A Method for Designing and Assessing Sensors for Chromaticity Constancy in High Dynamic Range Scenes

Sivalogeswaran Ratnasingam and Steve Collins, Department of Engineering Science,University of Oxford, Oxford, UK; and Javier Hernández-Andrés, Department of Optics, Sciences Faculty,University of Granada, Granada, Spain

Abstract

The dependence of an object's colour on the illuminant chromaticity makes it difficult to use colour as a reliable cue in machine vision applications, particularly in naturally illuminated high dynamic range scenes. To solve this problem the outputs from four logarithmic sensors with different spectral responses can be used to obtain a two dimensional description of an object's chromaticity that is independent of the illuminant. The spectral responses of these four sensors have been optimised. A simple test of colour separability then suggests that using the data from these sensors it is possible to match the ability of the human visual system to separate similar colours. A comparison of the performance of the proposed system when the reflectance and illuminantion data are both changed suggests that readily available data (Munsell reflectance spectra and CIE standard daylight spectra) can be used to design a generic system to separate colours that are described as matching each other. However, for applications that require discrimination between very well matched colours it may be necessary to use an application specific system designed using data relevant to the application.

Introduction

There are two challenges when imaging naturally illuminated scenes. The challenge that has been the focus of developments in the camera design community is the high dynamic range caused by variations in the intensity of the illuminant. A more subtle but potentially equally important problem is the changes in the illuminant spectrum. This makes it difficult to use colour or chromaticity information in applications such as scene segmentation and object recognition.

Imaging sensors are now available that have an input dynamic range of more than 5 decades or 100dB [1]. Now that high dynamic range imagers are available the more subtle problem can be addressed. One solution to the problems caused by variations in the illuminant spectrum is the colour constancy at a pixel algorithm proposed by Finlayson *et al*.[2]. Although implemented using the output from linear sensors, this algorithm is based upon sensors with a logarithm response to light in a narrow spectral range. Logarithmic imaging sensors have been developed with a wide dynamic range [3]. Other cameras are being developed with organic photodetectors [4]. This creates the opportunity to use photodetectors made from materials, such as cyanines, that have relatively narrow spectral responses. These developments mean that it is timely to investigate the possible advantages of designing an imaging sensor to reliably extract colour or chromaticity information from a scene despite variations in the spectrum of the illuminant.

The effects of changing the peak spectral wavelength of three sensors with narrow spectral responses on the ability to extract an illuminant independent descriptor of colour has been investigated by Romero *et al*.[5]. With these sensors it is possible to create histograms of the colour related descriptor that can be used to segment images based upon colour despite variations in the illuminant spectrum. The problem with any three sensor system is that some very different colours can give very similar descriptors [6]. However, this problem can be avoided by using systems containing four sensors.

In this paper a simple method of extracting illuminant independent features is described in section 2. The impact of optimising the spectral responses of the sensors with different spectral widths is described in section 3. As with any optimisation process there is the possibility that the results obtained depend upon the data used during optimisation. In section 4 the consequences of changing both the illumination and reflectance data is described for the first time.

Illuminant Independent Chromaticity Space

The 'colour constancy at a pixel' algorithm is based upon the assumption that the illuminant can be represented by a black-body with a colour temperature T and that the sensors have logarithmic responses to a single wavelength. Under these conditions the logarithmic response of sensor *i* for an illuminant with intensity *I* is given by

$$\log(R_i) = \log(I) - \frac{K_2}{\lambda_i T} + \log(K_1 \lambda_i^{-5} S_i) \qquad (1)$$

where λ_i is the wavelength at which sensor *i* responds, S_i is the reflectance of the surface being imaged at wavelength λ_i and both $K_1=2hc^2$ and $K_2=hc/k_B$ are constants. The first and second terms in equation (1) depend upon the illuminant intensity and colour temperature. Whilst the last term in this equation depends upon the reflectance of the surface. It is this term that needs to be extracted. To achieve this Finlayson *et al.* [7] have used the responses of four sensors. The intensity dependant term in equation (1) is cancelled by subtracting the response of one logarithmic sensor from those of the other three. The variations in the remaining three variables caused by changes in the illuminant spectrum are then removed by projecting these three variables into a two dimensional space. The result is a two dimensional feature space that represents the 'chromaticity' of the surface being imaged.

