A Biophysically-based Skin Model for Heterogeneous Volume Rendering

Qi Wang
Zhejiang University
Hangzhou 310058, China
wqnina1995@gmail.com

Fujun Luan
Adobe Research
San Francisco, CA 94103, USA
fluan@adobe.com

Yuxin Dai
Zhejiang A & F University
Hangzhou 311300, China
buttersdousha@163.com

Yuchi Huo
Zhejiang Lab; Zhejiang University
Hangzhou 311121, China; Hangzhou 310058, China
HUOYC@zhejianglab.com

Hujun Bao
Zhejiang University
Hangzhou 310058, China
bao@cad.zju.edu.cn

Rui Wang
Zhejiang University
Hangzhou 310058, China
ruiwang@zju.edu.cn

Abstract

Realistic human skin rendering has been a long-standing challenge in computer graphics, especially biological-based skin rendering has received more attention in recent years, since it provides a more realistic skin rendering appearance and a more intuitive way to adjust skin style. In this work, we present a novel heterogeneous biophysically-based volume rendering method for human skin that improves the realism of skin appearance while easily simulating various types of skin effects even skin diseases by modifying the biological coefficients textures. Specifically, we introduce a two-layer skin representation by mesh deformation that explicitly models the epidermis and dermis with heterogeneous volumetric medium layers that contain the corresponding spatially-varying melanin and hemoglobin, respectively. Further, to better facilitate skin acquisition, we introduce a learning-based framework that automatically estimates spatially-varying biological coefficients from an albedo texture, enabling biophysically-based and intuitive editing such as tanning, pathological vitiligo, freckle, etc. We illustrate the effects of multiple skin editing applications, and demonstrate superior quality to commonly used random walk skin rendering method with more convincing skin details regarding subsurface scattering.

Keywords: ***skin model, rendering, ray tracing, biophysics***

1. Introduction

Human skin typically exhibits fine-level, complex and translucent appearance with an extreme amount of details due to its biophysical nature and structural variations across the local bounding surface and volumetric, layered medium underneath it. A realistic biophysically skin model is thus required to faithfully reproduce these color, freckles, pimples, wrinkles, veins, pores, and scars with physically-based light transport simulation, as human observers are sensitive to such appearance details of human skin and (in particular) faces. In general, such skin tone and skin details are spatially-varying and heterogeneous, imposing difficulties in accurately modeling, rendering, and acquisition of skin appearance. As a result, it remains an important topic and long-standing challenge to generate photorealistic and convincing skin rendering in computer graphics.

Classical BSSRDF-based methods mainly focus on efficiently simulating the translucency of subsurface scattering either through extensive precomputation or approximation. Traditional skin models for one-layer BSSRDF-based shading [18, 14] leveraging diffusion approximation often tend to yield suboptimal quality (e.g., blurring of high-frequency details) due to its oversimplified assumptions. Multi-layered skin models [7, 9, 6] for subsurface scattering light transport simulation, on the other hand, are generally more reasonable to represent skin structure, but also suffers from the same oversimplified assumptions problem as single layer model. Different from these diffusion-based methods, random walk method[38] is now more commonly used in the film industry, which is a more directly simulated, physically correct, ray tracing approach to render skin ap-

pearance. However, the essence of this method is to use homogeneous volume path tracing to simulate the skin color, which lacks biological meaning, making it impossible to easily edit the skin appearance from a biological point of view.

To address the aforementioned challenges, we tackle this problem from the biophysical perspective and propose a novel biophysically-based, heterogeneous human skin model readily deployable in a GPU-based volumetric path tracer. Unlike previous layered skin models that approximate the translucent effects from diffusion-based BSSRDF or random walk methods that abandon the skin model and directly fit the color values, we explicitly model the two layers containing heterogeneous medium derived from biological mechanisms, representing spatially-varying melanin and hemoglobin in the corresponding epidermis and dermis layers, respectively. This enables intuitive and biological manipulation of the skin appearance, such as oxygenated blood, tanning, flushing, and simulated pathological vitiligo while maintaining the realism of the skin appearance. We further propose a learning-based framework to estimate these biophysical skin properties from a single captured albedo map, enabling automatic capturing of human skin appearance to facilitate acquisition.

Concretely, our high-level contributions include:

- A biophysically-based and heterogeneous human skin model with layers that explicitly represents spatially-varying melanin and hemoglobin distributions in medium with textures for realistic volumetric skin rendering.
- A learning-based framework that automatically estimates the biophysical parameters from a single captured albedo texture.
- Demonstrations like skin appearance editing, skin aging, and faithful tattoo, verifying the potential usage of the biophysically-based skin model.
- GPU-accelerated open source codes for biophysically-based skin rendering.

We take a step forward to the more convincing and realistic modeling of human skin from the biophysical viewpoint. Through a number of experiments and ablation studies, we demonstrate that our approach is robust to reproducing skin appearance details with high fidelity, as well as faithful biophysically-based editing of skin appearance. We compare our method with random walk method which demonstrate the superiority of heterogeneous skin modelling. Our code and data will be made publicly available.

2. Related Work

Human skin shading has always attracted much attention in computer graphics. Here, we focus on prior biology-based skin rendering models and methods that are most relevant to our proposed method. We refer to the book on skin shading models and rendering [2] for a more comprehensive survey.

2.1. Biophysically-based Models

Donner and Jensen [7] proposed a diffusion-based spectral BSSRDF for shading skin leveraging biological coefficients, which can physiologically control skin color. They decomposed the skin into two layers: epidermis and dermis, and also analyzed the concentration of melanin and hemoglobin in the corresponding layer. Subsequently, Donner *et al*. [9] proposed to use three 2D textures to represent spatially-varying biological coefficients: a melanin texture, a hemoglobin texture, and an inter-layer absorption texture. Different from Donner and Jensen [7], they also consider small amounts of β-carotene and hemoglobin in the epidermis, which makes their model closer to the real skin structure. However, the infinitesimally thin absorbing layer between the epidermis and dermis is not biologically correct, and additional physical equipment is required to capture these biological parameters. Gitlina *et al*. [12] also present a method related to skin rendering, while their work mainly focuses on spectral skin reflectance measurement and reconstruction. Their shading model refers to Donner and Jensen [7]. Similar to our work, Krishnaswamy and Baranoski [2] proposed a biophysically-based spectral model for skin, and rendered it with a volume path tracer. Their BioSpec model accounts for all components of light propagation in skin tissues and represents skin as a five-layer model, which guarantees the biological correctness of the model. Nevertheless, compared to our method, they only considered the homogeneous layers and cannot generate multicolored results. Further, its complex skin structure and overcomplicated biological coefficients are less artist-friendly for parameter manipulation or appearance editing. Jimenez *et al*. [20] extended static skin rendering to animation. They utilized specific hardware to obtain the melanin and hemoglobin content of the skin, and analyzed the changes in the distribution of these two pigments under different expressions. Similar to Jimenez, Iglesias *et al*. [15] analyzed changes in biological and optical parameters of skin at different ages and genders to simulate skin appearance of different ages. These two methods cannot recover biological parameters from a single albedo, also rendering with diffusion-based BSSRDF models, which is less practical and suboptimal quality. Recently, Aliaga *et al*. [1] proposed a method to recover biophysical skin properties from captured RGB albedo, which is similar in part to our method. However, they require captured high-quality albedo as input while our approach accept low-quality hand-drawn albedo. At the same time, they did not leverage the generated biophysical parameter textures for rendering, but

restore them to a new albedo and using Blender's random walk algorithm as the final renderer. This rendering method is not integrated with the skin biophysical model. While our approach explicitly models the biophysical layers of the skin and utilizing heterogeneous volume path tracing for rendering. Meanwhile, our method explicitly changes the biological parameters instead of restoring the albedo image through the neural network, which circumventing the robustness problem of the neural network.

