIAM AP. “A biophysical 3D morphable model of face appearance”. *Proc. of IEEE International Conference on Computer Vision (ICCV) Workshops*. 2017, 824–832 2, 4.
- [Ber19] BERNIS, ROY S. *Billmeyer and Saltzman’s principles of color technology*. John Wiley & Sons, 2019 4, 5.
- [BGT11] BASHKATOV, ALEXEY N, GENINA, ELINA A, and TUCHIN, VALERY V. “Optical properties of skin, subcutaneous, and muscle tissues: a review”. *Journal of Innovative Optical Health Sciences* 4.01 (2011), 9–38 3.
- [BK10] BARANOSKI, GLADIMIR VG and KRISHNASWAMY, ARAVIND. *Light and skin interactions: simulations for computer graphics applications*. Morgan Kaufmann, 2010 1, 2.
- [CBKM15] CHEN, TENN F, BARANOSKI, GLADIMIR VG, KIMMEL, BRADLEY W, and MIRANDA, ERIK. “Hyperspectral modeling of skin appearance”. *ACM Trans. Graph.* 34.3 (2015), 1–14 1–3.
- [CSLM19] CZEKALLA, CAROLIN, SCHÖNBORN, KARL HEINZ, LADEMANN, JÜRGEN, and MEINKE, MARTINA C. “Noninvasive determination of epidermal and stratum corneum thickness in vivo using two-photon microscopy and optical coherence tomography: Impact of body area, age, and gender”. *Skin pharmacology and physiology* 32.3 (2019), 142–150 3.
- [DHT*00] DEBEVEC, PAUL, HAWKINS, TIM, TCHOU, CHRIS, et al. “Acquiring the reflectance field of a human face”. *Proc. of SIGGRAPH*. 2000, 145–156 6.
- [DJ06] DONNER, CRAIG and JENSEN, HENRIK WANN. “A Spectral BSSRDF for Shading Human Skin.” *Rendering techniques*. Citeseer, 2006, 409–418 1–3, 12.
- [DWd*08] DONNER, CRAIG, WEYRICH, TIM, D’EON, EUGENE, et al. “A layered, heterogeneous reflectance model for acquiring and rendering human skin”. *ACM Trans. Graph.* 27.5 (2008), 1–12 2, 3.
- [Fai13] FAIRCHILD, MARK D. *Color Appearance Models*. John Wiley & Sons, 2013 7.
- [FJM*20] FRISVAD, JEPPE REVAL, JENSEN, SØREN ALKÆRSIG, MADSEN, JONAS SKOVLUND, et al. “Survey of Models for Acquiring the Optical Properties of Translucent Materials”. *Computer Graphics Forum* 39.2 (2020) 2.



Figure 10: Antera device comparisons. *Our reconstructions consistently exhibit less errors than in the previous work. From top to bottom the mean square error ([GGD*20] vs. ours) is: $(0.0142 \text{ vs. } 1.65 \cdot 10^{-6})$; $(0.1016 \text{ vs. } 2.31 \cdot 10^{-5})$; $(0.0585 \text{ vs. } 1.08 \cdot 10^{-5})$; and $(0.2239 \text{ vs. } 8.93 \cdot 10^{-5})$.*

- [FWW*22] FAWZY, ASHRAF, WU, TIANSHI DAVID, WANG, KUNBO, et al. “Racial and ethnic discrepancy in pulse oximetry and delayed identification of treatment eligibility among patients with COVID-19”. *JAMA internal medicine* 182.7 (2022), 730–738 10.
- [GAS*19] GEVAUX, LOU, ADNET, CYPRIEN, SÉROUL, PIERRE, et al. “Three-dimensional maps of human skin properties on full face with shadows using 3-D hyperspectral imaging”. *Journal of biomedical optics* 24.6 (2019), 066002 2.
- [GFT*11] GHOSH, ABHIJEET, FYFFE, GRAHAM, TUNWATTANAPONG, BOROM, et al. “Multiview face capture using polarized spherical gradient illumination”. *Proceedings of the 2011 SIGGRAPH Asia Conference*. 2011, 1–10 2.
- [GGD*20] GITLINA, YULIYA, GUARNERA, GIUSEPPE CLAUDIO, DHILLON, DALJIT SINGH, et al. “Practical Measurement and Reconstruction of Spectral Skin Reflectance”. *Computer Graphics Forum*. Vol. 39. 4. Wiley Online Library. 2020, 75–89 2, 8, 11.
- [GGR*21] GEVAUX, LOU, GIERSCHEIDT, JORDAN, RENGOT, JULIETTE, et al. “Real-time skin chromophore estimation from hyperspectral images using a neural network”. *Skin Research and Technology* 27.2 (2021), 163–177 2.
- [GHP*08] GHOSH, ABHIJEET, HAWKINS, TIM, PEERS, PIETER, et al. “Practical modeling and acquisition of layered facial reflectance”. *ACM Trans. Graph.* 27.5 (2008), 1–10 2.
- [GRB*18] GOTARDO, PAULO, RIVIERE, JÉRÉMY, BRADLEY, DEREK, et al. “Practical Dynamic Facial Appearance Modeling and Acquisition”. *ACM Trans. Graph.* 37.6 (Dec. 2018) 2.

