https://doi.org/10.2352/lim.2024.5.1.20

A Neural Approach for Skin Spectral Reconstruction

Fereshteh Mirjalili[1] and Giuseppe Claudio Guarnera[1,2]; [1]Department of Computer Science, Norwegian University of Science and Technology (NTNU), Gjøvik, Norway; [2]Department of Computer Science, University of York, York, England

Abstract

Accurate reproduction of human skin color requires knowledge of skin spectral reflectance data, which is often unavailable. Traditionally, spectral reconstruction algorithms attempt to recover the spectra using commonly available RGB camera response. Among various methods employed, polynomial regression has proven beneficial for skin spectral reconstruction. Despite their simplicity and interpretability, nonlinear regression methods may deliver sub-optimal results as the size of the data increases. Furthermore, they are prone to overfitting and require carefully adjusted hyperparameters through regularization. Another challenging issue in skin spectral reconstruction is the lack of high-quality skin hyperspectral databases available for research. In this paper, we gather skin spectral data from publicly available databases and extract the effective dimensions of these spectra using principal component analysis (PCA). We show that plausible skin spectra can be accurately modeled through a linear combination of six spectral bases. We propose a new approach for estimating the weights of such a linear combination from RGB data using neural networks, leading to the reconstruction of spectra. Furthermore, we utilize a daylight model to estimate the underlying scene illumination metamer. We demonstrate that our proposed model can effectively reconstruct facial skin spectra and render facial appearance with high color fidelity.

Introduction

Faithful representation of human skin color is crucial for virtual and augmented reality [1], skin cancer detection [2], forensics [3], prosthetic limb manufacturing [4], and the cosmetic industry [5]. Indeed, this requires access to detailed spectral reflectance data of the skin. Hyperspectral cameras are able to record the skin reflectance at high spectral and spatial resolution [6]. However, these cameras are significantly costly compared to standard consumer RGB cameras. Additionally, the time-consuming process of high-resolution hyperspectral acquisition renders it impractical for human subjects, who must remain still during the acquisition. On the other hand, despite being more affordable, RGB cameras utilize only three color sensors to capture the scene radiance. As a result, RGB imaging can only gather a limited information encoded in the scene. To address this, spectral reconstruction algorithms have been developed to infer the missing information from the RGB image of the scene [7]. Although the problem is heavily underconstrained, it is still feasible to upsample from three channels to more, given that most hyperspectral bands are highly correlated [8].

The simplest algorithms for spectral reconstruction tasks utilize regression methods. Linear regression, for example, maps RGB values to spectra through a linear transformation [7]. Extending this approach, nonlinear regression methods transform RGB values into a set of higher-order polynomial or root-polynomial terms, offering a more complex model for spectral reconstruction. These spectral reconstruction approaches are all pixel-based, meaning that they recover the spectrum of a pixel independent of the spatial location of that pixel in the image. In contrast, deep neural network approaches [9] leverage RGB patches for training. Indeed, they need a large amount of data and significantly more powerful processing and mapping architectures.

There is a very limited body of work focused on the spectral reconstruction of human skin from RGB images. One potential reason for this is the difficulty in acquiring high-quality hyperspectral data specifically for human skin. Although several hyperspectral image databases of outdoor and indoor scenes are publicly available for research [11], high-quality databases dedicated to skin are scarce, to our knowledge. This results in significant challenges for spectral reconstruction algorithms, especially for patch-based methods like deep neural networks, in accurately reconstructing large skin patches, such as human faces. In this paper, we introduce a simple paradigm for skin spectral reconstruction utilizing neural networks. Initially, we demonstrate the feasibility of accurately modeling a plausible skin spectrum with a set of spectral bases. Following this, we propose a neural network-based procedure to determine the set of weights required to linearly combine these bases for reconstructing the skin spectra. This method also includes estimating the spectrum of the scene illumination. We further validate the effectiveness of our method by demonstrating its application in rendering human faces across various skin tones.

Background

Most off-the-shelf RGB digital cameras contain three types of sensors, red, green and blue, each with a different sensitivity to a broad range of visible light spectrum. Each sensor $k = r,g,b$ integrates spectral data over the visible spectral range $\omega = [400-700nm]$; overall, camera sensors provide trichromatic color responses of the scene:

$$c_k = \int_{\omega} s(\lambda) l(\lambda) q_k(\lambda)\, d\lambda + \varepsilon_k, \qquad \{k = r,g,b\} \tag{1}$$

where $s(\lambda)$ denotes the spectral reflectance of the scene imaged by the camera, $l(\lambda)$ refers to the illumination spectrum, $q_k(\lambda)$ is the spectral sensitivity of the camera sensor of type k, ε_k is the acquisition noise, and $c = [c_r, c_g, c_b]$ is the trichromatic camera response. Given that spectral data are available at discrete intervals, the integration can be approximated by inner products as follows:

