

Estimation of Spectral Biophysical Skin Properties from Captured RGB Albedo

C. Aliaga¹ C. Hery¹ M.Xia^{1,2}
Meta Reality Labs Research¹ Cornell University²

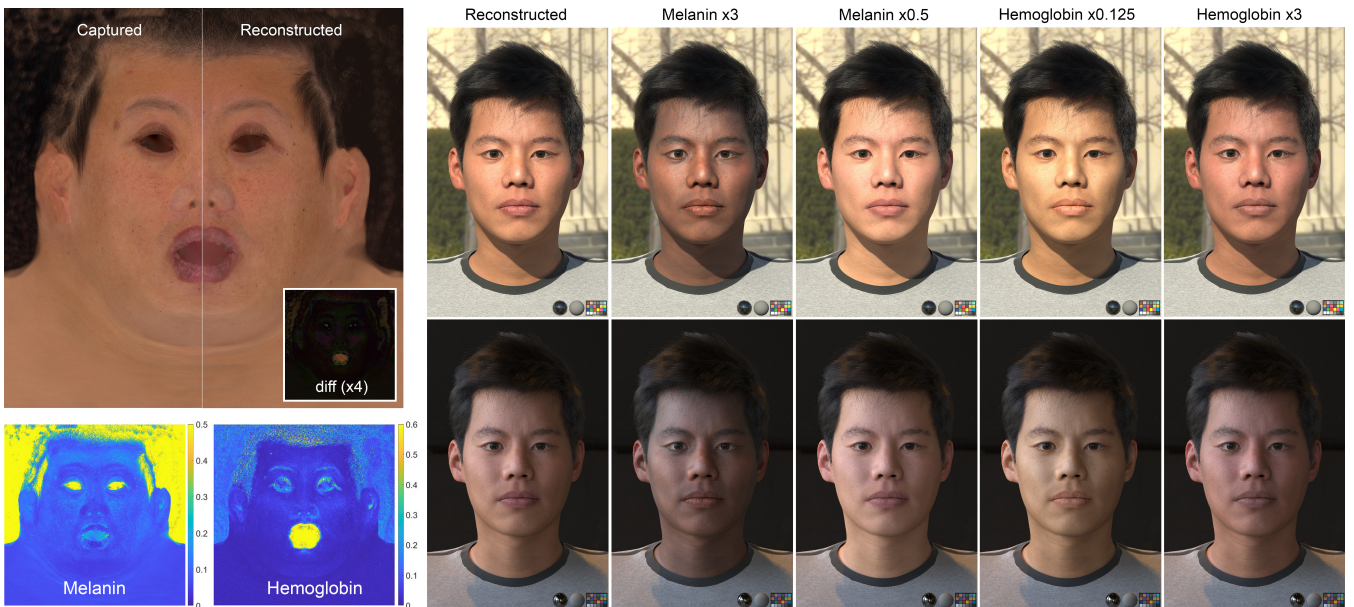


Figure 1: Our model takes RGB albedo captures as input, and reconstructs a set of biophysical properties of human skin. On the left, an example of a captured albedo vs the reconstruction of our model (top) using the estimated skin properties, of which melanin and hemoglobin are shown in the bottom. On the right, renderings using albedos modified through edits of some of the estimated parameters, under two different lighting setups (top and bottom rows).

Abstract

We present a new method to reconstruct and manipulate the spectral properties of human skin from simple RGB albedo captures. To this end, we leverage Monte Carlo light simulation over an accurate biophysical human skin layering model parameterized by its most important components, thereby covering a plausible range of colors. The practical complexity of the model allows us to learn the inverse mapping from any albedo to its most probable associated skin properties. Our technique can faithfully reproduce any skin type, being expressive enough to automatically handle more challenging areas like the lips or imperfections in the face. Thanks to the smoothness of the skin parameters maps recovered, the albedo can be robustly edited through meaningful biophysical properties.

CCS Concepts

• Computing methodologies → Reflectance modeling; Reconstruction;

1. Introduction

Creating convincing photo realistic replicas of human faces has been a long-standing goal of computer graphics, usually relying on skilled artists to fine tune a large number of spatially varying

shading parameters. As a crucial part of virtual humans, significant effort has been put on accurately modeling skin. Skin is not only challenging due to its complex interaction with light, but also because the human visual system has specifically evolved for strong facial perception. A number of models have targeted the effect of

the biophysical composition of the skin on its final appearance, often centered in its diffuse reflectance resulting from the complex interaction of light inside the dermal tissues [DJ06; KB04; BK10; CBKM15; IAJG15]. More recently, the interest in estimating skin properties from captures has grown, mostly focusing on inverting the biophysical properties of skin from its diffuse reflectance, leveraging biophysical constraints [AS17; GAS*19; ZDP*19; GGD*20; GGR*21].

We present a new model that, given a biophysically-inspired space of spectral skin albedos, makes use of a neural network to accurately recover, out of simple RGB albedos, the main components and structure of the skin. For this, our model builds upon a biophysical description of the main properties in human skin, which we translate into albedo by means of Monte Carlo simulations.

Thus, our high-level contributions can be summarized as follows:

- An expressive space of spectral skin albedo constrained by physical properties of real skin, in agreement with measurements reported in tissue optics and medical research. Such space is created through a biophysically-inspired model of human skin of practical complexity.
- A learned inverse mapping from skin albedo to its associated spectral biophysical skin properties, that enables the recovery of smooth, high resolution spatially varying maps of skin properties from RGB captured albedos.
- Altogether, a framework to faithfully reproduce captured albedos of a wide range of skin types with minimal error, also enabling editing capabilities in a robust way through the estimated biophysical properties.