- [GTB*13] GRAHAM, PAUL, TUNWATTANAPONG, BOROM, BUSCH, JAY, et al. “Measurement-based synthesis of facial microgeometry”. *Computer Graphics Forum*. Vol. 32. 2pt3. Wiley Online Library. 2013, 335–344 2.
- [Hal64] HALTON, JOHN H. “Algorithm 247: Radical-inverse quasi-random point sequence”. *Communications of the ACM* 7.12 (1964), 701–702 6.
- [IAJG15] IGLESIAS-GUITIAN, JOSE A., ALIAGA, CARLOS, JARABO, ADRIAN, and GUTIERREZ, DIEGO. “A Biophysically-Based Model of the Optical Properties of Skin Aging”. *Computer Graphics Forum* 34.2 (2015) 1–3.
- [INN07] IGARASHI, TAKANORI, NISHINO, KO, and NAYAR, SHREE K. *The appearance of human skin: A survey*. Now Publishers Inc, 2007 2.
- [IRB02] IMAI, FRANCISCO H, ROSEN, MITCHELL R, and BERNS, ROY S. “Comparative study of metrics for spectral match quality”. *Conference on colour in graphics, imaging, and vision*. Vol. 2002. 1. Society for Imaging Science and Technology. 2002, 492–496 5.
- [Jac13] JACQUES, STEVEN L. “Optical properties of biological tissues: a review”. *Physics in Medicine & Biology* 58.11 (2013), R37 3, 12.
- [JAG18] JARABO, ADRIAN, ALIAGA, CARLOS, and GUTIERREZ, DIEGO. “A radiative transfer framework for spatially-correlated materials”. *ACM Trans. Graph.* 37.4 (2018), 1–13 3.
- [JAM*10] JAKOB, WENZEL, ARBREE, ADAM, MOON, JONATHAN T, et al. “A radiative transfer framework for rendering materials with anisotropic structure”. *ACM Trans. Graph.* 29.4 (2010), 1–13 3.
- [JH19] JAKOB, WENZEL and HANIKA, JOHANNES. “A Low-Dimensional Function Space for Efficient Spectral Upsampling”. *Computer Graphics Forum* 38.2 (Mar. 2019) 2.
- [JM91] JACQUES, STEVEN L and MCAULIFFE, DANIEL J. “The melanosome: threshold temperature for explosive vaporization and internal absorption coefficient during pulsed laser irradiation”. *Photochemistry and photobiology* 53.6 (1991), 769–775 12.
- [JSB*10] JIMENEZ, JORGE, SCULLY, TIMOTHY, BARBOSA, NUNO, et al. “A practical appearance model for dynamic facial color”. *ACM Trans. Graph.* 29.6 (2010) 2.
- [KB04] KRISHNASWAMY, ARAVIND and BARANOSKI, GLADIMIR VG. “A biophysically-based spectral model of light interaction with human skin”. *Computer Graphics Forum*. Vol. 23. 3. Wiley Online Library. 2004, 331–340 1–3.
- [KB14] KINGMA, DIEDERIK P and BA, JIMMY. “Adam: A method for stochastic optimization”. *arXiv preprint arXiv:1412.6980* (2014) 6.
- [KLB*93] KRUSE, FRED A, LEFKOFF, AB, BOARDMAN, JW, et al. “The spectral image processing system (SIPS)—interactive visualization and analysis of imaging spectrometer data”. *Remote sensing of environment* 44.2-3 (1993), 145–163 6.
- [KRP*15] KLEHM, OLIVER, ROUSSELLE, FABRICE, PAPAS, MARIOS, et al. “Recent Advances in Facial Appearance Capture”. *Computer Graphics Forum* 34.2 (May 2015), 709–733 2.
- [LU14] LE MOAN, STEVEN and URBAN, PHILIPP. “Image-difference prediction: from color to spectral”. *IEEE Transactions on Image Processing* 23.5 (2014), 2058–2068 4.
- [LWC12] LISTER, TOM, WRIGHT, PHILIP A., and CHAPPELL, PAUL H. “Optical properties of human skin”. *Journal of Biomedical Optics* 17.9 (2012), 1–15 3.
- [Mac35a] MACADAM, DAVID L. “Maximum visual efficiency of colored materials”. *JOSA* 25.11 (1935), 361–367 2.
- [Mac35b] MACADAM, DAVID L. “The theory of the maximum visual efficiency of colored materials”. *JOSA* 25.8 (1935), 249–252 2.
- [MHP*07] MA, WAN-CHUN, HAWKINS, TIM, PEERS, PIETER, et al. “Rapid Acquisition of Specular and Diffuse Normal Maps from Polarized Spherical Gradient Illumination.” *Rendering Techniques*. 2007, 10 2.