$$c = Q^T s, \tag{2}$$

where Q is a matrix derived from the element-wise product of $l(\lambda)$ and $q(\lambda)$, and T is matrix transpose. All spectral reconstruction methods aim to accurately recover s using the information embedded in c. Among the early methods, linear regression establishes a relationship between c and Q through a single linear transformation matrix. Considering the entire database of N spectral reflectances S, and their corresponding RGB data C, one can express the relationship as follows:

$$CM = S, \tag{3}$$

where M denotes the linear transformation matrix, which is derived using error minimization criteria such as the least squares method. It is well established that recovering spectral reflectance from camera responses is inherently prone to the problem of metamerism: more than one reflectance can give rise to the same RGB values [12]. Furthermore, the task becomes even more challenging when dealing with complex surfaces, such as human skin, which comprises several layers of cutaneous tissue, each with distinct optical properties [13, 31]. In such cases, a simple linear transform will not suffice to accurately approximate the spectra. To address nonlinear relationships, polynomial [7] and root-polynomial [14] regressions, and neural networks [15, 16] are likely to provide better reconstruction performance. In the domain of skin spectral reconstruction, the literature suggests that most of the existing research has benefited from regression methods. In a pioneering work, Imai et al. [17] used second-order polynomial regression to recover the spectra for color reproduction on a CRT display. Xiao et al. [18] combined polynomial regression with principal component analysis (PCA) to reconstruct skin spectra, benefiting from access to comprehensive spectrophotometric measurements of human skin spectra and silicone skin samples. Recently, Li et al. [19] have aimed to enhance the accuracy of skin spectral reconstruction using polynomial regression by reevaluating the strategy for preparing the training dataset.

Method

We propose a straightforward neural upsampling procedure to transform measured RGB values of skin into spectral data, assuming these RGB values are acquired under broadband daylight illumination conditions. We will first outline the preparation of the training dataset tailored for this method. Additionally, for comparison purposes, we apply regression-based spectral reconstruction models on the same dataset.

Skin dataset

In this research, we use three publicly available databases of human skin spectral reflectance. The first database [18] includes 4,392 measurements from 482 subjects of Caucasian, Chinese, and Kurdish ethnicity, taken from the forehead, cheek, inner arm, and back of the hand, covering a spectral range of 360 to 740 nm at 10 nm intervals. The second database [20] contains measurements of the right inner forearm from 100 subjects, spanning 250 to 2500 nm at 3 nm intervals. Both databases have no age or gender constraints. Additionally, we use the ISET Hyperspectral Image Database [21], which has 24 hyperspectral images of human faces, covering 400 to 950 nm at 3.5 nm intervals. We selected two 10×10 pixel patches from the forehead and cheek, excluding specular highlights. We then clipped the data to 400 to 700 nm and interpolated it to 31 spectral bands at 10 nm intervals. Altogether, we have a skin spectral reflectance dataset of approximately 14,600 measurements, as illustrated in Figure 1 (left). For all spectra, we calculate the corresponding CIEXYZ tristimulus values and CIELAB color values, i.e., L*, a*, b*, and C*, using the CIE 1931 color matching functions, under the CIE standard D65 illuminant [22]. In Figure 1 (right), we also show the distribution of the calculated CIELAB color values for both non-facial and facial skin on the a*-b* and C*-L* chromatic planes. We observe that the color gamut of facial skin falls within the color gamut of non-facial skin. Consequently, we anticipate that a model trained on skin color data from non-facial regions could optimistically be applied to the reconstruction of facial skin color. We note that such a model is highly unlikely to work properly for facial hair, lips, and special features such as freckles, moles, and tattoos.

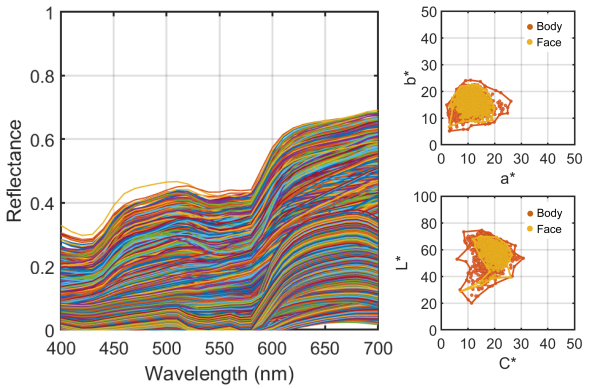

Figure 1. *Spectral curves of 14,600 skin reflectance measurements, gathered for our research from three publicly available databases [18, 20, 21]* ***(left)****, CIELAB gamuts of facial and non-facial skin colors* ***(right)****.*

Regression-based spectral reconstruction

We utilize linear regression (LR), 6th order polynomial regression (PR6), and 6th order root-polynomial regression (RPR6) in our skin spectral reconstruction experiment. To create the training set for these models, we make use of the gathered skin reflectance spectra dataset. For each spectrum, we calculate the corresponding RGB values using the CIE 1931 color matching functions following Equation 2. For illumination, we employ the new daylight equation proposed in [23]:

$$l(\lambda, T) = \mu_1 \lambda^{-5} f(\lambda) e^{-\mu_2 / T\lambda}, \tag{4}$$

where T refers to daylight correlated color temperature (CCT), $f(\lambda)$ denotes a filter correction function, $l(\lambda)$ is the daylight spectrum, and μ_1 and μ_2 are constants equal to 3.7418×10^{-16} Wm^2 and $1.4388 \times 10^{-2} mK$, respectively.