The process is composed of two steps. First, we precompute an albedo space in the form of a high dimensional tensor of skin tones resulting from all possible combinations of skin parameters. The details of our biophysical skin model are described in Section 3. Then, we learn the inverse mapping from an albedo to its associated skin properties in Section 4. After discussion of implementation details in Section 5, we show how to identically reconstruct the albedo and the results of manipulating the skin parameters in Section 6, with comparisons to related work in Section 7.

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 further details in skin and face appearance modeling [INN07; NMM*19], face appearance capture [KRP*15; WLL*09], or a broader collection of methods acquiring the optical properties of translucent materials [FJM*20].

Modeling. There has been a bunch of works focusing on biophysical skin components. Some of these isolated the distributions of melanin and hemoglobin in human skin [THM99; TOS*03]. After an early model first introduced multi-spectral skin color [AMD01], spectral techniques that explicitly drive the appearance through biophysical properties have been developed, spanning models of different complexity, ranging from two layers and the two most relevant components (melanin and haemoglobin) [DJ06], up to seven

layers and accounting for a large number of parameters related to the chemical composition and structure of human skin [CBKM15; IAJG15; KB04; BK10]. Or even a model for dynamic facial color [JSB*10], built from in vivo measurements of melanin and hemoglobin concentrations.

Capturing. Since early efforts in computer graphics to acquire the skin BRDF [MWL*99], a series of approaches focused on human face acquisition, progressively improving from the original Light Stage device [DHT*00]. Relying on such Light Stage-type scans, a number of models reconstructed the reflectance functions directly from captures [MHP*07; GHP*08; GFT*11; GTB*13], or synthesized high-resolution facial surface microgeometries [GTB*13]. Some other subsequent publications expanded on the facial appearance modeling and capture, including dynamic diffuse albedo encoding blood flow, specular intensity or high resolution normals [GRB*18; GRB*20], even single-shot [RGB*20], with editing capabilities [WMP*06]. Another method [DWD*08] used multi spectral photographs to recover spatially varying biophysical components over a layered skin model. A number of works opted to rely on Kubelka-Munk (KM) theory [KM31] to recover the skin parameter maps. An example [AS17] of this kind of models is tested and discussed in Supplementary Section 2.2. In skin research, 3D maps of human skin properties on the full face with shadows can be reconstructed from 3D hyper spectral imaging [GAS*19] via a two-flux KM model. Based on this, neural networks were trained to recover maps of oxygen rate, blood volume fraction, melanin concentration, bilirubin concentration, and epidermal thickness [GGR*21]. Also relying on hyperspectral imaging, a neural network and a more complex 7 layer skin model were presented [ZDP*19] to extract 2D maps of blood volume fraction, blood oxygen saturation, melanin content and epidermal thickness. It demonstrated the importance of the thickness of the superficial bloodless layer. However, both approaches [GGR*21; ZDP*19] lack albedos reconstructed by the model and the resulting estimated parameters, making it difficult to assess its application in a vfx or virtual reality context. Closest to our solution, a recent paper [GGD*20] reconstructs spectral skin properties through a novel measurement system, benefiting from narrow-band LEDs, based on a simplified skin model [JSB*10; DJ06].

Differently to previous works, our model is capable of faithfully reproducing with minimal error any skin type by a set of biophysical properties proven to be relevant. It relies on Monte Carlo simulations to accurately model light-skin interactions, avoiding assumptions (e.g. epidermal blood [GGD*20]) and corrections needed for KM approaches [GGR*21; Sau42]. Also, our model does not require hyperspectral captures, since it operates directly on simple RGB albedos (although it can be extended to spectral captures), and recovers smooth parameter maps that allow for intuitive edits in such parameter space.

3. A Space of Biophysically-Based Skin Albedos

In this section we describe how we create our albedo space, the details of the skin model and the data employed, together with balance of practical complexity and expressiveness.

3.1. Skin Model: Structure and Optical Properties

3.1.1. Skin Structure

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 fits our purpose. The epidermis is composed of two parts: the living epidermis and the *stratum corneum*. The latter is the outermost part, and exhibits properties (surface roughness and sebum production) that affect the specular reflectance of the skin. However, it has a minimal effect on the albedo of the skin, given its low absorption and relative small thickness (5 and 20 μm depending on the body location [CSLM19]).

On the other hand, we model the dermis as a single layer of averaged properties of scattering and absorption of its two sub-layers, the reticular dermis and the papillary dermis. It is modeled as a semi infinite medium, omitting the sub-dermal fat layer. This was decided to keep the generality of the model, with a focus on faces: next to the dermis can be not only fat, but other internal tissues like cartilage or muscles, heavily depending on the anatomical location and composition of the subject. Also, we empirically found that including the dermal thickness has a minimal effect over the resulting diffuse reflectance.

3.1.2. Absorption

Following well known multi-layered tissue optics models [Jac13], we describe the optical properties of each layer by its spectral absorption (μ_a) and scattering (μ_s) coefficients. 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 2 in Appendix A).

As a result, the total spectral absorption of epidermis is defined as $\mu_{a_e} = V_m (\phi_m \mu_{a_{eu}} + (1 - \phi_m) \mu_{a_{ph}}) + (1 - V_m) (\mu_{a_{\beta-c}} + \mu_{a_{base}})$, and the spectral absorption of the dermis is defined as $\mu_{a_d} = V_b (\phi_h \mu_{a_{hb}} + (1 - \phi_h) \mu_{a_{hbO_2}} + \mu_{a_{bil}} + \mu_{a_{\beta-c}}) + (1 - V_b) \mu_{a_{base}}$.

In the same spirit of biophysical 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 haemoglobin present in blood V_b in the dermis: oxygenated haemoglobin $\mu_{a_{hbO_2}}$, responsible for the saturated reddish tint, and deoxygenated haemoglobin $\mu_{a_{hb}}$, responsible for a desaturated purple color; ϕ_h being the haemoglobin 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.