Inspired by the results from the algorithm proposed by Finlayson *et al*. a study of the possible advantages of designed a camera with a logarithmic response specifically to extract chromaticity information from a high dynamic range scene has been undertaken here. A possible advantage of designing a camera specifically for this application is that processing the data obtained from the sensor should be as simple as possible. One possible simplification is to avoid the need to find a projection into the feature space. In order to avoid using a projection a method proposed by Marchant and Onyango [8] for three linear sensors has been adapted for use with four logarithmic sensors. To create two illuminant independent features (C_1 and C_2) the responses of four sensors (R_i) with peak spectral responses centered at λ_i are combined using the following equations

$$C_1 = \log(R_2) - \{\alpha \log(R_1) + \beta \log(R_3)\} \qquad (2)$$

$$C_2 = \log(R_3) - \{\gamma \log(R_2) + \delta \log(R_4)\} \qquad (3)$$

These two features are illuminant independent if the first and second terms in equation (1) are both cancelled from equations (2) and (3). Substituting equation (1) into these two equations shows that this is possible if,

$$\alpha+\beta=1,\ \delta+\gamma=1,\quad \frac{1}{\lambda_2}=\frac{\alpha}{\lambda_1}+\frac{\beta}{\lambda_3} \text{ and } \frac{1}{\lambda_3}=\frac{\gamma}{\lambda_2}+\frac{\delta}{\lambda_4}$$

The probable performance of a camera designed to extract chromaticity information using equations (2) and (3) has been assessed by calculating the outputs from a proposed sensor combination using illuminant spectra, surface reflectances and proposed sensor spectral responses, each represented on a 1nm scale. The output of any camera will be a digital rather than an analogue signal. The synthetic sensor responses have therefore been quantised to represent the effect of the analogue to digital converter (ADC) on the output of each sensor. In the model for the ADC the maximum ADC input was equated to the output from the sensor in a particular combination that gives the maximum output with CIE standard daylight illuminant (D65). To represent an n-bit ADC this maximum response is divided by 2^n to determine the sensor output equivalent to the change of one least significant bit of the ADC. The output from each sensor was then represented to an integer number of the output change equivalent to a least significant bit.

The two dimensional chromaticity space formed by C_1 and C_2 using the responses of sensors with Gaussian spectral responses at peak wavelengths in Table (1) is shown in Figure (1). This space has been formed using the reflectance spectra of Munsell colours [9] with CIELab L values between 47.8 and 50.2, illuminated by 14 different CIE standard daylight illuminants with correlated colour temperatures between 5000K and 9000K. The most important characteristics of this chromaticity space are that different colours are well separated and neighbouring colours are similar. A closer inspection of these results shows that there is a residual dependence upon the illuminant which means that each of the different Munsell colours creates a small cluster of points.

Table 1*. The wavelengths of peak spectral responses of the initial choice of sensors. Since all wavelengths could be equally important the sensors have been evenly spaced between 400 nm and 700 nm. Two of the corresponding calculated channel coefficients are also given.*

Sensor ID	1	2	3	4
Peak position (nm)	437.5	512.5	587.5	662.5
Channel coefficients	α=0.5638 and γ=0.5732			

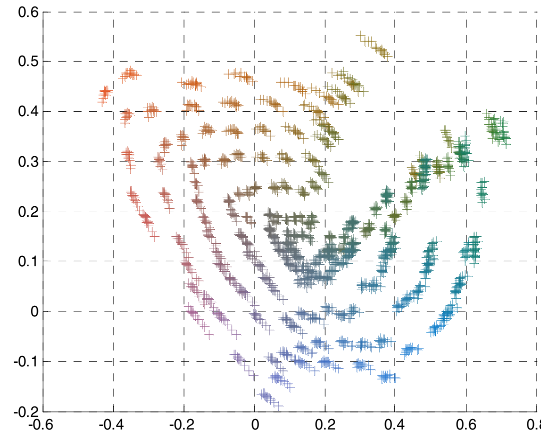

Figure 1*. The two dimensional feature space formed by 80nm wide sensors from 202 Munsell surfaces and 14 CIE standard daylights. Each cross is the colour of the relevant Munsell colour.*