To ensure our trained models are *versatile* and not specific to any illumination, we calculate skin RGBs across a range of natural daylight CCTs from warm to cool daylight. Empirical studies show that a shift of 5.5 Mired temperature units (reciprocal color temperature $\frac{10^6}{T}$) is required for the difference in light color to become perceptible [24]. Therefore, we sample the CCT range of 2500K to 15000K at intervals of 5.5 Mired units. To help our models maintain performance with changing light intensity, we augment the training data by scaling the skin reflectance spectra with constants from 0.1 to 2, simulating different exposure levels. Following [10], we prepare the training dataset for the regression models by dividing the spectra and their corresponding RGB values into four subsets. Two subsets are used for training, one for validation and optimizing regularization parameters, and the final subset for testing. Nonlinear regression models risk overfitting, where the model performs well on training data but poorly on unseen data. To prevent this, we use ridge regularization [10].

While regression-based spectral reconstruction is simple and interpretable, it can become suboptimal and time-consuming with larger datasets. Managing model complexity is crucial to prevent overfitting. Although regression models offer regularization strategies like ridge regularization, advanced regularization methods in neural networks may better manage overfitting in large, complex datasets. Therefore, we propose a neural approach for skin spectral reconstruction, demonstrating that skin spectral reflectance can be accurately modeled with a set of spectral bases.

Skin spectral bases

The reflectance spectra of most real-world surfaces, such as human skin, have been proven to be naturally smooth (see Figure 1, left). Consequently, those spectra can be accurately modeled using a limited number of spectral bases, often as few as six to nine [8]. Inspired by this fact, we perform principal component analysis (PCA) on the database of skin spectra to obtain the first six spectral bases (i.e., PCA eigenvectors) to represent all the spectra. Hence, with appropriate weighting factors, any skin spectrum can be modeled as a linear combination of the six spectral bases:

$$s'(\lambda) = \sum_{k=1}^{6} w_k b_k(\lambda), \tag{5}$$

where w_k represents the weight for the kth spectral basis, $b_k(\lambda)$ denotes the kth spectral basis, and $s'(\lambda)$ signifies the derived skin spectral reflectance. In this paper, we define the range of λ as being from 400 to 700 nm sampled at 10 nm intervals. Figure 2 (left) illustrates the six skin spectral bases derived using PCA. To demonstrate the effectiveness of this simple model, we calculate the weights w_k for two skin spectra from the database, one representing a light skin tone and the other representing a darker one, using the Penrose–Moore pseudo inverse method [26]. We then apply Equation 5 to reconstruct the spectra.

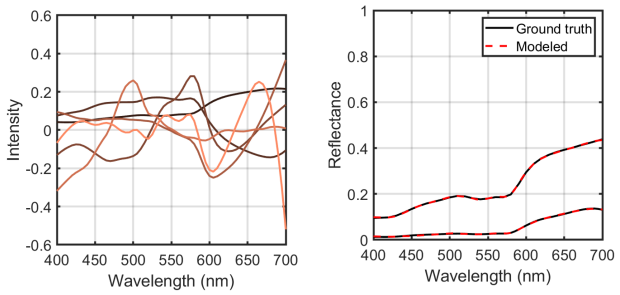

Figure 2. *The six skin spectral bases derived using PCA to represent the entire skin spectral data* ***(left)****, two pairs of ground truth skin spectra and their estimations using the six spectral bases following Equation 5* ***(right)****.*

Figure 2 (right) presents both the ground truth and the corresponding modeled skin spectra. The nearly perfect match between the ground truth and the reconstructed spectra indicates that a plausible skin spectrum can be accurately represented using six linear spectral bases. It is noteworthy that we explored both a lower and higher number of bases in our experiments. Utilizing fewer bases led to sub-optimal reconstructions, while increasing the number of bases did not enhance the model's performance.