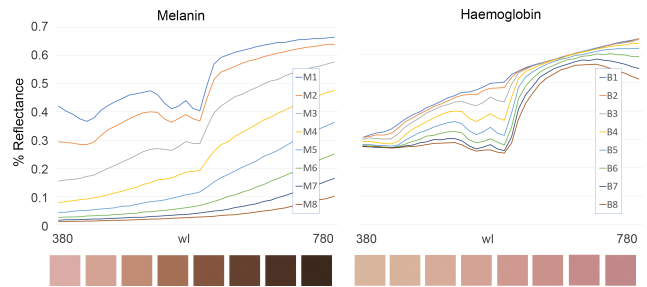


Figure 2: Melanin and haemoglobin have a highly non-linear relationship with the final albedo. In this example, we vary melanin and haemoglobin to show the effect on the spectral reflectance (top) and corresponding final albedos (bottom). Left: Melanin is cubically sampled with $M_n = (0.2, 0.7, 2.3, 5.4, 10.6, 18.2, 29.9, 43.1)\%$, while haemoglobin is fixed at 2%. Right: haemoglobin is quartically sampled with $B_n = (0.1, 0.2, 0.7, 2, 4.5, 9.6, 17.7, 30)\%$, while melanin is fixed at 0.8%. These spectral reflectance plots are qualitatively in agreement with previous models of computer graphics [KB04], optics [ZDP*19], color research [VG194], computer vision [AMD01], and dermatology [AP81].

3.1.3. Scattering

We treat epidermis and dermis as homogeneous mediums, the later being semi infinite. Both layers have an index of refraction of 1.4 resulting from the weighted (by thickness) sum of the corresponding sub layers: *stratum corneum* (1.53), living epidermis (1.34), papillary dermis (1.39) and reticular dermis (1.395). The interfaces of the two layers, as well as other details of multiple scattering in the tissue, are discussed in Section 3.2.

The reduced scattering coefficient μ'_s is modeled through a function of wavelength fitted by Jacques [Jac13], which is generic for a wide range of human tissues :

$$\mu'_s(\lambda) = a \left(f_{Ray} \left(\frac{\lambda}{\lambda_r} \right)^{-4} + (1 - f_{Ray}) \left(\frac{\lambda}{\lambda_r} \right)^{-b_{Mie}} \right) \quad (2)$$

where the wavelength λ is normalized by a reference wavelength $\lambda_r = 500$ nm, the scaling factor $a = \mu'_s(\lambda_r)$, f_{Ray} is the relative contribution of Rayleigh scattering, and b_{Mie} characterizes the wavelength dependence of the Mie scattering component. We rely on the coefficients reported in the optics literature [BGT11], leading to $a = 36.4$, $f_{Ray} = 0.48$, and $b_{Mie} = 0.22$ for human skin. In addition, we include the spectral dependency of the anisotropy factor g , which has been measured in the past [VJSS89], ranging from 0.73 at 380 nm to 0.84 at 780 nm. It is expressed as $g_e \sim g_d \sim 0.62 + \lambda 0.29 \times 10^{-3}$.

3.2. Layered Model for Albedo Generation

For computing the diffuse reflectance of a skin patch, we first attempted to use a Kubelka-Munk - based layering model [AS17]. However, it suffers from lack of expressiveness to robustly recover the parameters of a wide enough range of skin types. Details on such exploration can be found in the Supplemental, Section 2.2. Thus, we decided to shift to a brute force Monte Carlo random walk. At this stage, we focus on resolving the transports through the

Parameter	Description	Epidermis	Dermis
V_m	Melanin Volume Fraction	[0.001, 1]	-
V_b	Blood Volume Fraction	-	[0.001, 1]
t	Thickness [μm]	[10, 250]	2100
ϕ_m	Ratio of melanin types	[0.001, 1]	-
ϕ_h	Ratio of haemoglobin types	-	[0.001, 1]

Table 1: Ranges for the skin properties of our 5D albedo space.

stacking of the different layers of the skin structure (3.1.1). We conduct the walk in 2D, assuming symmetry in the azimuthal planes of each layer. The interface between the epidermis and dermis only considers the change in scattering and absorption parameters, since the index of refraction is measured to be mostly the same in both layers [LWC12].

For each skin tone, we run spectral simulations for the wavelengths comprised in the visible range, between 380 to 780 nm. We empirically found that steps of 10 nm is enough to produce stable, noise free albedos. For each wavelength, a million photons are launched, simulating the 2D random walk over the two-layered, semi infinite medium, to produce diffuse albedos (by opposition to the rendering step, where we make use of such uv-mapped albedos to conduct true 3D random walks, as outlined in Section 6). We treat the skin as an homogeneous medium with exponential decay, since the more recent statistical theories attempting to account for spatial correlations in some materials [WVH17; BRM*18; dEon19] do not have corresponding measured data of particle distributions for tissue. We use the multiple scattering parameters listed previously, including the wavelength dependent anisotropy factor g from the Henyey-Greenstein [HG41] phase function, expressed in its 2D form [dEo21]. Note that the simulation starts once the photon has crossed the outermost interface and scatters diffusely into the tissue, namely choosing the initial direction from a Lambertian distribution around the inverse surface normal, in agreement with what we later use at rendering time (Section 6).

We validated the biophysical model by plotting the spectral reflectance of different skin tones, showing that our results are in agreement with existing measurements and simulations found in the literature (see Figure 2).

3.3. Albedo Space Parametrization and Sampling

Our biophysically-based albedo space is created by varying the skin properties in the ranges listed in Table 1. Note that we consider epidermal thickness, which proved to be critical for parameter estimation [ZDP*19]. Accounting for this parameter over the face, along with varying values in melanin concentrations, helps us achieve local dark zones, such as moles, and generalizes over any skin type. We also allow melanin and haemoglobin 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.