The significance of the residual illumination dependence has been investigated using a simple test of the separability of perceptually similar colours. CIELab is one colour space that has been devised so that distances in the space are proportional to the perceptual differences between colours. Although the method to calculate CIELab co-ordinates has been agreed there are slightly different qualitative descriptions of the similarity of colours separated by different distances in CIELab space, for example, Abrardo *et al*. describe colours that differ by between 1.0 and 3.0 CIELab units as a very good colour match to each other, whilst colours separated by distances between 3.0 and 6.0 units are a good colour match to each other [11]. Since our aim is to differentiate between the chromaticities of very similar colours, the Munsell colours used to test the algorithm were taken for a small range of L values (47.8<L<50.2). From these colours two sets of 100 pairs of colours were chosen in such a way that the Euclidean distance between members of each pair in CIELab space was either in the range between 1.0 and 4.0 units or between 4.6 and 6.0 CIELab units. CIE standard test illuminant spectra (colour temperature 5100 K, 5520 K, 5570 K, 5620 K, 5670 K, 5720 K, 5770 K, 5900 K, 6100 K, 6400 K, 6600 K, 7100 K, 7900 K and 8800 K) were chosen in such a way that the distribution of the colour temperature is similar to that of the typical measured daylight [13]. All these data was sampled at 1 nm interval.

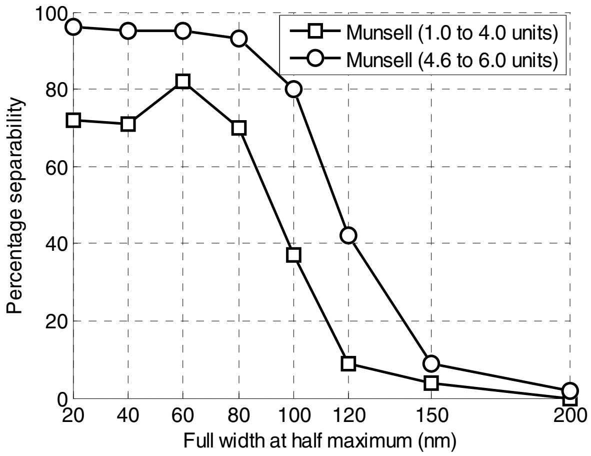

***Figure 2**. The separability results for the initial sensor combinations when the sensor responses are quantised to 10 bits. (Tested with Munsell and CIE daylight)*

A simple test of the separability of pairs of similar colour is to determine the smallest circle that encloses the cluster of points from each colour. Two colours are then considered to be separable if the distance between the centres of the two circles is larger than the sum of radii of the two circles. Using this definition of separability the performance of the algorithm was investigated for widths (full width at half maximum) of the sensor spectral responses from 20 nm to 200 nm. The results in Figure (2) indicate that even with the simple method of processing the sensor outputs most of the colours that are a good match to each other are separable and for some sensor widths it is even possible to separate colours that could be described as a very good match to each other. A common pattern for both sets of data is that the separability of colours decreases when the sensors are wider than 80 nm because the sensor responses become too correlated and any differences between the sensor responses are lost by quantisation.

Optimisation

The results obtained with equally spaced sensors were promising but an even better performance may be possible by optimising the choice of sensors to minimise the area of each colour in the feature space. This minimization has been achieved using a steepest descent algorithm and the sensor responses to 100 pairs of Munsell colours separated by between 6.0 and 7.1 CIELab units. In this optimisation these 100 pairs of samples were illuminated with 14 CIE standard daylight spectra (5000 K, 5500 K, 5550 K, 5600 K, 5650 K, 5700 K, 5750 K, 6000 K, 6200 K, 6500 K, 6700 K, 7000 K, 8000 K and 9000 K). The error measure used in this optimisation process was the ratio between the averages of the largest dimension of both clusters of a pair to the distance between corresponding cluster centres, all shown in figure (3), averaged over all pairs.

$$Error\ Metric = \frac{1}{100}\sum_{i=1}^{100}\frac{d_i^1 + d_i^2}{r_i} \qquad (4)$$

During the optimisation the parameters α, γ and the longest and shortest wavelengths of the peak spectral responses were taken as the independent parameters. This choice of independent variables made it possible to ensure that the values of the coefficients α and γ did not deviate too far from 0.5, so that all of the sensor responses made a non-negligible contribution to the feature space.

