Constraining a statistical skin colour model to adapt to illumination changes

Moritz Störring and Erik Granum

Computer Vision and Media Technology Laboratory
Institute of Electronic Systems, Aalborg University
Niels Jernes Vej 14, DK-9220 Aalborg, Denmark
{mst,eg}@cvmt.auc.dk
http://www.cvmt.dk

Abstract. This paper investigates how accurately the covariance matrix of chromaticity distributions of facial skin might be modelled for different illumination colours using a physics-based approach. Results are presented using real image data taken under different illumination colours and from subjects with different shades of skin. The orientation of the eigenvectors of the modelled and measured covariances deviate in average about 4°. The method has application in skin colour segmentation, e.g., to describe the confidence of a segmented skin area.

1 Introduction

Human face tracking systems are becoming more robust by fusing several cues such as motion and shape [3, 14, 10]. An often used cue is skin colour segmentation. Several approaches have been proposed, some statistically based, e.g. [6, 15], and some physics based, e.g. [9, 12].

A problem when using skin colour as a feature arises under varying lighting conditions. In particular changes in the spectral composition of scene illumination may result in failures of colour segmentation methods [5].

Yang *et.al.* [15] showed that the facial skin chromaticity distribution of an individual under a single light source may be approximated by a multivariate normal distribution in the red-green chromaticity plane. They proposed an adaptive statistical skin colour model updating the mean vector and the covariance matrix as the lighting changes. The model is used in a real-time face tracker and works under slightly changing indoor illumination conditions.

McKenna *et.al.* [6] use Gaussian mixtures to model the skin colour distribution in Hue-Saturation space. The parameters are updated over time in order to adapt to changes in illumination and viewing direction.

A problem with adapting a colour model during tracking is the lack of ground-truth of the region of interest [6], i.e. the colour model might adapt to image regions which do not belong to the skin coloured object.

In [12] skin chromaticities are modelled by a physics-based approach for different illuminations with a good approximation using knowledge about the camera parameters and assuming that commonly used in- and outdoor light sources can be modelled

by blackbody radiators [4]. The skin chromaticities for a variety of illuminations with different correlated colour temperatures (CCT) form a 'skin locus' which follows the curvature of the Planckian locus of blackbody radiators. This might be used to constrain the search area for skin colour in the chromaticity plane.

In Soriano *et.al.* [9] the search area is constrained by the skin locus and inside the skin locus a nonparametric skin colour model is learned and updated by histogram backprojection. They show a face tracking system working outdoors under changing illumination conditions. However, histogram backprojection has the drawback that the histogram might adapt to non skin coloured objects in the background. This effect might be avoided if the skin colour model would be constrained by physics-based knowledge about possible skin distributions.

In this paper a statistical model is constrained using a physics-based approach. In particular, it is investigated 1) How well may two dimensional skin chromaticity distributions be approximated by a normal distribution, i.e. by mean values and covariance matrices; 2) How accurately may the eigenspace of a skin chromaticity covariance matrix be modelled for illuminations with arbitrary CCTs by linking it with a physics-based approach.

Modelling of the covariance matrix has application in adaptive statistical models as used in [6, 9, 15]. It might give an indication on the current confidence about that the segmented skin area is the same as the one, e.g., in the previous frames or before an illumination change. The confidence measure might be obtained by matching the measured and modelled eigenspaces against each other. Confidence measures are especially of interest in systems where several tracking cues are fused [3].

2 Background

2.1 Skin chromaticities

Human skin reflections may be modelled with the Dichromatic Reflection Model [8] as surface and body/matte reflections. Most of the facial skin area shows 'pure' body reflections [2, 12]. Even under direct illumination highlights occur usually not on the entire face but only on some areas, e.g. the nose, forehead, or cheeks, which might be filtered [13]. In the following only body reflections are considered. Furthermore, it is assumed that the investigated skin areas are illuminated by a 'single' light source. This 'single' light source may be a mixture of several sources having different spectral compositions. The constraint is that the mixture is uniform for the investigated skin area. For spatially non-uniform illumination see e.g. [1].