Figure 3: Matrices of resulting skin albedos, built by varying melanin and blood concentrations in subsets (here [0.1, 43] and [0.1, 30] % for visualization purposes) of the ranges employed for our final space referenced in Table 1. Epidermal thickness, melanin type ratio and blood oxygenation significantly affect the appearance of the skin.

Thus, some parameters like bilirubin and β -carotene concentrations, remain fixed to common values of human skin measurements found in literature (see Table 2 in Appendix A). The remaining 5D parameter space is sampled as follows: melanin and haemoglobin are respectively selected cubically ($\sqrt[3]{V_m}$) and quartically ($\sqrt[4]{V_b}$), to better adjust to their non linear effect on the albedo (see Figure 2); epidermal thickness, melanin type ratio and haemoglobin type ratio all are treated uniformly.

4. Recovering the Skin Properties

Our biophysical model defines the forward mapping from the skin parameters to the skin albedo. For the inverse process, we need to characterize the mapping from the skin albedo back to skin parameters. The non-bijective nature of such mapping, where many combinations of skin properties can lead to the same albedo, makes this task challenging.

Metamerism. It is important to note that our model operates in spectral space to compute the reflectance corresponding to each skin parameter combination, but then such reflectance is integrated to RGB to learn the mapping. This is a design choice: it simplifies the capture, and provides good enough recovered skin parameter maps to both a) reconstruct the original albedo and b) perform plausible edits. However, operating in a higher dimensional space could alleviate metamerism, which could help to disentangle more accurately the skin parameters, generalize the model to any illuminant, or be robust to different levels of exposure, among others. Out of scope for this paper, further exploration can be done inspired by the literature in Spectral Reconstruction from RGB, where significant efforts have been done, specially in data-driven deep learning methods (a recent survey covers the wide variety of existing techniques [ZSR*21]).

4.1. Look-up Tensor (LUT).

We pre-compute a wide tensor of skin tones by following the sampling strategy outlined in Section 3.3. To estimate the skin parameters given an input albedo, we search the LUT on each texel of the albedo to find the best skin parameter set that minimizes a reconstruction L_2 error. Then, we can manipulate them and query the new corresponding albedo from the LUT. Some slices of the 5D tensor are shown in Figure 3. This approach is able to reconstruct the skin albedo faithfully with an error close to zero. However, the

inverted parameter maps are noisy and have many discontinuities, since the mapping from RGB to skin parameters is not smooth. In turn, editing operations over neighboring pixels in such extracted components can lead to unexpected abrupt changes in the reconstructed albedos (see Figure 4). We leave out of scope of this paper the assessment of different representations or data structures suitable for more efficient search strategies, that could dramatically improve the performance of this approach.




Figure 4: The LUT approach suffers from quantization in the estimated parameters maps even for large tensors. This issue does not reveal during reconstruction, but creates artifacts and even some perceptual color shifting when edits on the skin properties are performed. An example of a $\times 3$ edit over the reconstructed melanin concentration of a Type III skin is shown. From left to right: a) original image, and reconstructions (edited melanin $\times 3$) using a b) tensor of 55296 skin tones $(V_m, V_b, t, \varphi_m, \varphi_h) = (64, 32, 3, 3, 3)$, c) tensor of 256k skin tones $(64, 32, 5, 5, 5)$, d) learned inverse mapping using 600k points to train the network. Whereas even for densely sampled tensors quantization appears, the neural approach provides smooth maps for the estimated parameters, which results in clean edits.

4.2. Neural Skin Parameter Estimation

To overcome the limitations of the LUT approach, we opted to train an encoder-decoder network to recover smooth skin parameter maps. Both the encoder and the decoder are MLPs of four fully connected layers, with two hidden layers of 70 neurons each. The encoder maps a 3D skin albedo to a 5D skin parameter vector, while the decoder does the inverse.

Dataset. Using our biophysical skin model, we generate a data set of 600k pairs of 5D skin parameter vector and the corresponding 3D albedo, split 80% and 20% for training and validation re-

spectively. 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 schemes detailed in Section 3.3, and the corresponding albedos are computed using the biophysical model.

Training. We used the Adam optimizer [KB14], with an learning rate of 10^{-4} and a batch size of 4096 for an optimal number of 400 epochs to prevent over fitting. A grid search was performed to find the optimal hyper parameters listed, including the number of hidden neurons. The loss consists of three parts:

$$\mathcal{L} = \mathcal{L}_{param} + \mathcal{L}_{albedo} + \mathcal{L}_{cycle} \quad (3)$$

where \mathcal{L}_{param} denotes the parameter loss (encoder), which computes the $L2$ difference between the predicted skin properties from the input albedo and its corresponding ground truth properties. \mathcal{L}_{albedo} (decoder) computes the $L1$ differences between the predicted albedo and the ground truth albedo corresponding to the input skin properties. \mathcal{L}_{cycle} evaluates the full encoder-decoder cycle, computing the $L1$ difference between the predicted albedos from the predicted skin properties, and the ground truth albedos. We chose $L1$ loss for robustness against outliers in albedo. Note that, although in theory, a given albedo can invert to different sets of skin parameters, we use the skin parameter set that generates this albedo from the Monte Carlo simulation; this input serves as weak supervision and works well in practice.

5. Implementation Details

Spectral Downsampling. The albedo space is generated through spectral computations, but input albedos are in RGB space. To downsample the multi-band spectral values into RGB, we use an existing integration approach [MY19] (note that we could rely on other concurrent works [JH19] to this end). Most color spaces will work here (we use sRGB), considering that diffuse albedos have rather limited gamuts and dynamic ranges. We perform a change of illuminant for direct comparisons to previous work in Section 7, which we further detail in the Supplementary, Section 1.