Neural spectral reconstruction

By reducing the effective dimensionality of skin spectra, we enable the use of shallow neural networks for skin spectral reconstruction. We design a multilayer perceptron (MLP) and train it to learn the mapping from the space of skin RGB to the space of skin spectral bases. This involves finding a set of weights to linearly combine these bases to accurately reconstruct the skin spectra. Figure 3 provides a schematic overview of our MLP architecture designed for skin spectral reconstruction. The architecture of our MLP comprises three hidden layers with 128, 512, and 128 nodes, respectively. Both the input and hidden layers utilize the ReLU activation function. Our observations indicate that adding more hidden layers or nodes results in only a slight enhancement of training performance. The final layer consists of 6 nodes and employs the sigmoid activation function to provide the estimated weights. The MLP hyperparameters are optimized using the adaptive moment estimation (Adam) optimizer, aiming to minimize the mean squared error (MSE) between the predicted and ground truth weights. For training our MLP, we employ the same dataset used for the regression models. For each skin spectrum, in addition to RGB values, we determine a set of six weights w_k as outlined in Equation 5. We then split the dataset into two subsets: one for training and the other for validating the MLP. On average, the entire training duration is approximately 45 minutes, using a GPU-accelerated desktop PC equipped with an Intel(R) Core(TM) 14th Gen i9 3.60 GHz processor and 32GB of RAM.

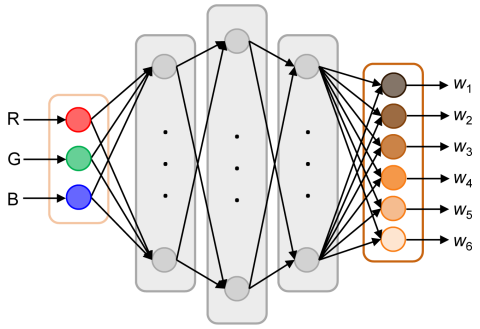

Figure 3. *Schematic representation of our fully-connected multilayer perceptron (MLP).*

Results

The spectral reconstruction algorithms recover an estimate of the reflectance spectrum for the given ground truth spectrum. To measure how closely the estimated spectrum matches the ground truth, we employ three widely used evaluation metrics: the Mean Relative Absolute Error (MRAE) [27], the Root-Mean Squared Error (RMSE) [28], and the CIEDE2000 color difference ΔE_{00} [29]. Figure 4 illustrates the spectral curves of three reconstructed skin spectra by two of the tested models, i.e., PR6 and the proposed MLP.

In Table 1, we present both the mean and the 99th percentile for MRAE, RMSE, and ΔE_{00} across 150 skin spectra that were not seen by the trained algorithms. We compare the performance of linear regression (LR), 6th order polynomial regression (PR6), 6th order root-polynomial regression (RPR6), and our MLP in accurately recovering these unseen skin spectra.

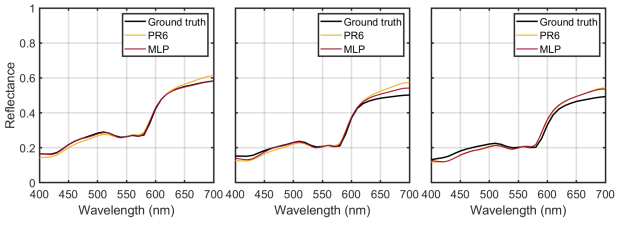

Figure 4. *Spectral reconstruction results for three skin spectra using the 6th order polynomial regression (PR6) and the MLP, shown from left to right, correspond to patches 3, 7, and 11 in Figures 5-6, respectively.*

It is noteworthy that our exploration extended to polynomial regression orders beyond 6. Despite a significant increase in training time, the enhancement in reconstruction performance was minimal. We also assess the performance of the tested models at three brightness levels i.e., original, half and double, to

ensure consistent performance despite variations in scene exposure. Essentially, the models should recover spectra with consistent shapes across different magnitudes under various exposure settings. To simulate changes in scene exposure, we scale the RGB values by constants of 1 (original), 0.5 (half), and 2 (double) following the method in [17]. After reconstructing the respective spectra, we then compare these spectra with the corresponding ground truth spectra, which are also scaled by the same constants.

Table 1: Evaluation statistics under original exposure.

	Mean			99th percentile		
Model	MRAE	RMSE	ΔE_{00}	MRAE	RMSE	ΔE_{00}
LR	0.064	0.021	2.546	0.119	0.042	3.776
PR6	0.039	0.012	1.274	**0.076**	**0.025**	**3.291**
RPR6	0.039	0.013	1.313	0.081	0.030	3.414
MLP	**0.033**	**0.010**	**1.011**	0.094	**0.025**	3.326

The mean results in Table 1 show that under original exposure, the proposed MLP demonstrates superior performance across all evaluated metrics, indicating its effectiveness in accurately reconstructing skin spectra. Furthermore, when considering the 99th percentile scores, the MLP model remains robust, providing reliable spectral reconstructions even in the tail end of the distribution. As expected, both PR6 and RPR6 markedly outperform the linear regression approach, with PR6 performing slightly better than RPR6. For a pictorial visualization of the results, Figure 5 showcases the RGB renderings of 24 pairs of ground truth and recovered skin spectra by the MLP and the best performing regression model i.e., PR6 in the form of skin patches. In each image, the ground truth and recovered colors are displayed within the same patch, with the inner circle representing the reconstructed spectrum and the outer frame indicating the ground truth spectrum. We see that under original exposure, the spectra reconstructed by the MLP yield RGB colors that more closely match the ground truth RGB colors when compared to those reconstructed by PR6.