In colour machine vision usually each pixel is represented by a 3D vector $\boldsymbol{C}_{RGB}$ containing red, green, and blue camera responses. For analysing colours independent of the scale of intensity, it is convenient to transform a colour vector $\boldsymbol{C}_{RGB}$ to its corresponding *chromaticity* $\boldsymbol{c}_{rgb}$. This is done by normalising the colour vector elements (R, G, B) with their first norm $(N = R + G + B)$:

$$c_r = \frac{R}{N}, \;\; c_g = \frac{G}{N}, \;\; c_b = \frac{B}{N} \tag{1}$$

2.2 Statistical modelling of skin chromaticities

The rg-chromaticities $\boldsymbol{c}_{rg,i}$ of a skin area with $i = 1 \ldots n$ pixels may simply be modelled by its mean value $\boldsymbol{\mu}_{rg}$ and covariance matrix $\boldsymbol{S}_{rg}$.

In figure 1 skin chromaticities of one subject are shown which were taken under four different CCTs. The asterisks $(*)$ are the mean values of the respective distributions. The dashed lines are 85% confidence ellipses calculated by the Mahalanobis-distance with a χ^2 for two degrees of freedom.

$$\chi^2 \geq (\boldsymbol{c}_{rg,i} - \boldsymbol{\mu}_{rg})^T \boldsymbol{S}_{rg}^{-1} (\boldsymbol{c}_{rg,i} - \boldsymbol{\mu}_{rg}) \tag{2}$$

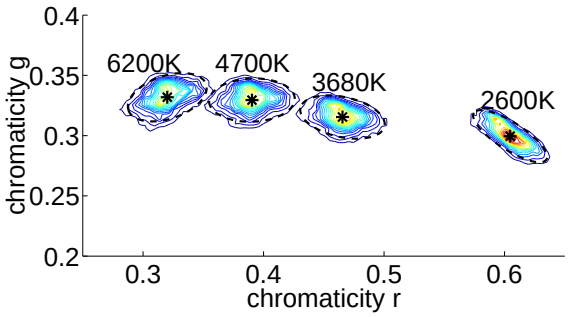

Fig. 1. Skin chromaticity distributions of an Asian subject under four different CCTs.

2.3 Physics-based modelling of skin chromaticities

The chromaticities of the body reflections of human skin can be modelled using reflectance curves of skin, the spectral sensitivities of the camera, and the spectral composition of the light source [7, 12]. The RGB values $\boldsymbol{C}_{RGB}$ are obtained by spectral integration and the corresponding chromaticities $\boldsymbol{c}_{rgb}$ by equation 1.

The reflectance curves of human skin may be modelled as a function of the melanin concentration if the epidermis and the blood content in the dermis [7, 12]. Negro skin has a high melanin concentration whereas Caucasian skin has a low melanin concentration. The melanin concentration for one subject is not constant but has a certain range of variation in space. The lower and upper limits of the blood content $b_{\min}$ and $b_{\max}$ are rather constant for all ethnic groups. The skin chromaticity distribution for an individual is approximated by a minimum and a maximum melanin concentration $m_{\min}$ and $m_{\max}$. In other words it is approximated by the area between the four chromaticities modelled with the reflectance curves using: $b_{\min}$ and $m_{\min}$; $b_{\min}$ and $m_{\max}$; $b_{\max}$ and $m_{\min}$; $b_{\max}$ and $m_{\max}$.

The chromaticities of general purpose light sources, e.g. daylight and fluorescent light, have only a small deviation from the Blackbody radiator with the same CCT.

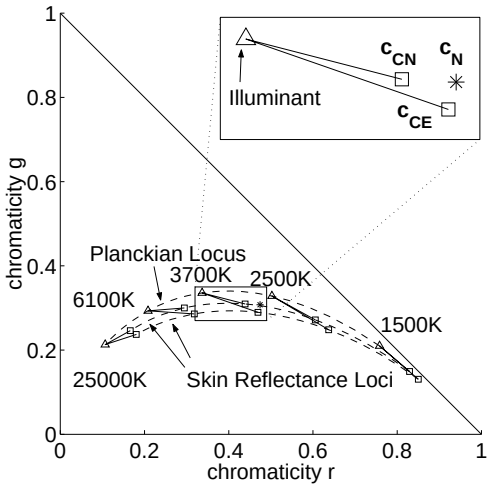

Fig. 2. Modelled illuminant- and body reflection-chromaticities of skin for a number of Planckian illuminants, from [12].