Using Albedo Maps in Rendering. The final 3D lit geometries of our virtual faces are rendered using our own skin materials inside *Blender Cycles* [Ble20]. Aside from the specular component of the skin, which we represent as a double lobe GGX [WMLT07], we follow, in the spirit of a state of the art technique in production [WVH17], a 3D random walk subsurface scattering solution that relies on a numerical albedo inversion around the mean free path and accounting for the anisotropy factor g (see Supplemental Section 2.1). At this stage, we simplify the model to be single layered, dermis and epidermis combined, as a semi infinite medium. Obviously multi-layered models could be employed, for instance directly consuming the chromophores estimations, but these are out of scope of this paper. See Figure 1, Figure 9 and the Supplemental material for examples of path traced renders under 3 different lighting environments.

6. Results

We perform estimations and manipulations of skin parameters over several skin types covering the Fitzpatrick scale [Fit88]. With the LUT approach, using the largest tensor (256k skintones) resulted in varying times from 2 to 5 hours for 2k by 2k images, or more than 7 hours for 4k by 4k images, using brute force multi threaded (12) search on the tensor in an Intel Xeon W-2135 at 3.70GHz. Instead, the learned mapping has the advantages of little memory consumption and efficient computation (less than 2 seconds on average for 2k by 2k images, on the same Intel Xeon CPU mentioned above). We refer the readers to Figure 4, Figure 7, Figure 8, Figure 9 and the Supplemental Material to see the neural results. In Figure 7, we observe that we trade a tiny bit of reconstruction correctness (imperceptible in final renders), in favor of spatial regularization and performance.

Editing the Skin Parameters 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. For skin types ranging from I to V, we perform large edits in haemoglobin and melanin content, with details explained in Figure 9. Note the edits are naive in order to cover similar ranges for all skin types. Figure 10 shows edits over the rest of the parameters of the model. The level of blood oxygenation translates into 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. Comparisons with Previous Work

We conduct a series of comparisons with the related and recent work from [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 5 is a comparison on more true albedos from [GGD*20], which were obtained via the Antera device (under D65 illuminant). In Figure 6 we process one of the D65' faces. This was a bit more complex, because of the conversion of the color spaces involved in some of their captures (which we outline in Supplemental). 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.

8. Discussion and Future Work

In this paper, we presented our framework to recover and robustly manipulate the key biophysical properties of human skin from a given RGB albedo. This is accomplished through the combination of an expressive biophysically-inspired skin model, together with an encoder-decoder structure that maps albedo to skin properties and back. We demonstrate how such learned mapping overcomes the limitations of a Look-Up Tensor approach, proving how the



Figure 5: Antera device comparisons. *Our reconstructions consistently exhibit less errors than in the previous work: MSE ([GGD*20], ours) of the skin patches from top to bottom: (0.0142, 1.65×10^{-6}), (0.1016, 2.31×10^{-5}), (0.0585, 1.08×10^{-5}), (0.2239, 8.93×10^{-5}).*

latter becomes impractical and leads to quantization artifacts and shifts in colors. On the other hand, the expressiveness and robustness of our method is assessed through estimating and manipulating biophysical parameters of a variety of skin shades covering the Fitzpatrick scale.

As future work, 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 weird 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.

Another direction to further improve our results would be to train the neural network to learn the mapping between skin parameters and spectral albedo directly, instead of RGB. Finally, we want to look at adding the ability to automatically remove or compensate the baked lighting that exists in some captures, through analyzing the parts of the image that show incoherence in the estimated skin parameters, with the goal of reconstructing the skin properties under uncontrolled lighting scenarios like normal pictures.

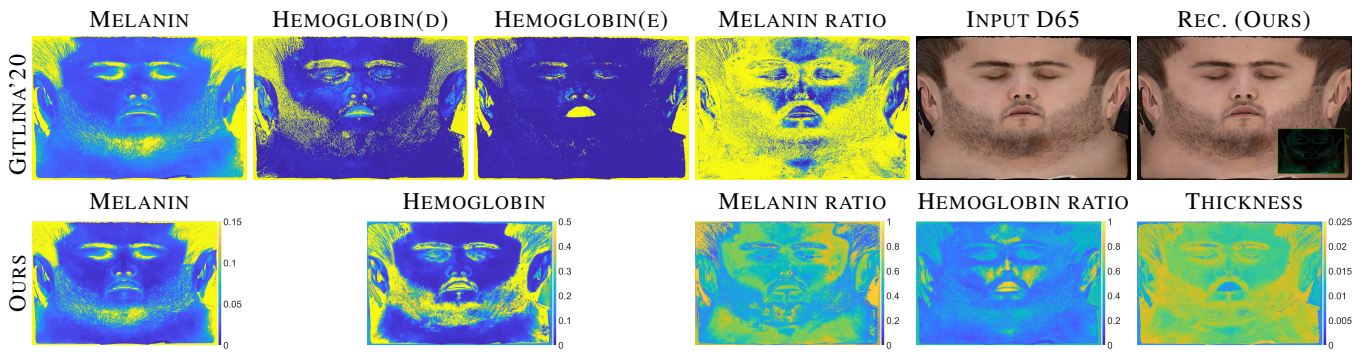


Figure 6: Albedo map from [GGD*20], comparison of their estimated parameters (top row) and ours (bottom row). On top right, we converted their albedo map from D65' illuminant to standard D65, and removed some baked lighting (see Supplemental, Section 1). Our model is able to reconstruct the albedo with $MSE = 0.1608$ (abs. error in the inset). $MSE = 0.0687$ when a mask over hair regions is applied as in Figure 7. While the common parameters with related work are similar, our new parameters predict well features like the reddish areas due to blood oxygenation, removing the need of extra blood in the epidermis. Or the epidermal thickness, where lips are identified as the thinnest area. The extra parameters available in our model might also explain the smoother melanin type ratio compared to [GGD*20], where it is the only extra parameter their model heavily relies on to enable a closer matching to the original albedo than previous works.