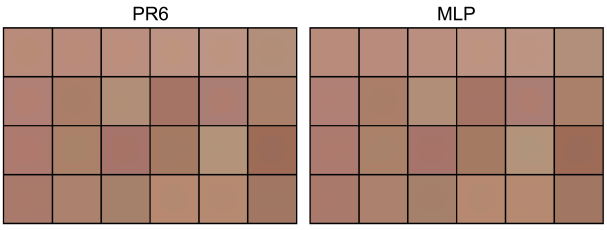

Figure 5. *Spectral reconstruction results for 24 skin spectra by the 6th order polynomial regression (PR6), and the MLP, under original exposure. The RGB renderings are performed for the CIE 1931 color matching functions and D65 illumination. In each patch, the inner circle corresponds to the reconstructed spectrum and the outer frame corresponds to the ground truth spectrum. The mean ΔE_{00} color difference errors for PR6 and MLP are 1.39 and 0.95, respectively.*

When the exposure varies, all the tested models perform slightly worse. However, the models retain their performance better when the exposure increases. The MLP performs considerably worse when the testing exposure is halved, while the performance of the LR slightly improves. This trend aligns with the findings reported in [14], which demonstrated that neural networks for spectral reconstruction tend to underperform with changes in image exposure. Although we have tried to overcome this by augmenting the training data, the MLP still exhibits this shortcoming. As mentioned previously, we augmented the training dataset not only by a set of scales to address the models' exposure variability but also by training our the models with RGB data calculated across a broad range of daylight CCTs. This approach enables us to reconstruct the skin spectra from RGB data regardless of the underlying illumination. To verify the illumination-independency of our proposed model, we reconstruct the 150 test skin spectra for two extreme daylight conditions, i.e., warm with a CCT of 3500K and cool with a CCT of 9000K. The reconstruction results for the same set of 24 test spectra are depicted in Figure 6. The MLP reconstruction errors, in terms of mean ΔE_{00} color difference are 1.065 and 1.019 for the CCTs of 3500K and 9000K, respectively. This indicates that the trained MLP can successfully reconstruct skin spectra under various daylight illumination conditions.

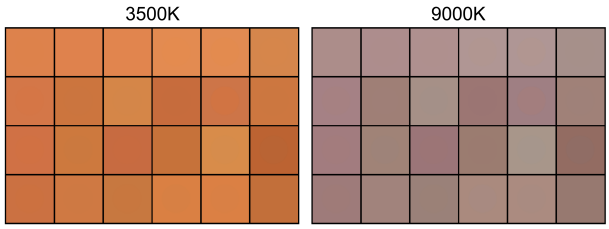

Figure 6. *Spectral reconstruction results for 24 skin spectra by the MLP, under daylight illumination with CCTs of 3500K and 9000K. In each patch, the inner circle displays the reconstructed color and the outer frame displays the ground truth color. The mean ΔE_{00} color difference errors for CCTs of 3500K and 9000K are 1.065 and 1.019, respectively.*

Facial spectral reconstruction

As discussed previously, we anticipate that a model trained on skin color data from non-facial regions such as inner forearm and back of the hand could be effectively applied to reconstruct the spectral reflectance of facial skin. To evaluate this hypothesis, we employ our trained MLP to reconstruct the skin reflectance spectra of facial images across various skin tones, including fair, light, and brown. To capture the facial photographs, our acquisition setup features a lighting arrangement with broadband cool white LED panels that simulate daylight with a CCT of 7000K, and a Canon EOS 850D DSLR camera. The LED panels are all cross-polarized with respect to the camera, allowing specular cancellation. The raw facial RGB images are inputted into the MLP to estimate the weights of the skin spectral bases for each pixel in the image, and subsequently reconstruct the spectrum of that pixel, thereby generating what is termed a hyperspectral facial image. After reconstructing the spectra, the next step involves estimating the underlying illumination of the face image. To achieve this, we employ the Nelder-Mead simplex method [30], in combination with the new daylight model proposed in [24]. We aim to minimize the ΔE_{00} color difference between the ground truth and reconstructed face images, through iteratively adjusting the CCT value in Equation 2, until the lowest ΔE_{00} is achieved. By following this estimation approach, we acknowledge that the estimated illumination of the image, while having a matching CCT, represents a metamer of the actual scene illumination with a different spectral power distribution (SPD), as illustrated in Figure 8.