Finlayson and Schaefer [4] measured 172 light sources, including daylights and fluorescent. They report that the illuminant chromaticities fall on a long thin band in the chromaticity plane which is very close to the Planckian locus of Blackbody radiators. Light sources will in the following be approximated by Blackbody radiators of the same CCT as the light source, which was successfully done in [12].

Figure 2 shows the rg-chromaticities of a number of Blackbody illuminants ranging from 1500K to 25000K (Planckian locus), plotted as triangles ($\triangle$). The figure is modelled for a Blackbody radiator with a CCT=3700K as canonical light source (white balance source). Two body reflection chromaticities, plotted as squares ($\square$), are modelled for each illuminant using the Caucasian reflectance curves for low and for high blood content, called $\boldsymbol{c}_{CN}$ and $\boldsymbol{c}_{CE}$. The solid lines indicate the corresponding illuminant and body chromaticities. The asterisks ($*$) shows the body reflectance chromaticity for Negro skin $\boldsymbol{c}_N$, i.e. with high melanin concentration.

3 Adapting statistical models to changing illumination

The method proposed in this paper uses physics-based knowledge to estimate how the statistical model will change as the illumination changes. A necessary condition for this is to find how these types of models can be related.

Firstly we consider how the distributions of the skin chromaticities change as illustrated in figure 1. They change position along the skin locus and the major and minor axes (eigenspace of the covariance matrix) change in orientation and aspect ratio.

The physics-based model describes an expected area for the chromaticities given the CCT and some biological parameters of the human skin. This area can also be described

by a major and minor axes, hence an eigenspace, such that if the biological parameters for a given individual are known, the eigenspace can be computed for any CCT.

The question is if the physics-based and the statistical model can be adequately related via their eigenspaces described above. If so, an initial 'calibration' of the two models can be established from a reference image, and parameters of the statistical model can then be estimated for any CCT.

To demonstrate and test if the eigenspace description may have practical use, we present a procedure including two steps, initialisation for calibration (section 3.1), and then estimation of eigenspaces for arbitrary CCTs (section 3.2).

3.1 Initialisation

Given is a reference image with a pre-segmented area of facial skin, and with known CCT of the illumination used. The corresponding skin chromaticities form a distribution M_{rg}. The eigenvectors $\boldsymbol{v}_M$ and -values $\boldsymbol{\lambda}_M$ of M_{rg} are calculated.

The next step is to find the physics-based model, which through maximum and minimum melanin and blood parameters defines an area with an eigenspace $\boldsymbol{v}_E$, $\boldsymbol{\lambda}_E$ that match $\boldsymbol{v}_M$, $\boldsymbol{\lambda}_M$.

An iterative procedure scans through relevant combinations of these parameters testing for the orientation deviation between $\boldsymbol{v}_M$ and $\boldsymbol{v}_E$. Among those configurations with a deviation below some threshold, here $0.8°$, we chose the best fitting aspect ration defined as

$$\min\left(\left|\frac{\lambda_{M,1}}{\lambda_{M,2}} - \frac{\lambda_{E,1}}{\lambda_{E,2}}\right|\right) \quad . \tag{3}$$

From the selected configuration we get the melanin parameters as the major result of the initialisation. To compensate for coarseness of our model (e.g. blood content independent of individuals/ethnic groups) we also introduce a diagonal matrix $\boldsymbol{k}$ relating the modelled and measured eigenvalues,

$$k_{11} = \frac{\lambda_{M,1}}{\lambda_{E,1}} \qquad k_{22} = \frac{\lambda_{M,2}}{\lambda_{E,2}} \quad . \tag{4}$$

3.2 Estimating eigenspaces for arbitrary CCTs

Using the parameters estimated from the reference image we get four skin reflectance curves. Together with the camera sensitivities, and the Blackbody spectrum of the CCT in question an expected area for the skin chromaticities can be computed. Hence, an eigenspace ($\boldsymbol{v}_E$, $\boldsymbol{\lambda}_E$) can be estimated for any CCT. This eigenspace can be used to estimate the expected covariance matrix $\boldsymbol{S}_E$ of the measurements for the CCT in question,

$$\boldsymbol{S}_E = \boldsymbol{v}_E \cdot \boldsymbol{\lambda}_E \cdot \boldsymbol{k} \cdot \boldsymbol{v}_E^{-1} \quad . \tag{5}$$

4 Test results

The method is tested using one reference CCT and 3 test CCTs. Images of 8 subjects having different ethnic backgrounds (China, Iran, Cameroun, Latvia, Greece, Spain, Denmark, and India) were captured, so that altogether 32 images were used. For each reference there are 3 test images used to compare the estimated statistical model with the measurements.