Figure 7: From top to bottom, captured albedos and reconstructions our Neural and LUT approaches for five subjects of different skin type ($\times 4$ absolute errors as insets). The learned mapping exhibits slightly higher reconstruction error than the discrete LUT approach. This was expected due to the inductive bias of the network, forcing smoothness in the recovered maps and thus enabling robust manipulation of the estimated properties.




Figure 8: Estimated skin parameters. *The overall melanin concentration is coherent to the skin type of each subject. Note how melanin spots are nicely isolated from small capillars, veins or reddish imperfections. Also, the color of the lips is mostly due to the high concentration of blood and the relatively thin epidermis in all cases.*

Acknowledgments

We thank Olivier Maury and Akira Orikasa for general lookdev support, Christoph Lassner for conversations about the neural model, and Tony Tung, Ronald Mallet and Yaser Sheikh for their face captures. Scans and actors for Subjects A and F were provided by Eisko [Eis21].

References

- [AMD01] ANGELOPOULO, E, MOLANA, RANA, and DANILIDIS, KOSTAS. “Multispectral skin color modeling”. *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*. Vol. 2. IEEE, 2001, II-II 2, 3.
- [AP81] ANDERSON, R ROX and PARRISH, JOHN A. “The optics of human skin”. *Journal of investigative dermatology* 77.1 (1981), 13–19 3.
- [AS17] ALOTAIBI, SARAH and SMITH, WILLIAM AP. “A biophysical 3D morphable model of face appearance”. *Proceedings of the IEEE International Conference on Computer Vision Workshops*. 2017, 824–832 2, 3.
- [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 2.
- [Bl20] BLENDER ONLINE COMMUNITY. “Blender: a 3D modelling and rendering package”. (2020). URL: <http://www.blender.org> 5.



Figure 9: Rendering results under our most neutral lighting scenario for edited skin parameters inferred by our model over different skin types (classified according to the Fitzpatrick scale). From left to right, original, followed by different manipulations per subject, three in blood concentration and three in melanin concentration (see similar renders under 2 additional lighting environments in Supplementary Figures 7 and 8). Note that these are straight algebraic edits on the recovered skin components, with no additional artistic tweaks or touch-ups involved. Though they demonstrate that our skin model describes an expressive latent space, these naive edits can result sometimes in semi non natural skins, and we suggest maybe more complex edits, some of them shown in Figure 10.



Figure 10: 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.

- [BRM*18] BITTERLI, BENEDIKT, RAVICHANDRAN, SRINATH, MÜLLER, THOMAS, WRENNINGE, MAGNUS, NOVÁK, JAN, MARSCHNER, STEVE, and JAROSZ, WOJCIECH. “A Radiative Transfer Framework for Non-Exponential Media”. *ACM Trans. Graph.* 37.6 (Dec. 2018). ISSN: 0730-0301 4.
- [CBKM15] CHEN, TENN F, BARANOSKI, GLADIMIR VG, KIMMEL, BRADLEY W, and MIRANDA, ERIK. “Hyperspectral modeling of skin appearance”. *ACM Transactions on Graphics (TOG)* 34.3 (2015), 1–14 2, 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.
- [dEo21] D’EON, EUGENE. “A Hitchhiker’s Guide to Multiple Scattering”. (2021) 4.
- [dEon19] D’EON, EUGENE. “A Reciprocal Formulation of Nonexponential Radiative Transfer. 2: Monte Carlo Estimation and Diffusion Approximation”. *Journal of Computational and Theoretical Transport* 48.6 (Sept. 2019), 201–262. ISSN: 2332-4325 4.
- [DHT*00] DEBEVEC, PAUL, HAWKINS, TIM, TCHOU, CHRIS, DUIKER, HAARM-PIETER, SAROKIN, WESTLEY, and SAGAR, MARK. “Acquiring the reflectance field of a human face”. *Proceedings of the 27th annual conference on Computer graphics and interactive techniques.* 2000, 145–156 2.
- [DJ06] DONNER, CRAIG and JENSEN, HENRIK WANN. “A Spectral BSSRDF for Shading Human Skin.” *Rendering techniques 2006* (2006), 409–418 2, 3, 11.
- [DWd*08] DONNER, CRAIG, WEYRICH, TIM, D’EON, EUGENE, RAMAMOORTHY, RAVI, and RUSINKIEWICZ, SZYMON. “A layered, heterogeneous reflectance model for acquiring and rendering human skin”. *ACM Transactions on Graphics (TOG)* 27.5 (2008), 1–12 2.
- [Eis21] EISKO. “Photoreal Digital Doubles”. (2021). URL: <https://www.eisko.com/8>.
- [Fit88] FITZPATRICK, THOMAS B. “The validity and practicality of sun-reactive skin types I through VI”. *Archives of dermatology* 124.6 (1988), 869–871 6.
- [FJM*20] FRISVAD, JEPPE REVAL, JENSEN, SØREN ALKÆRSIG, MADSEN, JONAS SKOVLUND, CORREIA, ANTÓNIO, YANG, LI, GREGERSEN, SØREN KS, MEURET, YOURI, and HANSEN, POUL-ERIK. “Survey of Models for Acquiring the Optical Properties of Translucent Materials”. *STAR* 39.2 (2020) 2.
- [GAS*19] GEVAUX, LOU, ADNET, CYPRIEN, SÉROUL, PIERRE, CLERC, RAPHAEL, TRÉMEAU, ALAIN, PERROT, JEAN LUC, and HÉBERT, MATHIEU. “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, BUSCH, JAY, YU, XUEMING, and DEBEVEC, PAUL. “Multi-view 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, HANSEN, JAN, LATTAS, ALEXANDROS, PAI, DINESH, and GHOSH, ABHIJEET. “Practical Measurement and Reconstruction of Spectral Skin Reflectance”. *Computer Graphics Forum.* Vol. 39. 4. Wiley Online Library. 2020, 75–89 2, 6, 7.
- [GGR*21] GEVAUX, LOU, GIERSCHENDORF, JORDAN, RENGOT, JULIETTE, CHEREL, MARIE, SÉROUL, PIERRE, NKENGNE, ALEX, ROBIC, JULIE, TRÉMEAU, ALAIN, and HÉBERT, MATHIEU. “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, FREDERIKSEN, SUNE, and DEBEVEC, PAUL. “Practical modeling and acquisition of layered facial reflectance”. *ACM SIGGRAPH Asia 2008 papers.* 2008, 1–10 2.
- [GRB*18] GOTARDO, PAULO, RIVIERE, JÉRÉMY, BRADLEY, DEREK, GHOSH, ABHIJEET, and BEELER, THABO. “Practical Dynamic Facial Appearance Modeling and Acquisition”. 37.6 (Dec. 2018) 2.
- [GRB*20] GOTARDO, PAULO, RIVIERE, JÉRÉMY, BRADLEY, DEREK, GHOSH, ABHIJEET, and BEELER, THABO. “Practical dynamic facial appearance modeling and acquisition”. (2020) 2.
- [GTB*13] GRAHAM, PAUL, TUNWATTANAPONG, BOROM, BUSCH, JAY, YU, XUEMING, JONES, ANDREW, DEBEVEC, PAUL, and GHOSH, ABHIJEET. “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 5.
- [HG41] HENYEV, LOUIS G and GREENSTEIN, JESSE L. “Diffuse radiation in the galaxy”. *The Astrophysical Journal* 93 (1941), 70–83 4.
- [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 (EUROGRAPHICS 2015)* 34.2 (2015) 2, 3.
- [INN07] IGARASHI, TAKANORI, NISHINO, KO, and NAYAR, SHREE K. *The appearance of human skin: A survey.* Now Publishers Inc, 2007 2.
- [Jac13] JACQUES, STEVEN L. “Optical properties of biological tissues: a review”. *Physics in Medicine & Biology* 58.11 (2013), R37 3, 11.
- [JH19] JAKOB, WENZEL and HANIKA, JOHANNES. “A Low-Dimensional Function Space for Efficient Spectral Upsampling”. *Computer Graphics Forum (Proceedings of Eurographics)* 38.2 (Mar. 2019) 5.
- [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 11.