Figure 11: Further examples of parameter manipulation. From left to right, a) original albedo, b) fully deoxygenated and c) oxygenated blood, epidermal thinning d) and thickening e) (minimum and maximum respectively), f) tanning (40% melanin increase, full pheomelanin), g) flushing (70% increase blood, fully oxygenated), and h) simulated vitiligo through edited melanin concentration.

- [MM02] MEGLINSKI, IGOR V and MATCHER, STEPHEN J. “Quantitative assessment of skin layers absorption and skin reflectance spectra simulation in the visible and near-infrared spectral regions”. *Physiological measurement* 23.4 (2002), 741–754.
- [MSHD15] MENG, JOHANNES, SIMON, FLORIAN, HANIK, JOHANNES, and DACHSBACHER, CARSTEN. “Physically meaningful rendering using tristimulus colours”. *Computer Graphics Forum* 34.4 (2015), 31–40.
- [MWL*99] MARSCHNER, STEPHEN R, WESTIN, STEPHEN H, LAFORTUNE, ERIC PF, et al. “Image-based BRDF measurement including human skin”. *Rendering Techniques*. Springer, 1999, 131–144.
- [NMM*19] NUNES, AUGUSTO LP, MACIEL, ANDERSON, MEYER, GARY W, et al. “Appearance modelling of living human tissues”. *Computer Graphics Forum*. Vol. 38. 6. Wiley Online Library. 2019, 43–65.
- [OYH18] OTSU, HISANARI, YAMAMOTO, MASAFUMI, and HACHISUKA, TOSHIYA. “Reproducing spectral reflectances from tristimulus colours”. *Computer Graphics Forum* 37.6 (2018), 370–381.
- [PJ] PRAHL, SCOTT and JACQUES, STEVEN L. *Tabulated Molar Extinction Coefficient for Hemoglobin in Water*. <https://omlc.org/spectra/hemoglobin/summary.html> 12.
- [RGB*20] RIVIERE, JÉRÉMY, GOTARDO, PAULO, BRADLEY, DEREK, et al. “Single-shot high-quality facial geometry and skin appearance capture”. *ACM Trans. Graph.* 39.4 (2020), 81–12.
- [Sai92] SAIDI, IYAD SALAM. “Transcutaneous optical measurement of hyperbilirubinemia in neonates”. PhD thesis. 1992.
- [SGD*22] SHI, CHUNHU, GOODALL, MARK, DUMVILLE, JO, et al. “The accuracy of pulse oximetry in measuring oxygen saturation by levels of skin pigmentation: a systematic review and meta-analysis”. *BMC medicine* 20.1 (2022), 267–280.
- [Smi99] SMITS, BRIAN. “An RGB-to-spectrum conversion for reflectances”. *Journal of Graphics tools* 4.4 (1999), 11–22.
- [SS06] SARNA, TADEUSZ and SWARTZ, HAROLD A. “The physical properties of melanins”. *The pigmentary system: physiology and pathophysiology* (2006), 311–341.
- [THM99] TSUMURA, NORIMICHI, HANEISHI, HIDEAKI, and MIYAKE, YOICHI. “Independent-component analysis of skin color image”. *JOSA A* 16.9 (1999), 2169–2176.
- [TOS*03] TSUMURA, NORIMICHI, OJIMA, NOBUTOSHI, SATO, KAYOKO, et al. “Image-based skin color and texture analysis/synthesis by extracting hemoglobin and melanin information in the skin”. *Proc. of SIGGRAPH*. 2003, 770–779.
- [TWF21] TÓDOVÁ, LUCIA, WILKIE, ALEXANDER, and FASCIONE, LUCA. “Moment-based Constrained Spectral Uplifting”. *Proc. of Eurographics Symposium on Rendering (EGSR)*. 2021, 215–224.
- [VJSS89] VAN GEMERT, MJC, JACQUES, STEVEN L, STERENBORG, HJCM, and STAR, WM. “Skin optics”. *IEEE Transactions on biomedical engineering* 36.12 (1989), 1146–1154.
- [vKri05] VON KRIES, JOHANN. “Die gesichtsempfindungen”. *Handbuch der physiologie des menschen* (1905) 7.
- [WVH17] WRENNINGE, MAGNUS, VILLEMIN, RYUSUKE, and HERY, CHRISTOPHE. *Path traced subsurface scattering using anisotropic phase functions and non-exponential free flights*. Tech. rep. Technical Memo, 2017.
- [XFM20] XIE, HAO, FARNAND, SUSAN P, and MURDOCH, MICHAEL J. “Observer metamerism in commercial displays”. *JOSA A* 37.4 (2020), A61–A69.
- [XYZ*17] XIAO, KAIDA, YATES, JULIAN M, ZARDAWI, FARAEDON, et al. “Characterising the variations in ethnic skin colours: a new calibrated data base for human skin”. *Skin Research and Technology* 23.1 (2017), 21–29.
- [ZDP*19] ZHEREBTSOV, EVGENY, DREMIN, VIKTOR, POPOV, ALEXEY, et al. “Hyperspectral imaging of human skin aided by artificial neural networks”. *Biomedical optics express* 10.7 (2019), 3545–3559.