4.1 Image acquisition

The images were captured with a JAI CV-M90 3CCD camera. Automatic gain control and automatic white balancing are switched off and gamma correction is set to one. The lens opening and the shutter speed are manually adjusted to make use of the dynamic range of the camera. The light sources are fluorescent lamps (Philips TLD 927, 940, 950, 965) with CCTs of 2600, 3680, 4700, and 6200K, respectively. The spectra of these lamps are provided from Philips and additionally measured with a J&M TIDAS spectrometer. The measured spectra were used to calculate the CCTs. The illuminance on the faces is approximately 2000 lux as it is recommended by JAI. There are no pronounced highlights on the faces. The number of pixels of a face is between 5000 and 15000. The camera is white balanced to the 3680K lamp. An example image is shown in figure 3. The images were hand segmented into facial skin areas, which are used in the following for evaluating the method.

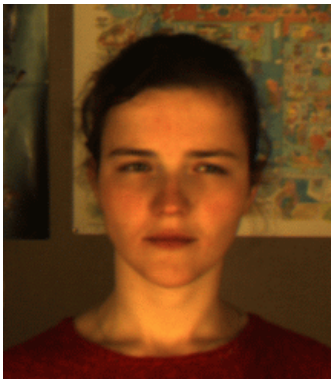

Fig. 3. Example Image of a face. Examples of colour images are available at http://www.cvmt.auc.dk/˜mst/ras.html

4.2 Modelling skin by mean and covariance

The measured skin distributions were tested for normality in the r and g chromaticity dimensions using a modified Kolmogorov-Smirnov test, which rejected the null hypothesis at a significance level of $\alpha = 0.2$. Furthermore, empirical quantile-quantile plots

were done, see figure 4 for some examples. It can be seen that non of the distributions is normal distributed.

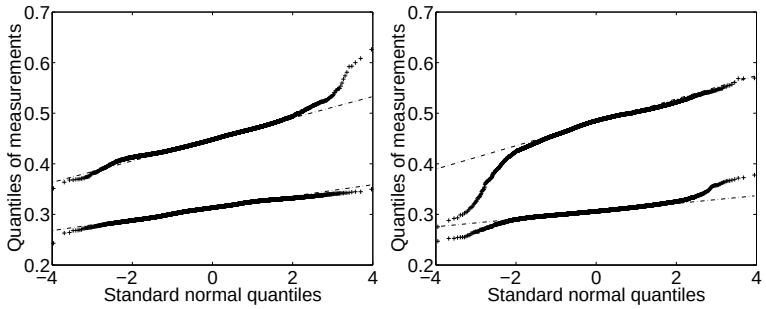

Fig. 4. Two examples of quantile-quantile plots of measured data (chromaticities) versus standard normal. The upper curves are in the red, the lower along the green chromaticity dimension.

However, the distributions might be approximated by their covariance matrix. The true positives inside the 85% confidence ellipses (shown in figure 1) are counted using equation 2 and the percentage is shown in figure 5 (*left*). The average is 88%.

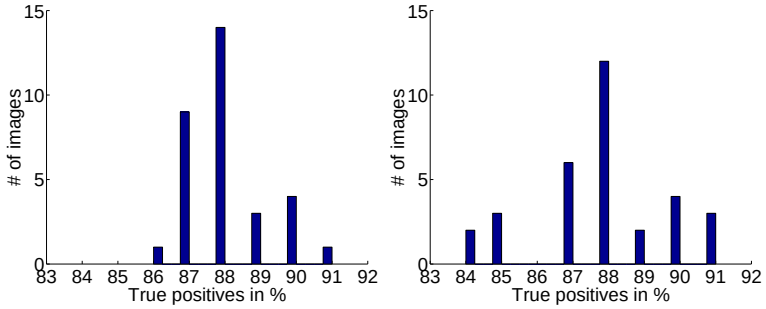

Fig. 5. Percentage of true positives using 85% confidence ellipses with the measured covariance (*left*) and modelled covariance (*right*).