- [JSB*10] JIMENEZ, JORGE, SCULLY, TIMOTHY, BARBOSA, NUNO, DONNER, CRAIG, ALVAREZ, XENXO, VIEIRA, TERESA, MATTS, PAUL, ORVALHO, VERÓNICA, GUTIERREZ, DIEGO, and WEYRICH, TIM. “A practical appearance model for dynamic facial color”. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 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 2, 3.
- [KB14] KINGMA, DIEDERIK P and BA, JIMMY. “Adam: A method for stochastic optimization”. *arXiv preprint arXiv:1412.6980* (2014) 5.
- [KM31] KUBELKA, PAUL and MUNK, FRANZ. “An article on optics of paint layers”. *Z. Tech. Phys* 12.593-601 (1931), 259–274 2.
- [KRP*15] KLEHM, OLIVER, ROUSSELLE, FABRICE, PAPAS, MARIOS, BRADLEY, DEREK, HERY, CHRISTOPHE, BICKEL, BERND, JAROSZ, WOJCIECH, and BEELER, THABO. “Recent Advances in Facial Appearance Capture”. *Computer Graphics Forum (Proceedings of Eurographics - State of the Art Reports)* 34.2 (May 2015), 709–733 2.
- [LWC12] LISTER, TOM, WRIGHT, PHILIP A., and CHAPPELL, PAUL H. “Optical properties of human skin”. *Journal of Biomedical Optics* 17.9 (2012), 1–15 4.
- [MHP*07] MA, WAN-CHUN, HAWKINS, TIM, PEERS, PIETER, CHABERT, CHARLES-FELIX, WEISS, MALTE, and DEBEVEC, PAUL. “Rapid Acquisition of Specular and Diffuse Normal Maps from Polarized Spherical Gradient Illumination”. 2007 2.
- [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 3, 4.
- [MWL*99] MARSCHNER, STEPHEN R, WESTIN, STEPHEN H, LAFORTUNE, ERIC PF, TORRANCE, KENNETH E, and GREENBERG, DONALD P. “Image-based BRDF measurement including human skin”. *Rendering Techniques '99*. Springer, 1999, 131–144 2.
- [MY19] MALLETT, IAN and YUKSEL, CEM. “Spectral Primary Decomposition for Rendering with RGB Reflectance”. *Eurographics Symposium on Rendering (EGSR 2019)*. Strasbourg, France: The Eurographics Association, 2019 5.
- [NMM*19] NUNES, AUGUSTO LP, MACIEL, ANDERSON, MEYER, GARY W, JOHN, NIGEL W, BARANOSKI, GLADIMIR VG, and WALTER, MARCELO. “Appearance modelling of living human tissues”. *Computer Graphics Forum*. Vol. 38. 6. Wiley Online Library. 2019, 43–65 2.
- [PJ] PRAHL, SCOTT and JACQUES, STEVEN L. *Tabulated Molar Extinction Coefficient for Hemoglobin in Water*. <https://omlc.org/spectra/hemoglobin/summary.html> 11.
- [RGB*20] RIVIERE, JÉRÉMY, GOTARDO, PAULO, BRADLEY, DEREK, GHOSH, ABHIJEET, and BEELER, THABO. “Single-shot high-quality facial geometry and skin appearance capture”. *ACM Transactions on Graphics (TOG)* 39.4 (2020), 81–1 2.
- [Sai92] SAIDI, IYAD SALAM. “Transcutaneous optical measurement of hyperbilirubinemia in neonates”. PhD thesis. 1992 3.
- [Sau42] SAUNDERSON, JL. “Calculation of the color of pigmented plastics”. *JOSA* 32.12 (1942), 727–736 2.
- [SS06] SARNA, TADEUSZ and SWARTZ, HAROLD A. “The physical properties of melanins”. *The pigmentary system: physiology and pathophysiology* (2006), 311–341 11.
- [THM99] TSUMURA, NORIMICHI, HANEISHI, HIDEAKI, and MIYAKE, YOICHI. “Independent-component analysis of skin color image”. *JOSA A* 16.9 (1999), 2169–2176 2.
- [TOS*03] TSUMURA, NORIMICHI, OJIMA, NOBUTOSHI, SATO, KAYOKO, SHIRAIISHI, MITSUHIRO, SHIMIZU, HIDETO, NABESHIMA, HIROHIDE, AKAZAKI, SYUUICHI, HORI, KIMIHIKO, and MIYAKE, YOICHI. “Image-based skin color and texture analysis/synthesis by extracting hemoglobin and melanin information in the skin”. *ACM SIGGRAPH 2003 Papers*. 2003, 770–779 2.
- [VGI94] VRHEL, MICHAEL J., GERSHON, RON, and IWAN, LAWRENCE S. “Measurement and Analysis of Object Reflectance Spectra”. *Color Research & Application* 19.1 (1994), 4–9 3.
- [VJSS89] VAN GEMERT, MJC, JACQUES, STEVEN L, STERENBORG, HJCM, and STAR, WM. “Skin optics”. *IEEE Transactions on biomedical engineering* 36.12 (1989), 1146–1154 3.
- [WLL*09] WEYRICH, TIM, LAWRENCE, JASON, LENSCH, HENDRIK PA, RUSINKIEWICZ, SZYMON, and ZICKLER, TODD. *Principles of appearance acquisition and representation*. Now Publishers Inc, 2009 2.
- [WMLT07] WALTER, BRUCE, MARSCHNER, STEPHEN R., LI, HONGSONG, and TORRANCE, KENNETH E. “Microfacet Models for Refraction through Rough Surfaces”. *Proceedings of the 18th Eurographics Conference on Rendering Techniques*. EGSR'07. Grenoble, France: Eurographics Association, 2007, 195–206. ISBN: 9783905673524 5.
- [WMP*06] WEYRICH, TIM, MATUSIK, WOJCIECH, PFISTER, HANSPETER, BICKEL, BERND, DONNER, CRAIG, TU, CHIEN, MCANDLESS, JANET, LEE, JINHO, NGAN, ADDY, JENSEN, HENRIK WANN, and GROSS, MARKUS. “Analysis of Human Faces Using a Measurement-Based Skin Reflectance Model”. *ACM Trans. Graph.* 25.3 (July 2006) 2.
- [WVH17] WRENNINGE, MAGNUS, VILLEMEN, RYUSUKE, and HERY, CHRISTOPHE. *Path traced subsurface scattering using anisotropic phase functions and non-exponential free flights*. Tech. rep. Technical Memo, 2017 4, 5.
- [ZDP*19] ZHEREBTSOV, EVGENY, DREMIN, VIKTOR, POPOV, ALEXEY, DORONIN, ALEXANDER, KURAKINA, DARIA, KIRILLIN, MIKHAIL, MEGLINSKI, IGOR, and BYKOV, ALEXANDER. “Hyperspectral imaging of human skin aided by artificial neural networks”. *Biomedical optics express* 10.7 (2019), 3545–3559 2–4.
- [ZSR*21] ZHANG, JINGANG, SU, RUNMU, REN, WENQI, FU, QIANG, and NIE, YUNFENG. “Learnable Reconstruction Methods from RGB Images to Hyperspectral Imaging: A Survey”. *arXiv:2106.15944* (2021) 4.