Appendix A: Specifications of Skin Chromophores

Parameter	Description	Value
$\mu_{a_{hbO_2}}$	Oxy-Hemoglobin absorption	$2.303 \frac{p_{hb} \epsilon_{hbO_2}}{w_{hb}}$
$\mu_{a_{hb}}$	Deoxy-Hemoglobin absorption	$2.303 \frac{p_{hb} \epsilon_{hb}}{w_{hb}}$
ϵ_{hbO_2}	Oxy-Hemoglobin Extinction	[Jac13]
ϵ_{hb}	Deoxy-Hemoglobin Extinction	[Jac13]
p_{hb}	Hemoglobin Concentration	150
w_{hb}	Molar weight of Hemoglobin	64500
$\mu_{a_{eu}}$	Eumelanin absorption	$6.6 \times 10^{11} \lambda^{-3.33}$
$\mu_{a_{phoe}}$	Pheomelanin absorption	$2.9 \times 10^{15} \lambda^{-4.75}$
p_{bil}	Bilirubin Concentration	0.05
w_{bil}	Molar weight of bilirubin	584.66
$p_{\beta-c_e}$	β -carotene Concentration (Epidermis)	2.1×10^{-4}
$p_{\beta-c_d}$	β -carotene Concentration (Dermis)	7×10^{-5}
$w_{\beta-c}$	Molar weight of β -carotene	536.8726

Table 3: Chromophore specifications. The absorption coefficient is defined in cm^{-1} ; the extinction coefficient ϵ_c in $\frac{\text{cm}^{-1}}{\text{moles/liter}}$; the concentration of the chromophore p_c in g/L; and the molar weight w_c in g/mol. The absorption of melanins $\mu_{a_{eu}}$ and $\mu_{a_{phoe}}$ is defined through a fit from [DJ06] to the measurements from [JM91] and [SS06] respectively. The 2.303 coefficient in $\mu_{a_{hb}}$ comes from deriving a factor of $\ln(10)$, since ϵ has been historically recorded in such base 10 nomenclature from measurements of old spectrometers in literature. Finally, oxy- and deoxy-hemoglobin extinction can be found tabulated in [PJ].




Figure 12: Ablation study: parameter maps for Subjects A and C. *The SPECTRAL+ model retrieves more even skin properties. The hemoglobin concentration and oxygenation maps better match the reddish areas of the face, whereas RGB and SPECTRAL wrongly interpret the shadows around the jaw as higher concentration of blood. The epidermal thickness and eumelanin ratio are also more uniform in the SPECTRAL+ approach, while keeping low values in the right areas, like the lips.*