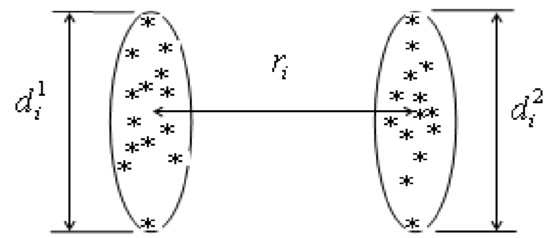

***Figure 3**. i^{th} cluster pair on the chromaticity space*

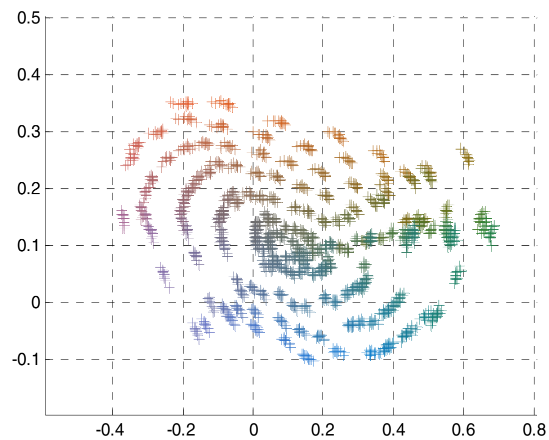

***Figure 4**. The two dimensional feature space formed by 80nm wide sensors from 202 Munsell surfaces and 14-CIE standard daylights after the sensor positions have been optimised. The data was quantised with 10-bits quantiser.*

***Table 2**. Optimised peak sensor positions and the channel coefficients with 80 nm full width half maximum sensors when quantising the channel responses with a 10 bits quantiser.*

Sensor ID	1	2	3	4
Peak position (nm)	462.4	539.5	607.6	661.3
Channel coefficients	α=0.5980 and γ=0.6090			

Figure (4) shows the two dimensional feature space formed by the optimised sensors listed in table 2. Again very different colours are well separated and compared to the corresponding equally spaced sensors the area covered by each Munsell colour is smaller. An important feature of this optimisation method is that the gap in the initial chromaticity space in figure (1) has been filled with colours in such a way that the colour variation across the chromaticity space is smooth. More importantly, the separability results in figure (5) show that optimisation has lead to a better separability performance, especially for the colours that have been described as very good matches to each other with narrow sensors.

The logarithmic sensors and the features have been chosen to deal with high dynamic range scenes. This aspect of the performance of the proposed sensors has been investigated by modeling the response of the sensors to illuminants with intensities that differ by a factor of 10^6, 120dB. The separability results for optimised sensors with different widths, such as those in figure (6), show that even with this extreme dynamic range the separability results are almost identical despite the very different illumination intensities.

Effect of changing illuminants and reflectances.

Results such as those in section (3) show that optimisation is expected to lead to an improvement in the proposed system's ability to separate similar colours, especially those that differ by between 1.0 and 4.0 CIELab units. Any optimisation process potentially suffers from the problem that it might lead to a solution that depends upon the data used in the optimisation process. This was a particular concern in this case because a scarcity of readily available reflectance data sampled at 1nm intervals meant that the only suitable readily available data set was the Munsell colours and these colours are known to be formed from a limited range of basis functions [12]. Similarly for convenience the illuminants used in the optimisation and initial assessments were the standard CIE illuminants that are calculated based upon a set of basis functions [10]. To investigate the impact of these choices the ability of the algorithm to separate non-Munsell colours illuminated with real measured daylights has been investigated. In particular measured daylight spectra [13] and a database of reflectances measured from a range of flowers [14] have been used.

Figure (7) shows the results obtained using the CIE standard illuminants but changing the reflectances of the surfaces. There is a small difference between the results obtained for the colours from the Munsell and floral databases that are a good match to each other when testing with narrow sensors, but this may be due to the slightly different range of CIELab differences that had to be used to obtain 100 pairs of colours in each case. A more significant difference is obtained for the colours that are described as very good matches to each other. In the particular case the Munsell data overestimates the performance that can be achieved with the floral data in the range of 40 to 100 nm width. Both Munsell data sets under estimate the performance of the algorithm with wide sensors compared to the floral data sets.

