

Skin chromaticity gamuts for illumination recovery

Stuart Crichton^{1,2}, Jonas Pichat³, Michal Mackiewicz⁴, Gui-Yun Tian², Anya Hurlbert¹

¹Institute of Neuroscience, Newcastle University, UK; ²School of Electrical, Electronic and Computer Engineering, Newcastle University, UK; ³ENSEIRB Bordeaux, France; ⁴School of Computing Sciences, University of East Anglia, UK

Abstract

Colour constancy algorithms range from image statistics-based pixel intensity manipulation to gamut-mapping methods, and are generally independent of specific image contents. In previous work, we have demonstrated that natural polychromatic surfaces possess distinct chromatic signatures in cone-contrast space that may be exploited for colour constancy, and that in human vision, colour constancy is improved for such objects. Here we set out to use the specific, recognisable, and ubiquitous content of human skin in colour images to drive a gamut mapping method for colour constancy. We characterise variations in the chromaticity gamut of varying types of, pre recognised, human skin (male, female; Caucasian, African, Asian) under varying illumination. We use a custom-built LED illuminator to produce daylight metamers, and a spectroradiometrically calibrated hyperspectral camera (Specim V10E) to acquire images and create a novel hyperspectral skin image database. We demonstrate that human skin gamuts in cone-contrast space are characterised by a set of features that can be used to differentiate between similar illuminations, whose estimate can then be used to colour correct an image.

Introduction

The mechanisms of colour constancy built into human vision enable object colours to remain roughly constant across changes in illumination. RGB-based camera systems are not natively colour constant, recording different RGB triplets for the same object under different illuminations, and it therefore remains a goal for colour image processing to develop and implement constancy algorithms that mimic human perception in correcting for changes in illumination. Over the past 40 years a number of approaches have been used to tackle the challenge of colour constancy in both computer vision and photography, ranging from image statistics-based pixel intensity manipulation to gamut mapping methods [3]. Gamut mapping methods typically use the entire image regardless of image content, while “white-balance” methods typically use a reference surface selected on the basis only of its pixel intensity values. “Max-RGB” or “normalise-to-white” algorithms, for example, assume that the pixels of highest intensity in each channel represent the highest possible reflectance in that channel, and therefore, would be white under a white illumination. In-camera “white-balance” methods may also require the user to set an arbitrary colour correlated temperature (CCT) value (at times annotated with more natural language labels, such as ‘cloudy’). If the highest-intensity pixels do not correspond to white reflectances, or the actual illumination does not match the selected correlated colour temperature, the algorithm will not adequately reproduce human perception of the scene. The inadequacies in these methods are often most noticeable in the colours of the most familiar objects, such as people and their skin. Here we propose to use the information contained in the images of these specific familiar objects to estimate the illumination incident on them, and, under the single source assumption, on the entire scene. We use a

modified gamut-mapping algorithm that operates on the chromaticity gamut of only a selected portion of the image.

Forsyth’s original gamut-mapping algorithm [4] assumes that the canonical gamut of a given scene is known. The algorithm takes the chromaticity gamut of the scene under an unknown illumination and calculates the most likely mapping that transforms its convex hull into the convex hull of the canonical gamut. The selected transform is the one that satisfies all of the possible point to point transforms between the two gamuts. The approach is illustrated in Figure 1 below.