Appendix A: Specifications of Skin Chromophores

Parameter	Description	Value
$\mu_{a_{hbO2}}$	Oxy-Haemoglobin absorption	2.303 $\frac{p_{hb}\epsilon_{hbO2}}{w_{hb}}$
$\mu_{a_{hb}}$	Deoxy-Haemoglobin absorption	2.303 $\frac{p_{hb}\epsilon_{hb}}{w_{hb}}$
ϵ_{hbO2}	Oxy-Haemoglobin Extinction	[Jac13]
ϵ_{hb}	Deoxy-Haemoglobin Extinction	[Jac13]
p_{hb}	Haemoglobin Concentration	150
w_{hb}	Molar weight of Haemoglobin	64500
$\mu_{a_{eu}}$	Eumelanin absorption	$6.6x10^{11} \lambda^{-3.33}$
$\mu_{a_{pheo}}$	Pheomelanin absorption	$2.9x10^{15} \lambda^{-4.75}$
p_{bil}	Bilirubin Concentration	0.05
w_{bil}	Molar weight of bilirubin	584.66
$p_{\beta-c_e}$	β -carotene Concentration (Epidermis)	$2.1x10^{-4}$
$p_{\beta-c_d}$	β -carotene Concentration (Dermis)	$7x10^{-5}$
$w_{\beta-c}$	Molar weight of β -carotene	536.8726

Table 2: Chromophore specifications. The absorption coefficient is defined in cm^{-1} ; the extinction coefficient ϵ_c in $\frac{cm^{-1}}{moles/liter}$; 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 $\mu_{a_{pheo}}$ 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 haemoglobin extinction can be found tabulated in [PJ].