2.2. Rendering methods

No matter what kind of skin biophysical model, a rendering method is the basis to generate the final result. The most classic ones are diffusion-based [32] methods, which approximate the solution for light transport equation in highly scattering medium. These methods account for subsurface scattering not only in skin but also in many kinds of translucent materials such as wax, marble, and so on. These diffusion-based methods are general methods so we briefly introduced here. A big breakthrough was introduced by Jensen *et al*. [18] who first use the diffusion-based dipole model to represent highly scattering medium. They improved the performance of dipole BSSRDF later in [17]. Donner and Jensen [6] further introduced a multi-layer model for simulating light scattering using multiple dipoles. Their method works in the frequency domain, which simplifies the convolutions to a geometric series that speed up the rendering. This method is used by most biophysical skin models. Donner and Jensen [8] then proposed a method that combines photo mapping [31] with a diffusion model that can simulate the effect of internal occluding objects in transparent materials. In order to break the assumption that the material surface is flat and semi-infinite, they introduced a new quadpole diffusion approximation that models diffusion at right-angled edges. Classic diffusion may cause errors when simulating the participating medium, and later BSSRDF-based models improved upon it via modified diffusion solutions [14]. Vicini *et al*. [35] presented a shape-adaptive subsurface scattering model using neural networks, enabling more realistic subsurface scattering compared to diffusion-based methods. Their method generates a sample on the surface from a reference distribution produced by a volumetric path tracer, along with the fraction of incident light that eventually leaves the surface. For diffusion-based methods, the flat and semi-infinite medium assumption leads to errors when rendering curved objects, such as the ears and nose for skin rendering. In contrast, our volumetric approach does not suffer from these limitations and is capable of producing more physically plausible results. Despite Vicini *et al*. [35] also lifted the BSSRDF assumptions, the quality of their results heavily depends on network predictions and the accuracy of local geometry approximated using a low-order polynomial, which can be less robust to complex geometric models.

To overcome the shortcomings of diffusion-based BSSRDF, the brute-force random walk (path traced subsurface scattering) methods [25, 38] are proposed. These two methods relax basically all assumptions and obtain more accurate results than the classic diffusion-based methods. For polychromatic rendering, random walk methods restore the corresponding homogeneous parameter of each color for homogeneous volume path tracing, which simulates a variety of colors. However, differs from heterogeneous volume path tracing, these methods does not take into account the influence between adjacent colors.

There are also some real-time BSSRDF-based models [22, 21, 19, 3, 13] that we will not discuss in details here, as they typically use simplified models for performance at the expense of downgraded quality and are less relevant to our approach.

3. Method

3.1. Biophysically-based Human Skin Model

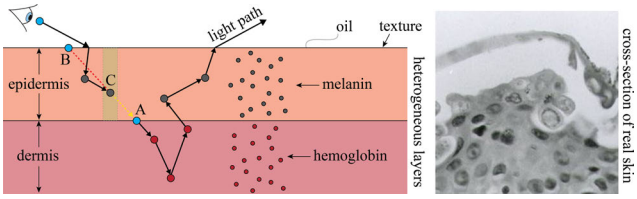

Figure 1. Illustration of our biophysical skin model and texture coordinate value calculation. Our model contains two heterogeneous skin layers, namely the epidermis layer (w/ melanin) and the dermis layer (w/ hemoglobin), both containing spatially-varying medium parameters, which produce changeable and complex skin color. When computing texture coordinate value for a specific particle C (melanin or hemoglobin) during volume path tracing, we sample a ray direction and intersect the layer boundary ($C \rightarrow A$), then back trace the ray to intersect the other layer boundary ($C \rightarrow B$). Finally, we use the texture coordinate values of point A, B and the segment length $|AC|$, $|BC|$ to interpolate the final texture coordinate value. On the right, we show a cross-section of real skin (source from Wiki Commons [5]). Note the distribution heterogeneity of melanin and hemoglobin within the skin layers.

Realistic rendering of human skin is a challenging task due to its subtle, translucent appearance with fine-level details from microgeometry and complex biological properties. Light paths in human skin are subsurface scattered as in medium with a high optical depth, sometimes even with hundreds of bounces. Numerous biological and medical studies have analyzed light interaction within human skin structure, outlining the biological characteristics of its main components and how these components affect the scattering and absorption of light. In this section, we first present our

biophysically-based human skin model that is built on top of previous biological research. Then we introduce a learning-based inference framework for biological skin properties leveraging captured albedo texture (Section 3.2). We show a summary of notations in Table 1.

Table 1. Summary of common notations used in the paper.

Symbol	Description
λ	Wavelength of light
σ_a^{eu}	Absorption coefficient of eumelanin
σ_a^{ph}	Absorption coefficient of pheomelanin
σ_a^{other}	Absorption coefficient of other physiological tissue
σ_a^{epi}	Absorption coefficient of epidermis
σ_a^{der}	Absorption coefficient of dermis
σ_s	Scattering coefficients of epidermis
υ	Volume fraction of the epidermis occupied by melanin
α	The proportion of eumelanin in melanin
τ	Volume fraction of the dermis occupied by hemoglobin
β	The proportion of oxygenated hemoglobin in hemoglobin

It is a consensus in the biological community that human skin is a layered structure with various small-scale elements, e.g., pigments and collagen. Although most biological and medical studies indicate that human skin model has at least five layers, in practice we generally consider only two layers that contribute most significantly to the skin color and scattering properties[11] (i.e., dermis and epidermis). To represent the scattering and absorption properties in human skin, we follow prior work and also build a two-layer skin model with epidermis and dermis (see Fig. 1 for illustration). However, unlike previous BSSRDF-based methods, we explicitly model the skin layers that contain heterogeneous participating medium, and derive their scattering and absorption coefficients from spatially-varying melanin and hemoglobin for Monte Carlo volume path tracing, which as we demonstrate later is capable of achieving physically-accurate light interaction and convincing skin appearance visually under various illumination. We will describe in detail how to formulate from biophysical properties the optical volume parameters of each skin layer for absorption and reduced scattering coefficients σ_a and σ_s' , respectively. This enables us to edit skin appearance in a more intuitive and biophysically-based manner, such as tanning and pathological control. Because the absorption properties of skin need to be considered separately in the two layers, we discuss them one by one in the following paragraphs.