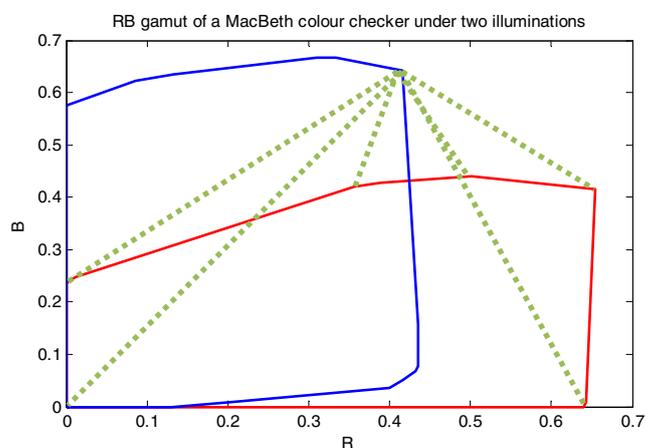


Figure 1. Visualisation of Forsyth’s algorithm

For scenes that contain a complete sampling of all possible natural surfaces, the canonical gamut is complete and unchanging, but for scenes that contain biased samplings, the canonical gamut may itself be unknown. Thus, like the grey-world and white-balance algorithms, the original gamut-mapping algorithm falters when the image contents fail to match the assumptions. We suggest that knowledge of the actual image contents may therefore be useful in constraining gamut-mapping methods.

Tominaga and Wandell [5] proposed a further development in the form of a gamut correlation algorithm for a specific set of natural images [6]. The algorithm compares incoming gamuts with memory gamuts (both in the RB color space) stored in an image database, choosing as a match the stored gamut which has the highest correlation to that of the test image. Figure 2, along with Equations 1-4, outline the correlation algorithm.

The restrictions on the incoming images, and on the database itself eliminates the problem of mismatch between the canonical collection of surfaces and the image contents. For general use, though, gamut-mapping methods will nonetheless always face the challenge of constructing a database that adequately accounts for all possible materials, configurations and illuminations encountered in unbounded images. One possible solution would be to focus the database on one specific feature class, or object type, and build it comprehensively.

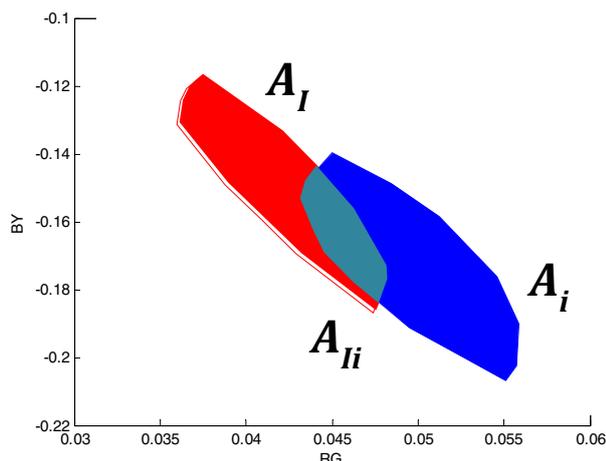


Figure 2. Depiction of the Tominaga-Wandell correlation algorithm

$$\text{Correlation;} \\ r_i = (A_{Ii} / (\sqrt{A_I A_i})) \quad (1)$$

$$\text{Where;} \\ A_{Ii} = A_I \cap A_i \quad (2)$$

$$A_I = \text{Area of test image gamut} \quad (3)$$

$$A_i = \text{Area of illuminant reference gamut} \quad (4)$$

Here we propose human skin as the object of focus, due to its recognisability and near ubiquitous nature in photographic images.

Previous work in computer vision has extensively examined skin colour as an identifying feature for face detection and tracking [7,8] or for image classification [9], and in this context, there is a driving goal to find illumination-independent representations of skin colour [e.g. 8]. Here, we do not propose a skin detection algorithm per se, but rather assume that skin detection has already been performed, for example, by a face detection algorithm, and that once identified, the characteristics of the observed skin colour may be used to drive a colour constancy algorithm for the entire image. A similar approach has been attempted in [10] and [11]. However, as with most previous work, these approaches are based in RGB or RGB derived spaces, e.g. HSV space [11] for scene classification. The use of general images in uncontrolled environments here leads to real problems in discovering the ground truth gamuts for the stated objects (skin, vegetation and sky-sea). Here we propose to analyse image data in a physiological cone-contrast space, reasoning that this space is likely to have been naturally optimised to encode skin colour variations. We also hypothesise that the human visual system may use skin colour to drive colour constancy, since rapid object-recognition processes are known to occur in human vision, which could mediate the initial step of identification and segmentation of skin areas [12].