Figure 7 shows comparisons between the ground truth facial images and their reconstructions using the MLP and the illumination estimation paradigm described above for three subjects. We can see that the face reconstructions are a close match to the ground truth photographs, indicating high fidelity in the reconstructions. Across all three subjects, the most significant dis-

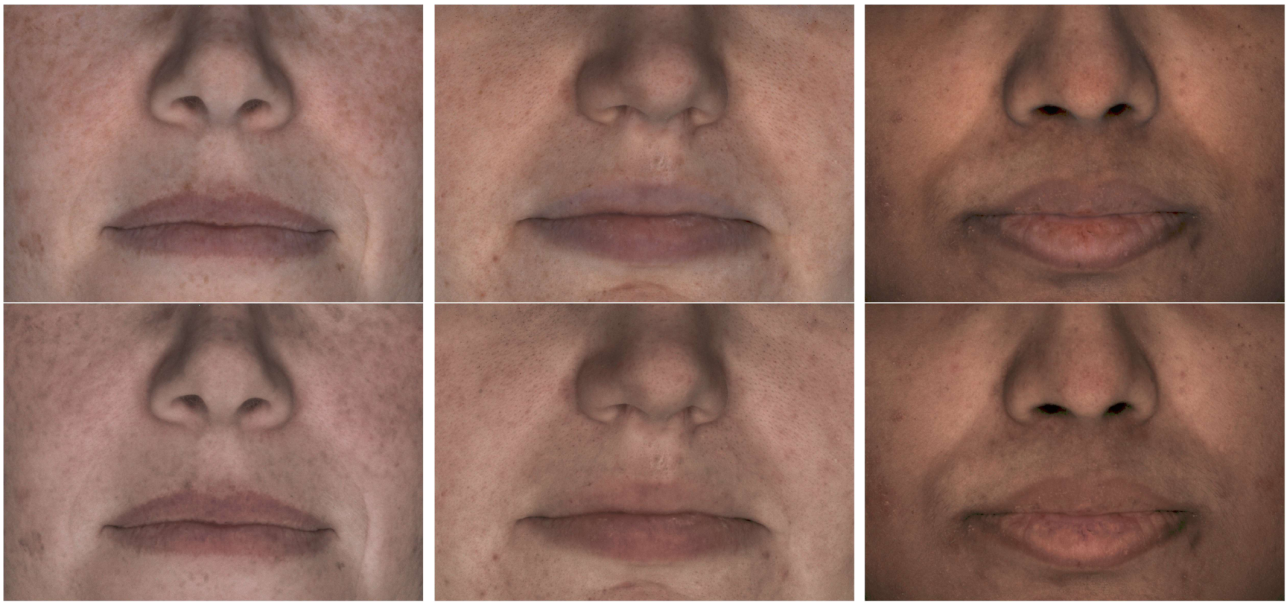

Figure 7. *Comparisons between the ground truth facial images* ***(top row)*** *and their reconstructions* ***(bottom row)*** *using our proposed model. The subjects represent 'fair,' 'light,' and 'brown' skin tones, from left to right.*

crepancies between the ground truth and the recovered images are observed in the lips and surrounding regions, as well as in areas with freckles, brown spots, and dimmer pixels. While the facial features are precisely detailed in the reconstructed images, the model falls short in accurately reproducing the colors in these areas. This limitation was anticipated, considering the training set does not include measured spectra for lips and certain skin features.

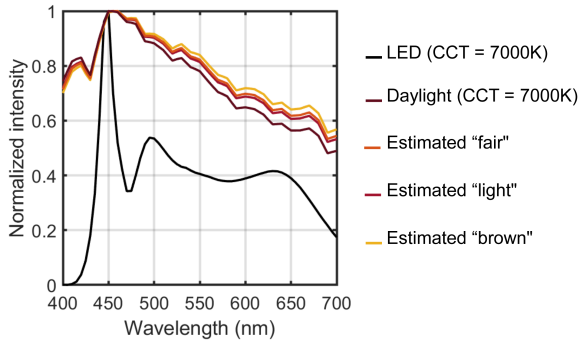

Figure 8. *The spectral power distribution (SPD) of the LED panels used in our facial acquisition setup, with a CCT of 7000K, is shown in black. The corresponding daylight spectrum, calculated using Equation 4, is depicted in dark brown. The estimated illumination spectra for the facial images of the subjects with 'fair', 'light', and 'brown' skin tones are illustrated in distinct colors.*

To effectively evaluate the color reproduction fidelity of our method, we select patches on the subjects' foreheads (excluding the lips). We then calculate the ΔE_{00} color difference between the ground truth and reconstructed patches. Figure 9 displays the renderings of these patches, along with their corresponding color reproduction error maps in terms of ΔE_{00}. The ΔE_{00} error maps reveal a mean ΔE_{00} of 1.24 for the subject with a fair skin tone, 1.18 for the subject with a light skin tone, and 1.71 for the subject with a brown skin tone. We can see that the model performs less effectively for the subject with a brown skin tone, and the one with pronounced facial features, like freckles. This is primarily due to the fact that the dataset used for training the MLP contains fewer spectra of darker skin tones compared to those acquired from subjects with lighter skin tones. This underscores the critical need for a well-balanced, high-quality hyperspectral skin database to ensure the model's robustness and accuracy across a wide range of skin tones and features.