Firstly, we consider absorption in the epidermis. Melanin is most responsible for absorption in the epidermis layer, and it is a mixture of two kinds of pigments, the red-yellow pheomelanin and the brown-black eumelanin which both control the intensity of skin color. In general, the absorption coefficients of eumelanin and pheomelanin are higher for shorter wavelengths and can be approximated by simple power functions (functions from [29]):

$$\sigma_a^{\text{eu}}(\lambda) = 6.6 \times 10^{10} \times \lambda^{-3.33}\ mm^{-1}, \tag{1}$$

$$\sigma_a^{\text{ph}}(\lambda) = 2.9 \times 10^{14} \times \lambda^{-4.75}\ mm^{-1}, \tag{2}$$

where σ_a^{eu}, σ_a^{ph} are absorption coefficients of eumelanin and pheomelanin, respectively. λ denote the wavelength in nanometers. The absorption spectra of other components in the epidermis can be approximated by (function from [34]):

$$\sigma_a^{\text{other}}(\lambda) = 0.0244 + 8.53e^{-(\lambda-154)/66.2}\ mm^{-1}. \tag{3}$$

Assuming that melanin and other tissues are uniformly distributed in the epidermis, the total absorption coefficient is given by:

$$\begin{aligned}\sigma_a^{\text{epi}}(\lambda) = \upsilon(\alpha\sigma_a^{\text{eu}}(\lambda) + (1-\alpha)\sigma_a^{\text{ph}}(\lambda)) \\ + (1-\upsilon)\sigma_a^{\text{other}}(\lambda),\end{aligned} \tag{4}$$

where υ is the volume fraction of the epidermis occupied by melanin that varies from 0.013 to 0.43. α is the proportion of eumelanin in melanin, which varies from individual to individual. Both eumelanin and pheomelanin are found in most skin types, and some types are more likely to contain higher amounts of eumelanin. Different studies report different values of α, and no study has measured all kinds of skins. So the value range of α is not strictly specified. In this work, we set the value range to be within 0 and 1.

Secondly, we discuss the absorption in the dermis. Hemoglobin has the largest contribution to the absorption coefficient, which is also made up of two types of hemoglobin — oxygenated hemoglobin and deoxygenated hemoglobin. These hemoglobins have slightly different absorption spectra. These absorption spectra cannot be fitted with a simple equation, so we look up the molar extinction coefficient for the corresponding wavelength directly from [27]. Note that the conversion of molar extinction coefficient to absorption coefficient is given by:

$$\sigma_a(\lambda) = \frac{(2.303)(\mathrm{e}(\lambda))(\mathrm{x})}{64,500}, \tag{5}$$

where $\sigma_a(\lambda)$ is the absorption coefficients corresponding to wavelength, $\mathrm{e}(\lambda)$ is the oxygenated hemoglobin or deoxygenated hemoglobin molar extinction coefficient for the wavelength of interest, x is the concentration of hemoglobin with unit g/liter, hemoglobin has a normal concentration of 150g/liter of blood, 64500 is the gram molecular weight of hemoglobin.

Although bilirubin concentration in a normal skin is negligible, we also take into account it in the dermis. We take the molar extinction coefficient of bilirubin for the corresponding wavelength directly from [10]. Similarly, we assume that hemoglobin and other tissues are uniformly distributed in the dermis and the absorption spectra of other components in the dermis is the same as the epidermis.

Consequently, the total absorption coefficient is given

by:

$$\sigma_a^{\text{der}}(\lambda) = \tau(\beta\sigma_a^{\text{oxy}}(\lambda) + (1-\beta)\sigma_a^{\text{deoxy}}(\lambda) + \sigma_a^{\text{bili}}) + (1-\tau)\sigma_a^{\text{other}}(\lambda), \quad (6)$$

where τ is the volume fraction of hemoglobin in the dermis roughly varies from 0.002 to 0.07. Nevertheless, this value range changes due to location and skin type, and τ is not strictly specified. β is the proportion of oxygenated hemoglobin in hemoglobin. β is 0.9-0.95 in arteries and more than 0.47 in veins. However, arteries and veins distribute distinctly in different skin locations, so we fix β to be 0.75 according to [7].

Table 2. Our spectral wavelengths (nm) and biological absorption and scattering coefficients (mm^{-1}) corresponding to R,G,B channels.

Description	R	G	B
Wavelength	700	546.1	435.8
Eumelanin Absorption	22.150	50.632	107.330
pheomelanin Absorption	8.875	28.864	84.291
Oxy-hemoglobin Absorption	0.1553	26.704	71.123
Deoxy-hemoglobin Absorption	0.9608	27.453	292.932
Bilirubin Absorption	0.00026	0.00017	0.1268
Other Absorption	0.02663	0.0472	0.1452
Epidermis Scattering	4.6483	6.2014	8.0584
Dermis Scattering	2.9329	4.0421	5.4101

Thirdly, we consider the scattering coefficients in the epidermis and dermis. Mie scattering caused by cell and collagen fibrils happens in the epidermis. Furthermore, this forward scattering is also wavelength dependent. In the dermis, microstructures and smaller collagen fibers are responsible for Rayleigh scattering. Light can hardly travel through the dermis, so we assume the dermis is semi-infinitely thick. Mie and Rayleigh scattering can be approximated by a function from [16]:

$$\mu_s' = a\left(\frac{\lambda}{500(\text{ nm})}\right)^{-b}, \quad (7)$$

where a is 68.7 in epidermis and 45.3 in dermis and b is 1.161 in epidermis and 1.292 in dermis[30].The spectral wavelength we select and corresponding biological absorption and scattering coefficients are all shown in Table 2. Our RGB spectral wavelength refer to the CIE 1931 color space. Spectrum of different biological parameters are plotted in Fig. 2.

Finally, we consider the thickness of the epidermis, which has significant effects on skin color. The thickness of the epidermis is approximately 0.027-0.15mm. Thicker epidermis makes melanin more influential and hemoglobin less influential, which makes skin less ruddy. Epidermis thickness varies with location and age. We fix the thickness to 0.1mm when rendering heterogeneous layers (for most results) since it is a general value sufficient for most adult faces.

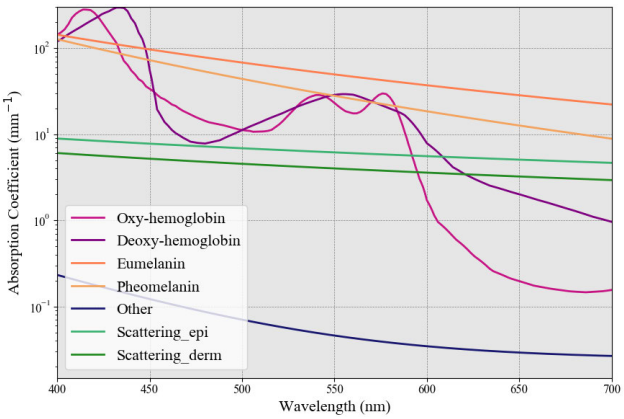

Figure 2. Spectral absorption and scattering coefficient plot of oxy-hemoglobin, deoxy-hemoglobin, eumelanin, pheomelanin, epidermis, dermis. As the order of magnitude of bilirubin absorption coefficient is too small, we do not illustrate it here.

3.2. Biophysical Skin Properties from Albedo

Utilizing Eqs. (4,6,7) and changing the thickness of the epidermis, two homogeneous layers that are capable of representing different skin types can be easily established. Combined with the volume rendering technique, we can render various skin colors. Nevertheless, spatially-variant skin details cannot be expressed by the assembly of two different homogeneous layers. Therefore, we propose a novel method to generate spatially-varying absorption and scattering coefficients of each layer leveraging the albedo texture, which can reproduce the complex variations of skin pigmentation.