We specifically explore the idea that the distribution of chromaticities, in a known skin sample, may provide sufficient information to estimate the unknown illumination on the scene, as in [13]. We start from our previous observation [1] that the chromaticity distributions of natural polychromatic surfaces such as fruits and vegetables form regular signatures in the

physiological cone-contrast space. These signatures exhibit invariant features (e.g., hue angle) under changing illumination, which may mediate their constancy to the human visual system. The signatures also vary in predictable ways under changing illumination, suggesting that signatures of familiar objects may be exploited to recover information about an unknown illumination, and thereby aid colour constancy for unfamiliar objects under the same illumination. Empirically, we have shown that colour constancy is indeed improved for naturally chromatically variegated familiar objects, compared with chromatically uniform familiar or unfamiliar objects [2]. Here we observe that, in human vision, one's own skin provides an omnipresent, familiar reference surface. We then aim to determine whether features of the skin chromaticity distributions, as represented in the physiological colour space used by the human visual system, contain sufficient information to estimate unknown illuminations. The idea that human skin may provide a reference surface for chromatic adaptation, and thereby mediate colour constancy, is not new [14]. Previous implementations of the idea have relied, though, on mean skin chromaticities only, effectively using the mean chromaticity to provide the white point against which other image chromaticities are referenced. Here we explore whether the additional information contained in the inherent spread of chromaticities of bare skin provides additional support for colour constancy.

Outline of the method

As previously mentioned, we examine the feasibility of illumination estimation from skin chromaticity gamuts on the assumption that skin regions have already been identified. The first step is to form a reference set of skin chromaticity gamuts. To obtain accurate representations of the chromaticity variations in cone-contrast space, and to obtain ground-truth information for both reflectance and illumination spatial variations, we use hyperspectral imaging to create a hand image database. We then characterize the chromaticity gamuts of each skin sample in terms of a set of features, creating a reference feature look up table (FLUT). Lastly, we devise and test two illumination estimation methods, explained in the cost function algorithm section, with a set of test images. These test images include some used to create the reference FLUT, as well as 'blind' images taken under illuminations not used in the creation of the FLUT, in order to test whether our set of features can be used to closely estimate and match (in terms of CCT) the unknown illumination in a scene to one of our known illuminants. The workflow is as follows:

1. The first stage is the creation of the FLUT, consisting of the aforementioned chromaticity gamut features across different illuminations for a range of skin types. As a single example, this would entail saving the differing gamut areas for the same patch of skin across a range of illuminations.
2. A new image, under an unknown illumination, is then received. The skin type is identified, and a section of skin is processed in order to produce its chromaticity gamut.
3. The features from this skin patch under unknown illumination are then compared to the features stored within the feature LUT, using the

cost functions which are described in the following sections.

4. The reference illumination with the lowest feature error between itself and the incoming gamut is then selected from the database as the estimate.

In the following sections, we elaborate on each of these steps in the workflow.

Hyperspectral Hand Image Database

To obtain accurate records of the variation in colour across a single skin sample we use a photometrically calibrated hyperspectral camera (Specim V10E) to record image irradiance spectra at 2-nm intervals at each pixel. We have initially acquired hyperspectral images of hands from 8 different human subjects, spanning a range of geographical origins (see Table 1), and are continuing to expand the database with additional subjects. For each subject, we have acquired hyperspectral images of the hand under 39 distinct illuminations, produced by a custom-built LED illuminator that provides diffuse light over a large surface area (Mackiewicz et al., CGIV 2012, accepted). The future potential regarding the use of hyperspectral imaging in colour constancy research is shown by [15]. These illuminations range from CIE D type to F type and into other highly chromatic lights. Figure 3 below illustrates the locations of the white-surface chromaticities of the illuminations used in this database on the CIE 1931 xy chromaticity diagram. In this report, we compare results for skin gamuts from ten CIE D type of illuminants, ranging from D40 (yellowish) to D250 (bluish). This range ensures that each illuminant has a large degree of similarity to its neighbours, posing a particular challenge for a method to differentiate between possible illuminations, whilst also encompassing a noticeable change in chromaticity.