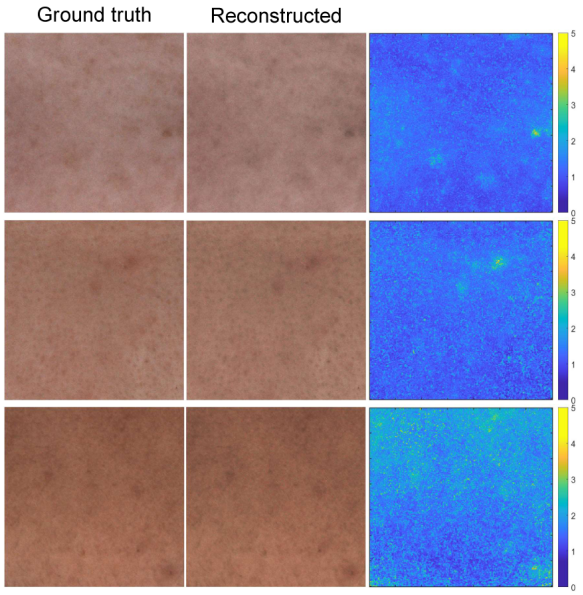

Figure 9. *Comparisons between ground truth and their reconstructed skin patches, selected from the forehead of the subjects with 'fair' (top row), 'light' (middle row), and 'brown' (bottom row) skin tones. The color reproduction error maps are also displayed on the left.*

Conclusion

In this paper, we introduce a new method for high-fidelity spectral reconstruction of facial RGB images. Our approach eliminates the need for additional information such as color chart measurements, the camera's spectral sensitivities, or the scene's

illumination spectrum. We achieve close matches to ground truth facial photographs across a range of skin tones, demonstrating the method's effectiveness. We recognize that access to a high-quality facial hyperspectral database would further enhance the robustness of our results. We are optimistic about achieving this in future work.

Acknowledgments

This work was supported by the Spectraskin project N-288670, funded by the Research Council of Norway.

References

[1] T.C. Peck, J.J. Good, A. Erickson, I. Bynum, G. Bruder, "Effects of Transparency on Perceived Humanness: Implications for Rendering Skin Tones Using Optical See-Through Displays," IEEE Trans. Vis. Comput. Graph., 28(5), 2179-2189 (2022).

[2] M. Fijalkowska, M. Koziej, E. Zadzinska, B. Antoszewski, A. Sitek, "Assessment of the Predictive Value of Spectrophotometric Skin Color Parameters and Environmental and Behavioral Factors in Estimating the Risk of Skin Cancer: A Case-Control Study," J. Clin. Med., 11(11), 2969 (2022).

[3] K.R. Scafide, D.J. Sheridan, J. Campbell, V.B. Deleon, M.J. Hayat, "Evaluating change in bruise colorimetry and the effect of subject characteristics over time," Forensic Sci. Med. Pathol., 9(3), 367-376 (2013).

[4] X. Hu, W.M. Johnston, R.R. Seghi, "Measuring the color of maxillofacial prosthetic material," J. Dent. Res., 89(12), 1522-1527 (2010).

[5] C. ArceLopera, T. Igarashi, K. Okajima, "Colorimetric Analysis of Makeup Styles and Their Relation with Visual Quality Perception of the Skin," J. Vis., 14(10), 461 (2014).

[6] G.N. Stamatas, C. Balas, N. Kollias, "Hyperspectral image acquisition and analysis of skin," in SPIE BiOS, 2003.

[7] V. Heikkinen, R. Lenz, T. Jetsu, J. Parkkinen, M. Hauta-Kasari, T. Jaaskelainen, "Evaluation and unification of some methods for estimating reflectance spectra from RGB images," J. Opt. Soc. Am. A, 25(10), 2444-2458 (2008).

[8] J.P.S. Parkkinen, J. Hallikainen, T. Jaaskelainen, "Characteristic spectra of Munsell colors," J. Opt. Soc. Am. A, 6(2), 318-322 (1989).

[9] B. Arad, R. Timofte, O. Ben-Shahar, et al., "NTIRE 2020 Challenge on Spectral Reconstruction from an RGB Image," in 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 1806-1822, 2020.

[10] Y.-T. Lin, G.D. Finlayson, "On the Optimization of Regression-Based Spectral Reconstruction," Sensors, 21(16), 5586 (2021).

[11] J. Zhang, R. Su, Q. Fu, et al., "A survey on computational spectral reconstruction methods from RGB to hyperspectral imaging," Sci. Rep., 12, 11905 (2022).