The majority of changes in skin color, such as freckles and red rashes, are due to the proportion of melanin and hemoglobin in the epidermis and dermis. The main cause of freckles is the deposition of melanin in the skin epidermis because of endocrine disorders or external stimuli. That is, the proportion of melanin in specific areas of the epidermis increases. Red rash initiated by disease dilates capillaries that increase hemoglobin content in the dermis, which makes the skin area look reddish. The relationship of melanin and hemoglobin to the medium parameters can be calculated by the equations introduced before (Eqs. 4,6,7), so by acquiring the unique biological parameters at each point in the epidermis and dermis, the spatially-varying medium parameters is able to calculate and the color variation of the skin can be simulated biologically and physically correctly.

To obtain the spatially-varying biological coefficients textures, we first build a lookup table (LUT) using a two-layer homogeneous model with discrete biological parameters introduced above. The LUT contains discrete skin colors and corresponding biological coefficients. Using this LUT, we look up the corresponding biological coefficients

for each pixel in the albedo texture, thus obtaining three coefficient textures corresponding to the melanin fraction, hemoglobin fraction, and melanin type blend. Utilizing Eqs. (4,6,7), we generate two absorption textures of epidermis and dermis representing the medium absorption in these layers. Each channel of the absorption textures represents one of the three wavelengths of R,G and B. These two textures only contain absorption coefficients because scattering coefficients are fixed in each layer for every light wavelength.

After obtaining absorption textures, we can calculate the absorption coefficient of any point in the heterogeneous medium. We assume that each layer is a **local-homogeneous** medium, which means a tiny cylinder centered on a point on the skin is a homogeneous medium. For points in the skin, their texture coordinate value can be calculated by the equation:

$$UV_C = \left(\frac{l_{|AC|}}{l_{|AC|+|BC|}}\right) UV_A + \left(\frac{l_{|BC|}}{l_{|AC|+|BC|}}\right) UV_B, \quad (8)$$

where $l_{|AC|}$ is the distance from the current medium point to the intersection of the ray direction and the layer boundary. Similarly $l_{|BC|}$ is the distance to the inverse ray direction intersection of the layer's boundary. $l_{|AC|+|BC|}$ is the sum of $l_{|AC|}$ and $l_{|BC|}$. UV_A is the first intersection's texture coordinate value and UV_B is the second intersection's texture coordinate value (see Fig. 1). Eq. (8) represents a linear interpolation between these two texture coordinate values. These layers are relatively thin compared to the whole model and can be regarded as parallel planes locally, thus making our smooth interpolation of texture coordinate values sufficient. We can obtain medium parameters from the corresponding coefficient texture using this texture coordinate value. Thus a two-layer heterogeneous medium presented by two absorption textures can be rendered by a Monte Carlo volume path tracer.

Although the LUT-based method is able to estimate biological skin properties from albedo texture, the brute-force looking up process is relatively slow, and its results are not smooth because of the discrete nature of the undersampled LUT. Therefore, we then present a network structure PNNet combining a multi-layer perception (MLP) and a convolutional neural network (CNN) structure to regress albedo values to biological coefficients. We found that some works also use a similar approach for parameter mapping. Different from [12], all biological coefficients are predicted at once by our PNNet, we do not require multiple networks to predict biological parameters separately (see Fig. 3). The input to PNNet is an albedo texture patch, and the output is the biological coefficient texture patch. The detailed implementation will be introduced later. We show that our PNNet can predict accurate and smooth results, please refer to the Sec. 4 for details.

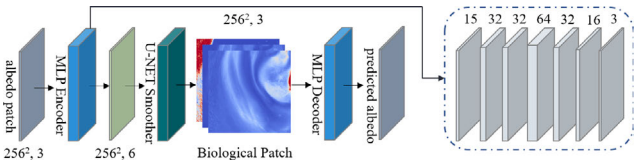

Figure 3. Illustration of our PNNet structure. Firstly, the albedo patches are passed as input to a 7-layer MLP network, then the outputs of MLP are concatenated with an albedo patch as input to the U-NET. Finally, the U-NET predicts the corresponding biological coefficients as texture maps. Note that we also train a decoder to revert the albedo patch from biological coefficient textures but it is not trained jointly. The structure of the decoder is the inverse of the encoder except that the output is 3 channels.

4. Implementation Details

Volumetric path tracing. We implement a GPU-based volume path tracer. We use three channels of R,G,B to render our result, which we found sufficient in reproducing most skin appearance, despite that sampling multiple light wavelengths can produce more accurate results [7]. An efficient wavelength-dependent estimator and a spectral tracking estimator [26] are both implemented in our renderer.

Next, we introduce our biological coefficients with typical values that calculating the medium parameters used by our volume path tracer. Table 3 summarizes four coefficients of our model. Note that the "t" coefficient is fixed to 0.001 when rendering heterogeneous layers. We use Blender to resize the geometry model along normal according to the thickness that guarantees an equidistant scaling. The gap between the outer and inner models represents the epidermis, while the inner model represents the dermis as we assume that the dermis is semi-infinitely thick. Note that models should be water tight. When rendering homogeneous layers, using Eqs.(4,6,7) directly gives the medium parameters of two models for specific thicknesses. When rendering heterogeneous layers, we use the LUT and PNNnet to obtain spatially-vary medium parameters that will be discussed later. Note that our typical values exceeds the range mentioned above, thus a wider range of colors can be represented, which enables our method to handle abnormal skin albedos.

Table 3. Typical values of coefficients to our two-layer model. We did not consider the hemoglobin content of the epidermis and $\beta-$carotene content which are generally less than 5% in regular skin because these pigments have little effect on skin color and greatly increase the complexity of our method. The oil content on skin surface is controlled by the specular roughness.

Parameter	Description	Typical range
υ	Melanin fraction	0 - 1
α	Melanin type blend	0 - 1
τ	Hemoglobin fraction	0 - 1
t	Thickness of epiderms	0.00027 - 0.0025

Lookup table. The settings of LUT are important for heterogeneous rendering since it will affect the medium parameter estimation. Our model for rendering LUT is a cube with a side length of twenty centimeters as the outer model, and we shrunk it along normal by 0.001 as the inner model since the thickness is fixed to 0.001 while rendering heterogeneous layers. After the model is determined, we then choose biological coefficients to obtain the medium parameters. We employ a relatively dense discretization for υ and τ (101 bins) coefficients which have a noticeable effect on skin color, while applying a coarse discretization for α (11 bins) to reduce the time to generate and look up the LUT since its contribution is subtle to the overall skin appearance. Next, we render the LUT using a classic volume path tracer with 1024 samples per pixel and a max depth of 95. Since skin is a highly-scattering material, setting the max depth too low will result in a loss of energy. We only focus on the result for the central part of the cube, so only the 2*mm* square patch centered at the cube is rendered. The average of pixel values in that patch is the final color corresponding to the cube coefficients.

We implement a CPU parallel look-up method to generate three biological coefficient textures and compute two absorption textures as a preprocessing operation before heterogeneous volume rendering; thus only texture coordinate calculation is needed during rendering. It takes approximately ten minutes to process a 2K resolution texture with a 112,211 data LUT using an Intel i7-7700 CPU.