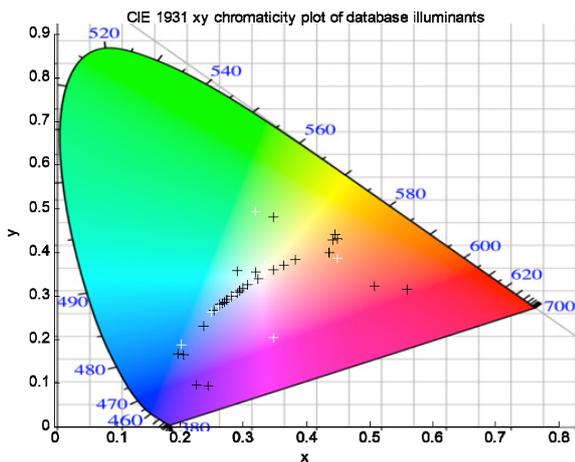


Figure 3. The database illuminations in CIE 1931 xy space

Table 1: Subject skin types.

Subject	Geographical/racial origin
1	Caucasian
2	Caucasian
3	Tanned Caucasian
4	Indian sub continent
5	North African
6	East Asian
7	West African

The choice to base this work upon hyperspectral imaging instead of traditional RGB imaging is driven by the level of information which is captured by a hyperspectral imaging system in comparison to that captured by a traditional digital camera with a Bayer or Foveon RGB sensor. As stated earlier, the underlying inspiration of this work is the regularity of chromatic signatures of natural objects in physiological colour space. In order to reconstruct the physiological data accurately, we need to maximise the amount of spectral information captured before conversion into cone-contrast space. It is also possible to characterise and calibrate conventional RGB cameras with 8-bits per channel in order to recover estimates of cone-contrast coordinates, but the estimates are less accurate, and the ground-truth information about surface reflectance functions and illumination power spectrum is lost.

Skin chromaticity gamuts in cone-contrast space

From the hyperspectral image database, we extract from each image a rectangular patch of human skin (here, 25x25 pixels, although the exact size of the patch does not affect the results). To obtain the skin chromaticity gamut, we convert the irradiance spectrum at each pixel into CIE XYZ tristimulus values and thence into cone-contrast coordinates, using the Eskew et al. formula [16] and the method described in [1]. We then find the convex hull of the set of surface chromaticities in cone-contrast space. In the cone-contrast equiluminant hue plane, the two axes plot the cone-opponent contrasts with respect to the neutral adaptation point (0,0), i.e. the RG (“red-green”) axis plots the contrast of the (L-M) cone-opponent excitation with respect to the (L-M) excitation of the neutral adaptation point, and the BY (“blue-yellow”) axis plots the contrast of the (S-(L+M)) cone-opponent excitation with respect to the (S-(L+M)) excitation of the adaptation point. In the algorithm, we compute the average chromaticity of a known neutral patch in the image to provide the neutral adaptation point, although an equally plausible adaptation point would be the average of the entire scene. The human visual system is known to adjust its neutral adaptation point dynamically in response to a space-time averaging of light signals over the scene. Figure 4 illustrates two sample skin chromaticity gamuts in cone-contrast space obtained from the same human hand in the same scene, under two distinct illuminations, together with the irradiance spectra from the same hand location in the two images.

Note that the chromaticity gamuts form regular clusters in the cone-contrast space, occupying similar territory at similar orientations under two very different illuminations. A similar regularity in LMS cone space is predicted by a physics-based model of human skin colouring [17].