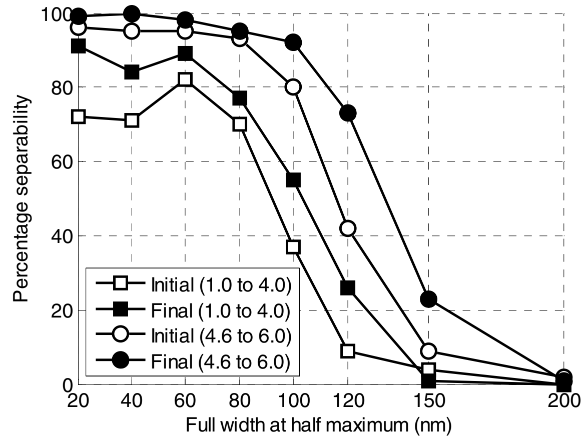

***Figure 5**. The separability results with both the initial and optimised sensor combinations when the sensor responses are quantised to 10 bits. (Tested with Munsell and CIE test daylight spectra)*

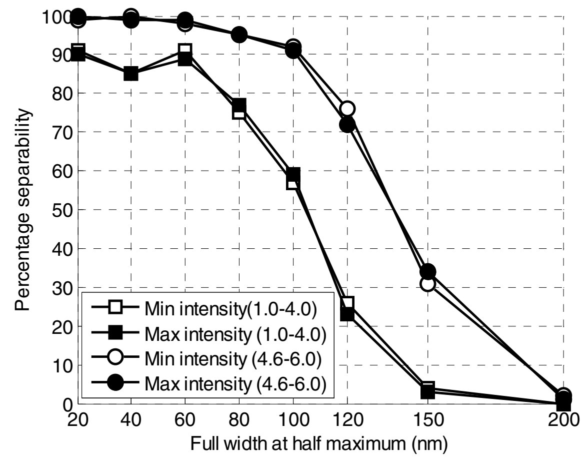

***Figure 6**. The colour separability results when the sensor responses are quantised to 10 bits for two illuminant intensities that differ by a factor of 120dB (or 10^6). This result shows that as expected this approach is able to deal with a wide dynamic range of inputs.*

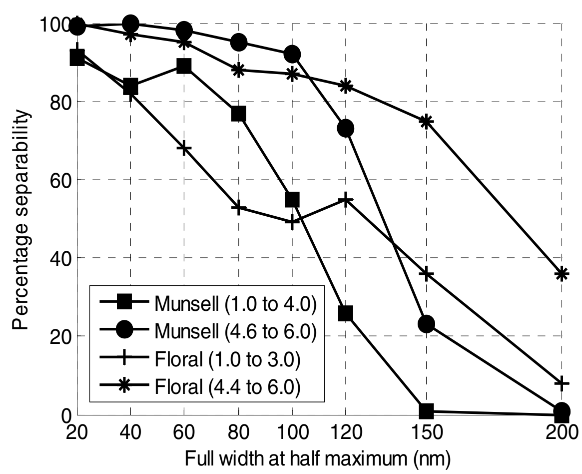

***Figure 7**. Separability results of the optimised sensor set with 10-bits quantised response. CIE standard test daylight was applied in this test.*

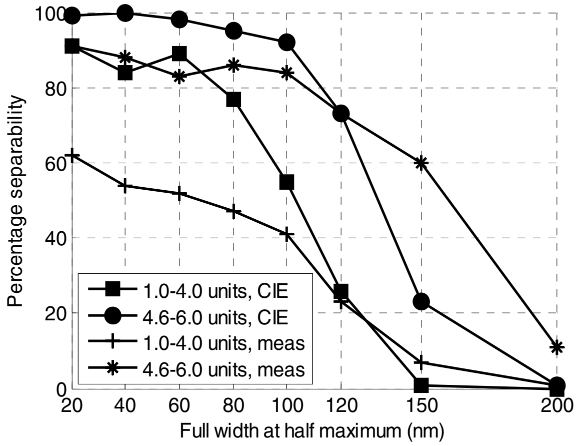

***Figure 8**. Results of the optimised sensors with 10-bits quantisation showing the effect of changing the illuminant on the separability results from Munsell colours.*

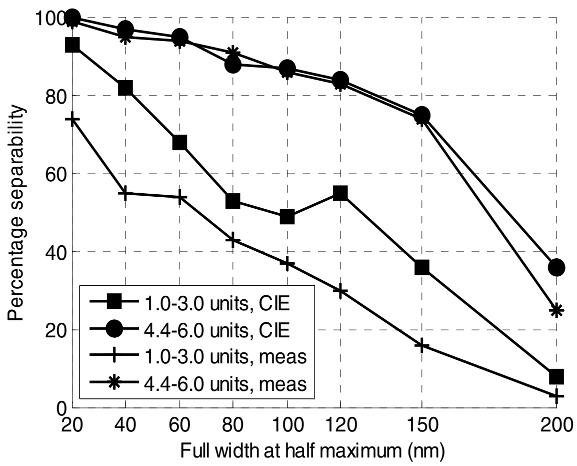

***Figure 9**. Results of the optimised sensor with 10-bits quantisation showing the effect of changing the illuminant on the separability results of the floral samples.*