PNNet training. We train a PNNet combining MLP and CNN to recover smooth biological coefficients. The MLP consists of 5 fully connected hidden layers and the CNN is a classic U-NET structure [28] with 4 upsampling/downsampling layers. Albedo patches first pass through MLP that is responsible for pixel-by-pixel prediction of biological coefficients to generate three biological coefficient vectors, which are reshaped into three textures and used as input to CNN. CNN is responsible for smoothing the results of MLP with the origin albedo texture and outputting the final biological parameter textures. We found that the predicted results using only MLP were not smooth enough, and more accurate results can be obtained by following the MLP with a CNN smoother. We also tried adding MLP decoder for joint training, but the results did not improve. Our training dataset contains 13462 pairs of patches generated by slicing albedo textures and corresponding biological textures using LUT. The resolution of each patch is 256*256. The resolution of albedo textures is generally high, so choosing a relatively large patch can contain more skin details.

We use the Adam optimizer [24], with a learning rate of 10^{-3} for MLP and 10^{-4} for CNN. The batch size is 128 of 100 epochs. Although more epochs result in better convergence, we find that 100 epochs are good enough to provide smooth and accurate predictions. The loss function is defined as:

$$\mathcal{L}(\mathbf{I}, \mathbf{I}') = \lambda_1 \mathcal{L}_2 + \lambda_2 \mathcal{L}_{\text{percept}}, \tag{9}$$

where $\mathcal{L}_2$ denotes the biological coefficients loss which computes the L2 loss between the predicted biological coefficients and the ground truth. $\mathcal{L}_{\text{percept}}$ denotes the VGG-16 perceptual loss [23]:

$$\mathcal{L}_{\text{percept}}(\mathbf{I}, \mathbf{I}') = \sum_{j=1}^{4} w_j^{\text{percept}} \|F_j(\mathbf{I}) - F_j(\mathbf{I}')\|_2^2, \tag{10}$$

where $F_1, ..., F_4$ are the feature maps corresponding to the output of layers of conv1_1, conv1_2, conv3_2 and conv4_2. We also experimented with a variety of different losses including SSIM loss, L1 loss, L2 loss, perceptual loss and their combinations, however we got the best results using L2 and perceptual loss.

To verify the accuracy of the generated biological textures, we also train a decoder to remap the biological textures back to albedo texture. The structure of the decoder is simple, which is the reverse of MLP. The training data is similar to the previous network, except that the biological coefficient textures are generated by the trained PNNet.

5. Experiments and Results

5.1. Ablation Study

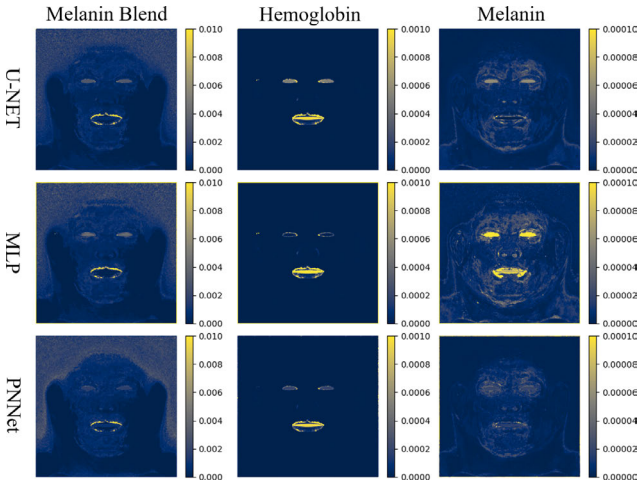

Figure 4. From top to bottom, we show error maps of biological coefficient textures belong to an Asian skin type albedo predicted by our CNN, MLP, and PNNet that combined both. Note that the error of PNNet predictions is smaller than the other two networks.

Here we compare the biological coefficient textures predicted by our three networks of an Asian skin type (see Fig. 4). Each row shows the error map against the corresponding LUT-based textures (ground truth). The PNNet result presents the smallest error comparing the other two alternatives, demonstrating the effectiveness of our network design choice. Table 4 shows the average PSNR of our evaluation

Table 4. Quantitative comparisons of our evaluation dataset for different network structures. The PSNR, SSIM and LPIPS value of PNNet is the best in all three biological coefficients.

Network	↑ PSNR			↑ SSIM[36]			↓ LPIPS[39]		
	Melanin Blend	**Hemoglobin**	**Melanin**	**Melanin Blend**	**Hemoglobin**	**Melanin**	**Melanin Blend**	**Hemoglobin**	**Melanin**
U-NET	28.509	40.292	45.133	0.720	0.886	0.968	0.365	0.252	0.189
MLP	32.010	46.599	57.316	0.725	0.913	0.986	0.280	0.139	0.112
PNNet	**32.685**	**49.215**	**58.176**	**0.786**	**0.926**	**0.988**	**0.250**	**0.096**	**0.095**

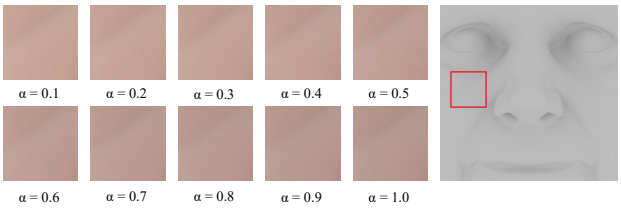

Figure 5. Change in skin color as the melanin type blend α increases from 0.1 to 1.0. The melanin fraction is fixed at 0.02, and the hemoglobin fraction is fixed at 0.025. Thickness t is 0.002. The higher the value of α, the darker the skin color. Note that α doesn't have that much effect on skin color. The rightmost picture illustrates the rendering area of the head model.

dataset. PNNet achieved the best results across biological coefficients consistently.

5.2. Results

In this section, we present our experiment results from three aspects: color representation capability of homogeneous layers with varying biological coefficients; network-generated biological coefficients textures; rendered spatially-varying skin color with heterogeneous layers.

Fig. 5 shows the effect of changing melanin type blend coefficient α from 0.1 to 1. The rightmost picture indicates the rendering area, same as in Fig. 6 and 7. High α values present a higher ratio of eumelanin which gives darker skin color. Certainly, the influence of α on skin color is more subtle than υ and τ, since absorption coefficients are not much different between eumelanin and pheomelanin, and the lightly pigmented skin type has low amounts of total melanin.

Figure 6. Change in skin color as the thickness t (in decimetre) increases from 0.001 to 0.005. The thin epidermis layer gives a reddish and pink appearance to the skin, while the thick epidermis makes the skin looks dark and brown. The melanin fraction υ is fixed at 0.025, the hemoglobin fraction is 0.02, and α is 0.5.

Then, in Fig. 6, we illustrate the effects of changing thickness t from 0.001 to 0.005 from left to right. Although the typical value of t ranges up to 0.0025, we break through this limit, to more clearly show the effect of thickness, such t values are not used in practical rendering. Low values of thickness t generate reddish skin color, which hemoglobin contributes significantly to skin color because more light energy enters the dermis. As the thickness of the epidermis increases, the light entering the dermis has less energy that melanin dominates skin color. In practice, the epidermis is not particularly thick, so the skin always shows reddish except for the African skin type, where most light energy is absorbed by melanin.

Fig. 7 shows results of changing melanin fraction υ and hemoglobin fraction τ. From left to right, υ increases from 0.005 to 0.4, and from top to bottom, τ increases from 0.001 to 0.1. The Caucasian skin type is close to the left part of the figure, while the result on the right is close to the African skin type, which indicates the effect of melanin on the skin color. The bottom part of the figure shows a more reddish color to demonstrate the role of hemoglobin. However, the redness of the skin is not obvious when the skin is darker, as we analyzed before. We demonstrate that even with a fixed epidermal thickness, our model can represent a wide variety of skin types.