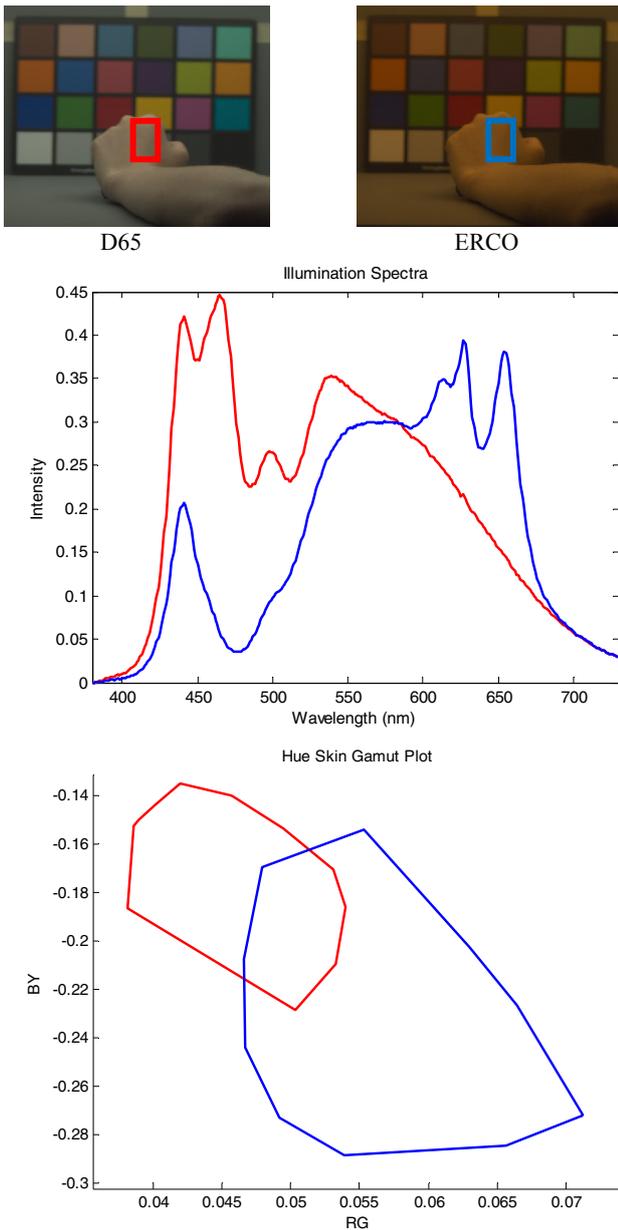


Figure 4. The variation of gamut characteristics across two illuminations. (a) Images of the hand patches under two illuminations, (b) the spectral power distributions of the two illuminations and (c) the gamuts of the same patch under two illuminations.

We analyse and describe the characteristics of the chromaticity gamuts using the following measurements: angle and radius of the centre of mass and extreme points of the convex hull, and planar area of the convex hull. Although here we focus on the RG-BY hue plane only, in the full description, we compute and store these measurements in each of the three planes of physiological colour space: RG-BY (hue), RG-Lum and BY-Lum, where Lum is the luminance contrast ($L+M$) relative to the neutral adaptation point. Figure 5 illustrates these features on an example gamut:

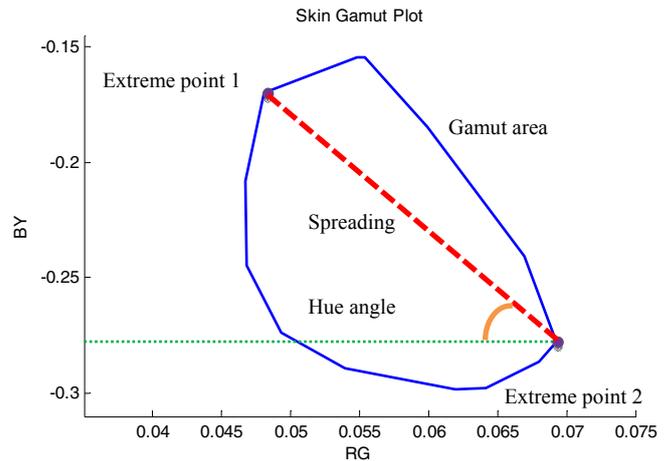


Figure 5. Characteristics of the chromaticity gamut.