4.3 Physics based modelling of eigenspaces

The images taken at a CCT=3680K were used to estimate the parameters of the method as it was described in section 3.1. Then the covariance matrices for the CCT=2600, 4700, and 6200K, respectively, were modelled as described in section 3.2.

Figures 6 and 7 show example results with measured and modelled confidence ellipses. To allow comparison of the measured and the modelled covariances, their eigenvectors and eigenvalues are given in table 1. The average deviation between measured

and modelled orientation is about $4°$ and maximal $9.8°$. The average deviation between the lengths of the vectors is 6 %.

Figure 5 (*right*) shows the true positives inside the 85% modelled confidence ellipses. The average is 87%. The area of the modelled ellipses is in average not bigger than the area of the measured ones.

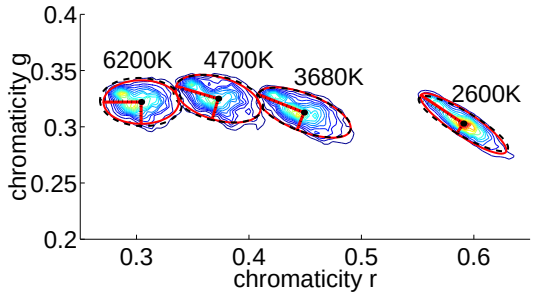

Fig. 6. Skin chromaticities for a Caucasian subject under four different CCTs. The solid lines are the measured ellipses and the dashed lines are the modelled ellipses.

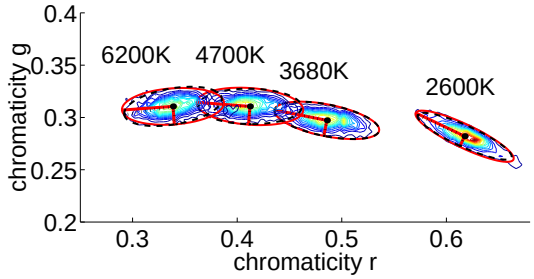

Fig. 7. Skin chromaticities for a Negro subject under four different CCTs. The solid lines are the measured ellipses and the dashed lines are the modelled ellipses.

5 Discussion

The proposed method aims to estimate the eigenspaces of an individual's skin chromaticity distribution for illuminations with arbitrary CCTs. We have tested this for 8 very different skin types, and with 3 illuminations for each. The average deviation between measured and estimated orientations of the eigenspaces is about $4°$ and the maximum about $10°$. The average deviation between the lengths of the vectors is about 6%.

Table 1. Deviations between measured and modelled data for 3 test CCTs and 8 subjects. The angle error is the angle between the eigenvectors. The eigenvalue error is the deviation between the square roots of the measured and modelled eigenvalues in %. The numbers in brackets are from the initialisation data.

illumi-	angle error		eigenvalue errors in %			
nation	in degree		first		second	
CCT	mean	max	mean	max	mean	max
2600	2.7	7.5	4.8	8.7	6.1	15.7
(3680)	(0.3)	(0.8)	(0.0)	(0.0)	(0.0)	(0.0)
4700	3.7	5.2	3.8	8.1	6.0	16.5
6200	4.5	9.8	6.5	12.2	8.0	16.9

The deviations are due to measurement noise during the image acquisition and the approximations in modelling. The light spectra are approximated by blackbody radiators which very coarsely match spectra of fluorescent lamps and, thus, introduce a small error, especially in the green chromaticity [12]. Furthermore, the assumption that the range of blood concentrations is constant is convenient, but may as well contribute to the error.

If 85% confidence ellipses are used to approximate the distributions, the number of true positives of the measured pixels is in average 88% for the measured covariances and 87% for the modelled covariances. Hence, the skin distributions may for practical purposes be approximated with their covariances, which was also reported in [15].

The method requires knowledge about an (approximate) CCT of the current illumination. This may be estimated from the mean value of the measured skin chromaticity distribution or, more accurately, by methods proposed in [11, 13].