[12] Y.-T. Lin, G.D. Finlayson, "Physically Plausible Spectral Reconstruction," Sensors, 20(21), 6399 (2020).

[13] R.R. Anderson, J.A. Parrish, "The Optics of Human Skin," J. Investig. Dermatol., 77(1), 13-19 (1981).

[14] Y.-T. Lin, G.D. Finlayson, "Exposure Invariance in Spectral Reconstruction from RGB Images," in International Conference on Communications in Computing, 2019.

[15] A. Ribes, F. Schmitt, "Reconstructing Spectral Reflectances with Mixture Density Networks," in Proceedings of the First European Conference on Colour in Graphics, Imaging, and Vision, 486-491, Poitiers, France, 2002.

[16] T. Stiebel, S. Koppers, P. Seltsam, D. Merhof, "Reconstructing Spectral Images from RGB-Images Using a Convolutional Neural Network," in 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 1061-10615, 2018.

[17] F.H. Imai, N. Tsumura, H. Haneishi, Y. Miyake, "Principal component analysis of skin color and its application to colorimetric color reproduction on CRT display and hardcopy," J. Imaging Sci. Technol., 40, 422-430 (1996).

[18] K. Xiao, Y. Zhu, C. Li, D. Connah, J.M. Yates, S. Wuerger, "Improved method for skin reflectance reconstruction from camera images," Opt. Express, 24(13), 14934-14950 (2016).

[19] S. Li, K. Xiao, P. Li, "Spectra Reconstruction for Human Facial Color from RGB Images via Clusters in 3D Uniform CIELab* and Its Subordinate Color Space," Sensors (Basel), 23(2), 810 (2023).

[20] C. Cooksey, D. Allen, B. Tsai, "Reference Data Set of Human Skin Reflectance," J. Res. (NIST JRES), 2017.

[21] T. Skauli, J. Farrell, "A collection of hyperspectral images for imaging systems research," Proc. SPIE 8660, Digital Photography IX, 86600C (2013).

[22] R.W.G. Hunt, M.R. Pointer, "Measuring Colour," Wiley, 2011.

[23] R. Deeb, G.D. Finlayson, E. Daneshvar, "Locus Filters: Theory and Application," J. Imaging Sci. Technol., 2023, pp 1-11.

[24] R. Deeb, G. Finlayson, "Locus filters," Opt. Express, 30(8), 12902-12917 (2022).

[25] D. Connah, J. Hardeberg, "Spectral recovery using polynomial models," in Color Imaging X: Processing, Hardcopy, and Applications, vol. 5667, San Jose, CA, USA, 65-75, SPIE, January 2005.

[26] J.C.A. Barata, M.S. Hussein, "The Moore–Penrose Pseudoinverse: A Tutorial Review of the Theory," Gen. Appl. Phys., 42, 146-165 (2012).

[27] B.L. Bowerman, R.T. O'Connell, A.B. Koehler, "Forecasting, Time Series, and Regression: An Applied Approach," 2005.

[28] K. Tyagi, C. Rane, Harshvardhan, M. Manry, "Chapter 4 - Regression analysis," in Artificial Intelligence and Machine Learning for EDGE Computing, R. Pandey, S.K. Khatri, N. Kumar Singh, P. Verma, Eds., Academic Press, 53-63, 2022.

[29] G. Sharma, W. Wu, E.N. Dalal, "The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations," Color Res. Appl., 30, 21-30 (2005).

[30] J.A. Nelder, R. Mead, "A Simplex Method for Function Minimization," The Computer Journal, 7(4), 308-313 (1965).

[31] Y. Gitlina, G.C. Guarnera, D.S. Dhillon, J. Hansen, A. Lattas, D.K. Pai, A. Ghosh, "Practical Measurement and Reconstruction of Spectral Skin Reflectance," Comput. Graph. Forum, 39(4), 75-89 (2020).

[32] Y.-T. Lin, G.D. Finlayson, "A Rehabilitation of Pixel-Based Spectral Reconstruction from RGB Images," Sensors (Basel, Switzerland), 23, 2023.

Author Biography

Fereshteh Mirjalili *received her Ph.D. in color science from the Amirkabir University of Technology (Tehran Polytechnic), Iran, in 2014. She has worked as a postdoctoral fellow at the Department of Computer Science, Norwegian University of Science and Technology (NTNU) since 2018, where she is now a guest researcher. Her research interests lie in the domain of material appearance measurement and perception.*

Giuseppe Claudio Guarnera *received his Ph.D. degree in computer science from the University of Catania (Italy), with a doctoral dissertation in computer vision and pattern recognition. He is a senior researcher at the Department of Computer Science, Norwegian University of Science and Technology (NTNU), and a lecturer in Computer Vision and Graphics at the University of York (UK). He has been a guest researcher at the Department of Psychology, Justus-Liebig University Giessen (DE) since 2017. His research interests include computer vision, computer graphics and applications of visual perception in computer graphics.*