Variations in skin chromaticity gamuts with illumination

Figure 6 illustrates, for two subjects, the set of chromaticity gamuts for seven selected illuminations (D40, D50, D80, D100, D150, D200 and D250) in the RG-BY plane. The results

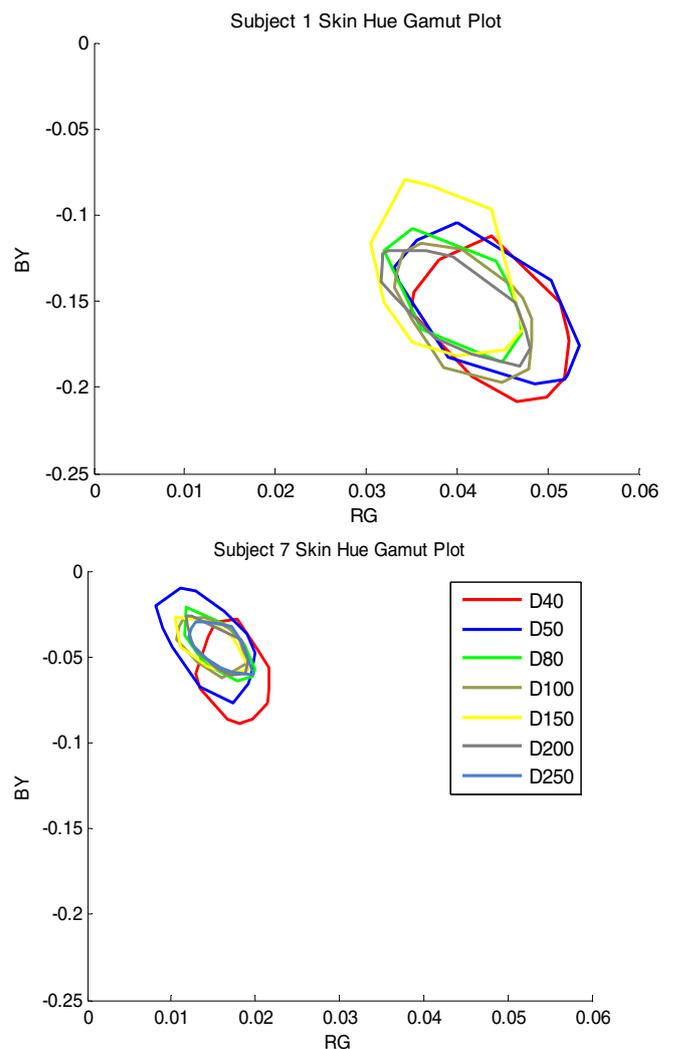


Figure 6. Skin gamuts for two subjects across 10 daylight illuminations (Legend is correct for both plots).

presented here draw on a greater set of 10 gamuts for each of the 7 subjects in Table 1 (with missing values of D150 for Subject 6, and D80, D120 and D150 for Subject 7). On visual inspection of Figure 6, it is clear that despite the difference in mean chromaticity of the two skin samples, for both subjects the hue angle is similar and varies little across illumination, while the gamut area varies systematically with illumination. Quantitatively, although the range of variation in hue angle across illuminations is small (76-87 degrees), it is significant ($F(9,69) = 2.22, p < 0.05$). The variation in hue area is highly significant ($F(9,69) = 5.25, p < 0.0005$), as illustrated in Figure 7. From the analysis of these collected features the features which give the greatest cue to illumination (i.e. those which change in the most visible/ predictable manner) were used to create the cost functions laid out within the next section.