In Fig. 8, we show biological textures predicted by our PNNet, taking albedo as input. These textures correspond to the content of biological coefficients, and these biological coefficients will be used to compute the spatially-varying absorption coefficients. Note that our PNNet only takes approximately 3.5 seconds to predict three 2K resolution biological coefficient textures with an NVIDIA Tesla-V100 GPU while LUT-based method takes approximately 10 minutes as mentioned above.

In Fig. 9, we compare textures of the decoder-generated albedo to the origin albedo. The biological coefficients generated by PNNet can be well mapped to the original albedo with the decoder, illustrating the accuracy of PNNet.

In Fig. 10, we show various parameter manipulation results to demonstrate the effectiveness of our explicit biological editing approach. Thanks to explicit modeling and heterogeneous volume rendering, the plotted results naturally represent changes in the corresponding biological parameters.

Fig. 11 shows images generated by heterogeneous volume rendering leveraging spatially-varying absorption coefficients under different environment maps. The specular reflection is modeled with Torrance-Sparrow microfacet

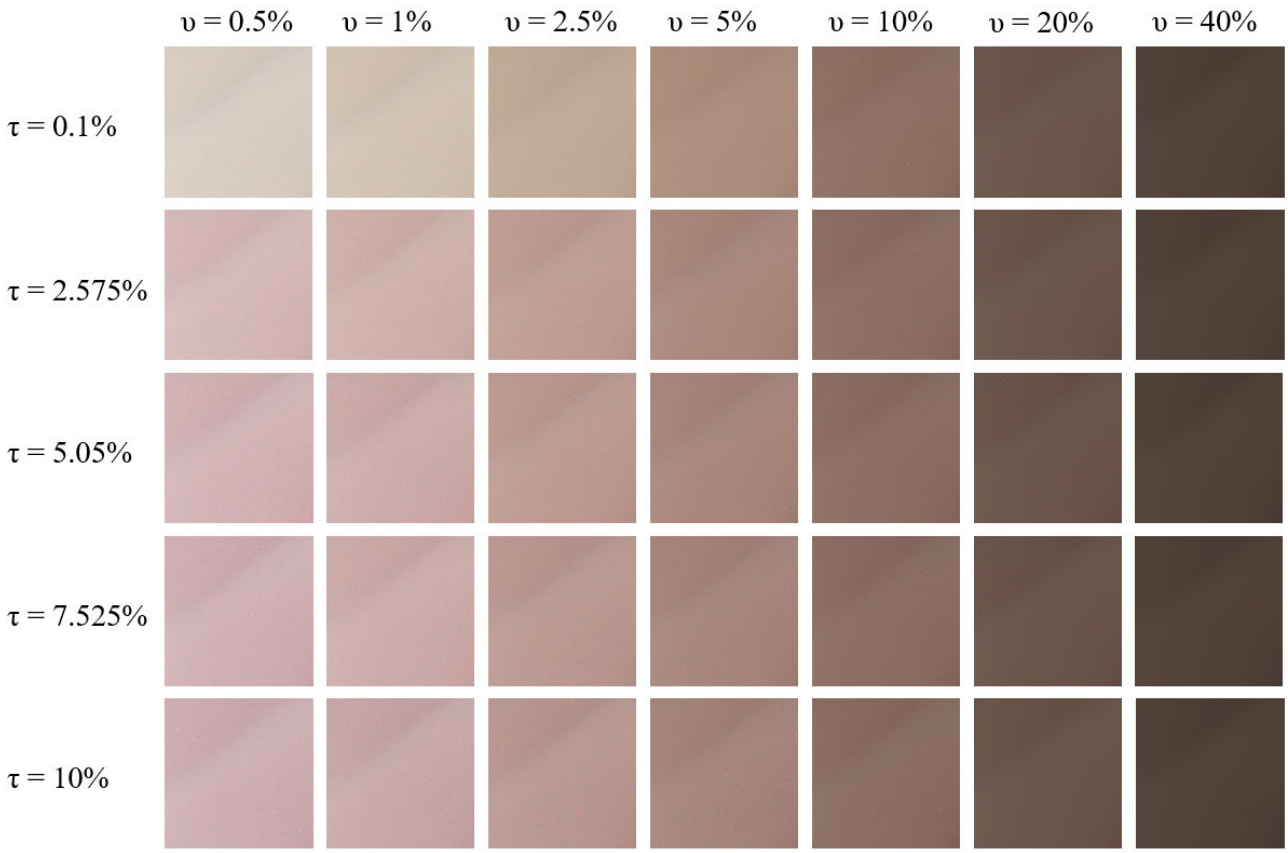

Figure 7. Changing in skin color as melanin fraction υ increases from 0.005 to 0.4 from left to right and hemoglobin fraction τ increases from 0.001 to 0.1 from top to bottom. The melanin type blend α is kept constant at 0.5. The thickness t is fixed at 0.002. Note that υ is not sampled uniformly as υ increases, the skin darkens rapidly. Also, we do not consider specular reflections here and show results with specular reflections later. As can be seen, our two-layer homogeneous model can represent a variety of kinds of skin.

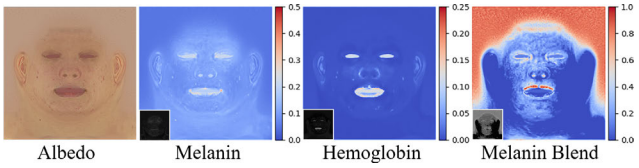

Figure 8. From left to right, we show the origin albedo, melanin fraction, hemoglobin fraction, and melanin type blend texture maps. Difference map $\times 4$ are shown in the bottom left corner of each predicted texture. PNNet predictions contain rich details, and the predicted values are basically in the typical range except for the lips, which physiologically contain more hemoglobin.

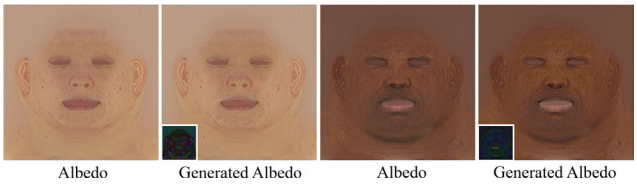

Figure 9. From left to right, we show the origin albedo texture and decoder-generated albedo texture of two skin types and the difference map $\times 4$ in the bottom left corner of generated albedos.

BRDF [33], which has been proved to be practical for fitting skin highlights [37]. Please zoom in to focus on skin details.

Apart from classic skin rendering, our method can be easily extended to render tattoo effects. In Fig. 12, we show the result of an Asian skin type hand model and tattoo editing on the arm, demonstrating the generality and flexibility of our model. Although tattoos cannot be represented by biological parameters directly, we can still simulate the changes in the dermis pigmentation caused by tattoos, which can be considered as a physiological change. Note that the effect of subsurface scattering around the tattoo is well reproduced. Another tattoo effect on the face is shown in Fig. 13.