Figures 1, 6, and 7 show the distributions of an Asian, a Caucasian, and a Negro subject. A model for each individual is estimated. It can be seen that they differ from each other. If this difference is significant it might be useful to track and distinguish the skin of multiple faces in an image.

6 Conclusions

This paper demonstrated that the eigenspace of a two dimensional skin chromaticity distribution can be modelled for different illuminations using a physics-based approach. The average orientation error is about $4°$, the average deviation between the lengths of the vectors is about 6 %.

The performance seems to be within a useful range to allow significant support of a statistical based approach by improving the robustness of the colour based segmentation in face tracking systems.

In future work the skin reflection model could be improved by estimating also the blood concentration as a parameter, and by that possibly eliminate the correction matrix k.

Acknowledgements

We would like to thank all the persons who participated in the image acquisition. The work was funded in part by the Danish National Research Council project “The Staging of Virtual 3-D Spaces”.

References

1. Hans J. Andersen and Erik Granum. Classifying illumination condition from two light sources by colour histogram assessment. *Journal of the Optical Society of America A*, 17(4):667–676, April 2000.
2. R. R. Anderson, J. Hu, and J. A. Parrish. Optical radiation transfer in the human skin and applications in *in vivo* remittance spectroscopy. In R. Marks and P. A. Payne, editors, *Bioengineering and the Skin*, chapter 28, pages 253–265. MTP Press Limited, 1981.
3. James L. Crowley and Francois Berard. Multi-modal tracking of faces for video communications. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 640–645, Puerto Rico, June 1997.
4. Graham D. Finlayson and Gerald Schaefer. Solving for colour constancy using a constrained dichromatic reflection model. *International Journal of Computer Vision*, 42(3):127–144, May 2001.
5. Brian Funt, Kobus Barnard, and Lindsay Martin. Is machine colour constancy good enough. In Hans Burkhardt and Bernd Neumann, editors, *5th European Conference on Computer Vision*, pages 445–459, Freiburg, Germany, June 1998.
6. Stephen J. McKenna, Yogesh Raja, and Shaogang Gong. Tracking colour objects using adaptive mixture models. *Image and Vision Computing*, 17(3–4):225–231, 1999.
7. Tomoyuki Ohtsuki and Glenn Healey. Using color and geometric models for extracting facial features. *Journal of Imaging Science and Technology*, 42(6):554–561, December 1998.
8. Steven A. Shafer. Using color to separate reflection components. *COLOR Research and Application*, 10(4):210–218, 1985.
9. Maricor Soriano, Birgitta Martinkauppi, Sami Huovinen, and Mika Laaksonen. Skin detection in video under changing illumination conditions. In *International Conference on Pattern Recognition*, volume 1, pages 839–842, Barcelona, Spain, September 2000.
10. Martin Spengler and Bernt Schiele. Towards robust multi-cue integration for visual tracking. In B. Schiele and G. Sagerer, editors, *International Workshop on Computer Vision Systems*, volume 2095 of *LNCS*, pages 93–106, Vancouver, Canada, July 2001.
11. Moritz Störring, Hans J. Andersen, and Erik Granum. Estimation of the illuminant colour from human skin colour. In *4th IEEE International Conference on Automatic Face- and Gesture-Recognition*, pages 64–69, Grenoble, France, March 2000.
12. Moritz Störring, Hans J. Andersen, and Erik Granum. Physics-based modelling of human skin under mixed illuminants. *Journal of Robotics and Autonomous Systems*, 35(3–4):131–142, June 2001.
13. Moritz Störring, Erik Granum, and Hans J. Andersen. Estimation of the illuminant colour using highlights from human skin. In *First International Conference on Color in Graphics and Image Processing*, pages 45–50, Saint-Etienne, France, October 2000.
14. Jochen Triesch and Christoph von der Malsburg. Self-organized integration of adaptive visual cues for face tracking. In *4th IEEE International Conference on Automatic Face- and Gesture-Recognition*, pages 102–107, Grenoble, France, March 2000.
15. Jie Yang, Weier Lu, and Alex Waibel. Skin-color modeling and adaptation. In Roland Chin and Ting-Chuen Pong, editors, *3rd Asian Conference on Computer Vision*, volume 1352 of *LNCS*, pages 687–694, Hong Kong, China, January 1998.