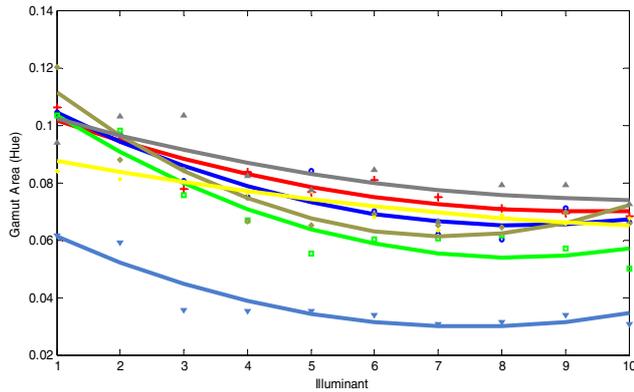


Figure 7. Variation of gamut area for each subject across the range of illuminations. Each colour represents a distinct subject.

Cost function algorithm for illumination estimation

From the analysis of skin gamut features across a range of illuminations the best features for estimating the illumination were retrieved. Thus in order to estimate illumination we employ a cost function that calculates, for each candidate (A_i) and incoming gamut pair (A_1), the distance between their respective centres of mass (CoM) and extreme points ($ExtP1$, $ExtP2$), the fractional intersectional area of the two gamuts ($IntArea$), and the area of the uniquely non-overlapping gamut regions ($UniqueA$, $UniqueB$), as follows from Equation 2;

$$A_{ii} = A_1 \cap A_i \quad (2)$$

$$UniqueA = A_1 - A_{ii} \quad (5)$$

$$UniqueB = A_i - A_{ii} \quad (6)$$

$$\Delta(ExtP1) = \sqrt{((P1x_1 - P1x_i)^2 + (P1y_1 - P1y_i)^2)} \quad (7)$$

$$\Delta(ExtP2) = \sqrt{((P2x_1 - P2x_i)^2 + (P2y_1 - P2y_i)^2)} \quad (8)$$

$$\Delta(COM) = \sqrt{((RG_1 - RG_i)^2 + (BY_1 - BY_i)^2)} \quad (9)$$

$$Original\ Cost\ Function = 1 - A_{ii} + \Delta(ExtP1) + \Delta(ExtP2) + \Delta(COM) + UniqueA + UniqueB \quad (10)$$

$$COM\ Cost\ Function = \Delta(COM) \quad (11)$$

We select the memory gamut that minimises the above cost function.

Table 2: The chromaticity dE between the actual and estimated illuminations.

Subject	Test	Original Cost Function Chroma dE	CoM Cost Function Chroma dE
Subject 1	D40	0.000	0.000
	D50	0.000	0.031
	D60	0.000	0.088
	D80	0.000	0.049
	D100	0.022	0.000
	D120	0.035	0.013
	D140	0.085	0.000
	D150	0.000	0.000
	D200	0.061	0.000
	D250	0.000	0.032
	D65	0.028	0.028
Subject 4	D40	0.000	0.074
	D50	0.000	0.043
	D60	0.074	0.000
	D80	0.000	0.049
	D100	0.000	0.027
	D120	0.000	0.013
	D140	0.000	0.016
	D150	0.000	0.049
	D200	0.000	0.000
	D250	0.000	0.022
	D65	0.043	0.086

Table 2 illustrates the results of the cost function for an incoming image under 11 different illuminations (the original 10 plus the additional unknown illumination, D65) using a distinct surface patch on the human hand from those used to construct the database. To determine whether the additional information contained in the distribution improves colour constancy, we compare the performance of this cost function with a cost function that calculates only the distance between the two centres of the mass of the incoming gamut and the candidate gamut. Over all subjects, and all illuminations, the error in illumination estimation is significantly increased for the centre-of-mass-matching algorithm compared with the gamut-matching algorithm (Table 2 depicts results for two individual subjects; for subject 4, the difference in estimation error for the two methods is significant on its own, $F(1,21) = 4.4; p < 0.05$).

Conclusion and future work

We have demonstrated that the surface chromaticity gamut of human skin possesses a typical, predictable structure in physiological cone-contrast space across a range of different skin types. Skin chromaticities form an oriented cluster, the hue angle (in the RG-BY chromaticity plane) of which remains roughly similar across skin types, but varies in a similar way across illuminations. The shape and spread of the gamut varies in a similar way across a range of daylight illuminations for all skin types we studied here. Using a basic cost function minimisation for database searching, we demonstrate that the area and extreme points of the chromaticity gamut in this chromaticity space provide additional information that aids in estimating the illumination, in comparison with the mean chromaticity of the distribution, even after adaptation to the neutral point.