In Fig. 14, we demonstrate that our biological model enables pathological editing with various disorder effects caused by skin disease. Note that, no artist involvement is required since this can be achieved by modifying biological

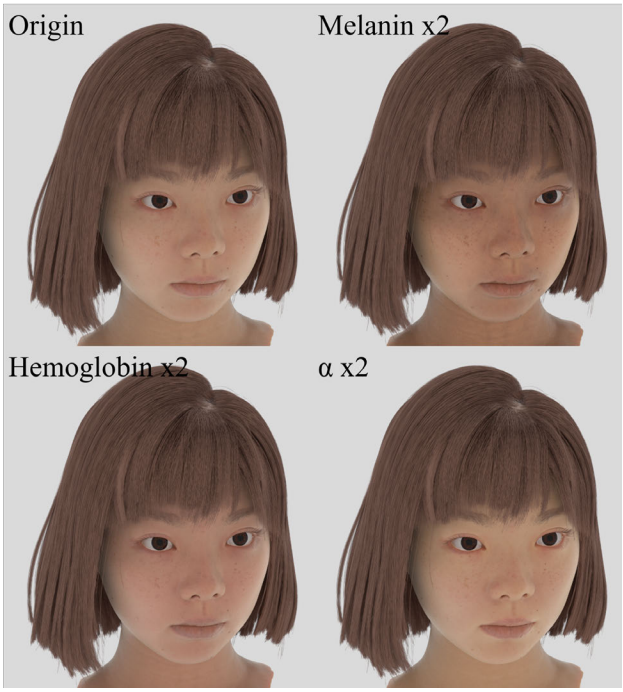

Figure 10. Biological parameter manipulation results under ambient light. Increased content of melanin or hemoglobin makes skin appear darker or redder respectively, while the melanin type blend α influence subtle to skin appearance because the spectra of eumelanin and pheomelanin are relatively similar.

Figure 11. Heterogeneous volume rendering results for Asian skin type. The corresponding environment map is in the insight of each result. The facial details of skin are well preserved with our model under various illumination conditions with faithful specular highlights and subsurface scattering. Zoom in is recommended to observe skin details.

coefficients with our method. Please zoom in to focus on pathological effect.

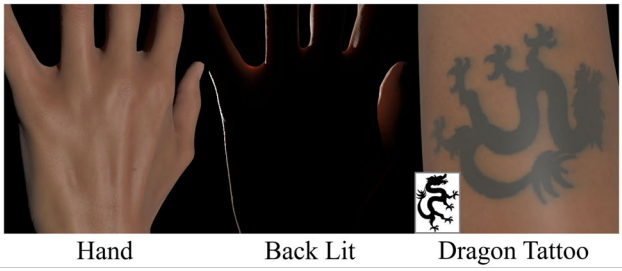

Figure 12. Rendering results of Asian skin type hand model. From left to right, original hand, transmittance effect of back illuminated and dragon tattoo on arm.

Figure 13. Another tattoo effect on the cheek.

Figure 14. Rendering results of pathology effects. We show the birthmark, freckle and vitiligo effect on the cheek by modifying biological coefficient textures without any artist involved.

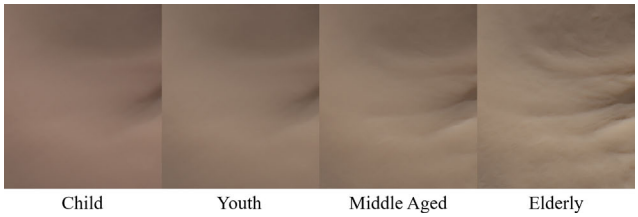

Figure 15. Rendering results of changing thickness and model details to simulate skin changes with age. We focus on eye corner which produce visible wrinkles with age. All results share the same albedo texture.

Fig. 15 illustrates the ability of our method to represent skin of different ages by changing the thickness of epidermis (older skin type with thinner epidermis) and model details. The "Child" and "Youth" results share the same epidermis thickness and model, while we modify the hemoglobin content of "Child" algebraically to present rosy skin effect.

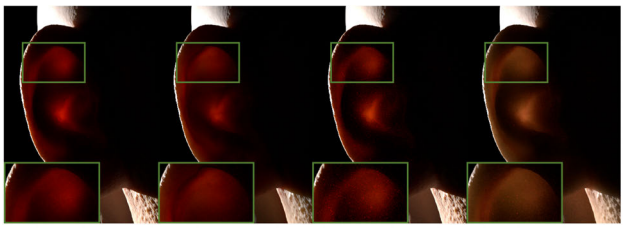

Figure 16. Transmittance effect of the ear. From left to right, we show the result of our method, followed by BSSRDF-based method, homogeneous method and Disney BSDF with subsurface method.

In Fig. 16, we show the transmittance effect at the ear. From left to right, we show the rendering results of our method, BSSRDF-based method [8], two-layer homogeneous method and Disney BSDF [4] with subsurface method. Compared to BSSRDF-based and Disney method that assumes semi-infinitely flat model which breaks in curved parts such as the ear, our method guarantees a more accurate transmitted intensity.

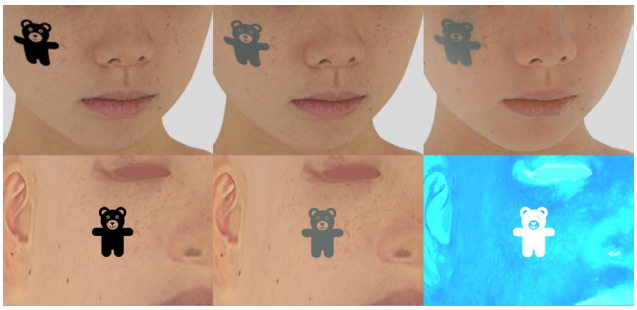

Figure 17. Comparison between random walk and our method on rendering tattoo. The texture corresponding to the rendering result is in the second line. Our method achieves photorealistic tattoo rendering by generating a dermis absorption coefficient texture. Even compared with random walk with modified texture, our method also has stronger subsurface scattering effect.

We also compare to widely used random walk method by rendering tattoos on the cheek. Tattooing is the penetration of black ink into the dermis, which presents cyan patterns under the combined interaction of hemoglobin and other pigments. As shown in Fig. 17, traditional random walk method only revert the original albedo, while our method realistically simulates the color presented by the ink in the dermis. We also changed the color of the tattoo pattern to cyan to render a relatively realistic tattoo effect using random walk, however, the patterns still have distinct hard borders while our method has obvious subsurface scattering effect. Random walk only uses homogeneous rendering to fit each albeo's pixel color, but our approach explicitly model the heterogeneous nature of the skin utilizing absorption texture. Both results exclude the specular term for a fair comparison. We use Blender to generate the result of random walk.

6. Discussion and Limitations

Due to the absorption coefficients calculation of the internal points of the layers, our implementation requires more intersection test and texture coordinates mapping than classic volume path tracing, which slow down rendering. As a result, we sample three light wavelength corresponding to RGB in order to obtain convergent results within a reasonable time. Compared to methods with a densely sampled wavelength, our work may introduce extra errors. Further, our biological two-layer model is unable to simulate some specific skin effects such as carotenemia (caused by β-carotene) since it cannot be modeled with hemoglobin and melanin concentrations. In the future, the two-layer model can consider more biological substances in the skin. However, this will increase the complexity of the LUT and the network can be more difficult to train.

7. Conclusion

Our work advances the realism of human skin modeling and rendering from the biophysical perspective. Unlike previous BSSRDF-based and random walk methods, we explicitly model the skin using layers of heterogeneous medium, and derive their volumetric scattering parameters from biological properties of spatially-varying melanin and hemoglobin, enabling controllable and biophysically-based appearance editing of human skin such as simulated pathological skin condition. To facilitate appearance acquisition, we further present a learning-based framework that automatically infers such biophysical parameters from a single albedo. Our model reproduces faithful and consistent skin appearance across different lighting, skin type and colors.