In this work, we also created a hyperspectral image database of skin types under a variety of tunable illuminations which we anticipate will be useful for other applications in colour constancy and object recognition. In future work on the chromaticity gamut-matching algorithm, we will examine in more detail the variation in the gamut characteristics with respect to location of the skin surface patch, remove the dependence on the pre-identification of skin type, and test the estimations of illumination by colour-correcting images taken under unknown illuminations.

Acknowledgements

We gratefully acknowledge the support of the EPSRC (project grant EP/H022325/1; Newcastle University EPSRC DTA) and the assistance of Dr Angela Owen and Bradley Pearce.

References

- [1] Y Ling, M Vurro, A C Hurlbert, "Surface chromaticity distributions of natural objects under changing illumination", Proceedings of the 4th European Conference on Colour in Graphics, Imaging and Vision (CGIV), pg. 263-267, (2008).
- [2] Vurro, M, "The role of chromatic texture and 3D shape in colour discrimination, memory colour and colour constancy of natural objects", PhD Thesis, Newcastle University, (2010).
- [3] Hurlbert, A.C., Computational Models of Colour Constancy. In "Perceptual Constancy: Why things look as they do", edited by V. Walsh and J. Kulikowski, Cambridge University Press, pg. 283-322, (1998).
- [4] D. Forsyth, "A novel algorithm for color constancy", Int. J. of Comput. Vis., vol. 5, no. 1, pg. 5-36, (1990).
- [5] S.Tominaga, B.A. Wandell, "Natural scene-illuminant estimation using the sensor correlation", Proceedings of the IEEE 90(1), pg. 42-56, (2002).
- [6] S. Tominaga: Natural Image Database and Its Use for Scene Illuminant Estimation, Journal of Electronic Imaging, Vol. 11, No. 4, pg. 434-444, (Oct. 2002).
- [7] P. Kakumanu, S. Makrogiannis, N. Bourbakis, "A survey of skin-color modeling and detection methods", Pattern Recognition, Volume 40, Issue 3, pg. 1106-1122, (March 2007).
- [8] J. Yang, W. Lu, and A. Waibel., "Skin-color modelling and adaptation". In Proc. ACCV, pg. 687-694, (1998).
- [9] M.J. Jones and J.M. Rehg, "Statistical Color Models with Application to Skin Detection," Proc. IEEE Conf. Computer Vision and Pattern Recognition, vol. 1, pg. 274-280, (1999).
- [10] S. Battiato, A. Bosco, A. Castorina, G. Messina, "Automatic Image Enhancement by Content Dependent Exposure Correction", (2004).
- [11] Naccari, F. et al., "Natural scenes classification for color enhancement", IEEE Transactions on Consumer Electronics, 51(1), pg. 234-239, (2005).
- [12] Cadieu C, Kouh M, Pasupathy A, Connor CE, Riesenhuber M, Poggio T. A model of V4 shape selectivity and invariance. J Neurophysiology 98, pg. 1733-1750, (2007).
- [13] Montojo, J., "Face-based chromatic adaptation for tagged photo collections", (2009).
- [14] M.D. Fairchild, "Chromatic adaptation to image displays", TAGA 2, pg. 803-824, (1992).
- [15] Foster, D.H., "Color constancy", Vision Research, 51, pg. 674-700, (2011).
- [16] Eskew R T, McLellan J S, Giulianini F, "Chromatic detection and discrimination" in "Colour Vision: From Genes to Perception", Cambridge, UK: University Press, pg. 345-368, (1999).
- [17] Cotton, SD, Claridge E, "Do all human skin colours lie on a defined surface within LMS space?", Technical Report CSR-96-01, School of Computer Science, The University of Birmingham, (1996).