References

[1] C. Aliaga, C. Hery, and M. Xia. Estimation of spectral biophysical skin properties from captured rgb albedo. *arXiv e-prints*, pages arXiv–2201, 2022. 2

[2] G. V. Baranoski and A. Krishnaswamy. *Light and skin interactions: simulations for computer graphics applications*. 2010. 2

[3] B. Burley. Extending disney's physically based brdf with integrated subsurface scattering. In *ACM SIGGRAPH*, 2015. 3

[4] B. Burley and W. D. A. Studios. Physically-based shading at disney. In *ACM SIGGRAPH*, volume 2012, pages 1–7, 2012. 11

[5] P. Daszak, L. Berger, A. A. Cunningham, A. D. Hyatt, D. E. Green, and R. Speare. Emerging infectious diseases and amphibian population declines. *Emerging infectious diseases*, 5(6):735, 1999. 3

[6] C. Donner and H. W. Jensen. Light diffusion in multi-layered translucent materials. *ACM Transactions on Graphics (ToG)*, 24(3):1032–1039, 2005. 1, 3

[7] C. Donner and H. W. Jensen. A spectral bssrdf for shading human skin. *Rendering techniques*, 2006:409–418, 2006. 1, 2, 5, 6

[8] C. Donner and H. W. Jensen. Rendering translucent materials using photon diffusion. In *Acm siggraph 2008 classes*, pages 1–9. 2008. 3, 11

[9] C. Donner, T. Weyrich, E. d'Eon, R. Ramamoorthi, and S. Rusinkiewicz. A layered, heterogeneous reflectance model for acquiring and rendering human skin. *ACM transactions on graphics (TOG)*, 27(5):1–12, 2008. 1, 2

[10] H. Du, R.-C. A. Fuh, J. Li, L. A. Corkan, and J. S. Lindsey. Photochemcad‡: a computer-aided design and research tool in photochemistry. *Photochemistry and photobiology*, 68(2):141–142, 1998. 4

[11] L. Gevaux, J. Gierschendorf, J. Rengot, M. Cherel, P. Séroul, A. Nkengne, J. Robic, A. Trémeau, and M. Hébert. Real-time skin chromophore estimation from hyperspectral images using a neural network. *Skin Research and Technology*, 27(2):163–177, 2021. 4

[12] Y. Gitlina, G. C. Guarnera, D. S. Dhillon, J. Hansen, A. Lattas, D. Pai, and A. Ghosh. Practical measurement and reconstruction of spectral skin reflectance. In *Computer Graphics Forum*, volume 39, pages 75–89, 2020. 2, 6

[13] E. Golubev. Efficient screen-space subsurface scattering using burley's normalized diffusion in real-time. *Retrieved Aug*, 29:2019, 2018. 3

[14] R. Habel, P. H. Christensen, and W. Jarosz. Photon beam diffusion: A hybrid monte carlo method for subsurface scattering. In *Computer Graphics Forum*, volume 32, pages 27–37, 2013. 1, 3

[15] J. A. Iglesias-Guitian, C. Aliaga, A. Jarabo, and D. Gutierrez. A biophysically-based model of the optical properties of skin aging. In *Computer Graphics Forum*, volume 34, pages 45–55, 2015. 2

[16] S. L. Jacques. Optical properties of biological tissues: a review. *Physics in Medicine & Biology*, 58(11):R37, 2013. 5

[17] H. W. Jensen and J. Buhler. A rapid hierarchical rendering technique for translucent materials. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 576–581, 2002. 3

[18] H. W. Jensen, S. R. Marschner, M. Levoy, and P. Hanrahan. A practical model for subsurface light transport. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 511–518, 2001. 1, 3

[19] J. Jimenez and D. Gutierrez. Faster rendering of human skin. 2008. 3

[20] J. Jimenez, T. Scully, N. Barbosa, C. Donner, X. Alvarez, T. Vieira, P. Matts, V. Orvalho, D. Gutierrez, and T. Weyrich. A practical appearance model for dynamic facial color. In *ACM SIGGRAPH Asia 2010 papers*, pages 1–10. 2010. 2

[21] J. Jimenez, V. Sundstedt, and D. Gutierrez. Screen-space perceptual rendering of human skin. *ACM Transactions on Applied Perception (TAP)*, 6(4):1–15, 2009. 3

[22] J. Jimenez, D. Whelan, V. Sundstedt, and D. Gutierrez. Real-time realistic skin translucency. *IEEE Computer Graphics and Applications*, 30(4):32–41, 2010. 3

[23] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711, 2016. 7

[24] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7

[25] J. Křivánek and E. d'Eon. A zero-variance-based sampling scheme for monte carlo subsurface scattering. In *ACM SIGGRAPH 2014 Talks*, pages 1–1. 2014. 3

[26] P. Kutz, R. Habel, Y. K. Li, and J. Novák. Spectral and decomposition tracking for rendering heterogeneous volumes. *ACM Transactions on Graphics (TOG)*, 36(4):1–16, 2017. 6

[27] S. Prahl. Optical absorption of hemoglobin. *http://omlc. ogi. edu/spectra/hemoglobin*, 1999. 4

[28] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241, 2015. 7

[29] I. S. Saidi, S. L. Jacques, and F. K. Tittel. Mie and rayleigh modeling of visible-light scattering in neonatal skin. *Applied optics*, 34(31):7410–7418, 1995. 4

[30] E. V. Salomatina, B. Jiang, J. Novak, and A. N. Yaroslavsky. Optical properties of normal and cancerous human skin in the visible and near-infrared spectral range. *Journal of biomedical optics*, 11(6):064026, 2006. 5

[31] P. Shirley. A ray tracing method for illumination calculation in di use-specular scenes. In *Proceedings of Graphics Interface*, volume 90, pages 205–212, 1990. 3

[32] J. Stam. Multiple scattering as a diffusion process. In *Eurographics Workshop on Rendering Techniques*, pages 41–50, 1995. 3

[33] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *Josa*, 57(9):1105–1114, 1967. 9

[34] M. Van Gemert, S. L. Jacques, H. Sterenborg, and W. Star. Skin optics. *IEEE Transactions on biomedical engineering*, 36(12):1146–1154, 1989. 4

[35] D. Vicini, V. Koltun, and W. Jakob. A learned shape-adaptive subsurface scattering model. *ACM Transactions on Graphics (TOG)*, 38(4):1–15, 2019. 3

[36] Z. Wang, E. P. Simoncelli, and A. C. Bovik. Multiscale structural similarity for image quality assessment. In *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402, 2003. 8

[37] T. Weyrich, W. Matusik, H. Pfister, B. Bickel, C. Donner, C. Tu, J. McAndless, J. Lee, A. Ngan, H. W. Jensen, et al. Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics (ToG)*, 25(3):1013–1024, 2006. 9

[38] M. Wrenninge, R. Villemin, and C. Hery. Path traced subsurface scattering using anisotropic phase functions and non-exponential free flights. Technical report, Tech. Rep. 17-07, Pixar. https://graphics. pixar. com/library . . . , 2017. 1, 3

[39] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a percep-

tual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 8