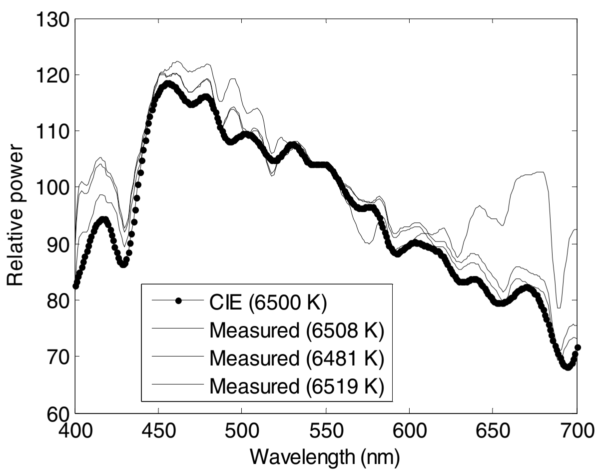

***Figure 10**. Spectra of CIE (6500 K) and measured daylight with correlated colour temperature 6508 K, 6481 K and 6519 K.*

The results in figures (8) and (9) show the effect of changing the illuminant from the CIE standard to 34 daylights measured between 09.00 and 16.00 GMT on six days spread throughout 1996. The difference between these two figures is that figure (8) compares the separability of Munsell colours under both illuminants whilst figure (9) shows the separability of the floral data. In the results presented in figure (8) the CIE standard daylight spectra data overestimates the performance with narrow sensors. Comparisons of spectra such as the one in figure (10) suggest that this difference in results might be caused by the fact that the fine structure in the measured daylight spectra is different from that in the CIE standard daylight spectra. When testing with CIE standard spectra and floral data the colour samples described as good match to each other correctly estimates the performance that can be achieved with the measured spectra. However, the performance is over estimated with the floral data set in which the sample pairs are described as very good match to each other. Comparing the results presented in figures (8) and (9) it can be seen that with harder data set and with narrow sensitivity functions, the CIE spectra overestimates the performance that could be achieved with measured daylight spectra.

In figure (11) the results obtained with the Munsell and CIE standard test illuminants are compared to the results obtained with the floral data and the measured daylight spectra. A comparison of results for the two sets of colours that are a very good match to each other suggests that for applications that require very fine colour discrimination it will be necessary to develop a system using representative data. Given the difficulty of separating these colours this is not surprising. In contrast to these results when the sensors FWHM is less than 120nm the results for the pairs of colours that are good matches are very similar. This suggests that to as long as the sensor responses are narrower than 120nm it might be possible to design and develop a generic system to separate colours that are good matches to each other.

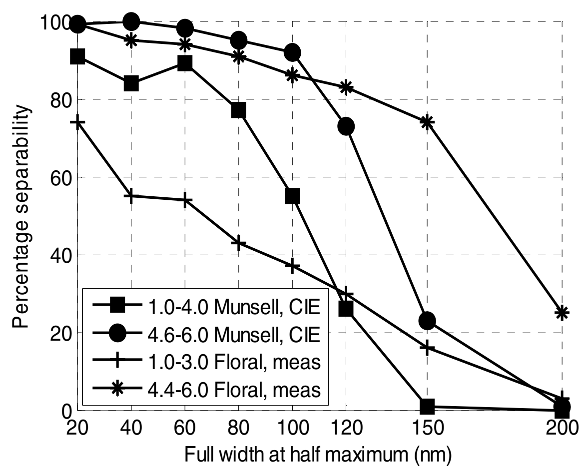

***Figure 11**. The separability results for optimised sensor combinations with 10-bits quantised response. The results shown are for the Munsell and CIE standard test daylights and floral with measured daylights which might represent data from a possible application.*

Concluding Remarks

A simple method to extract two illuminant independent features from the responses of four logarithmic sensors has been proposed. The illuminant independent chromaticity space formed with these two features was shown to have similar colours as neighbours and different colours widely separated. The significance of the small area in this space occupied by each colour was investigated using a simple method to determine the separability of perceptually similar colours. The separability of two sets of pairs of colours that can be described as either good or very good colour matches to each other have been reported. Using logarithmic sensors with peak spectral responses that are uniformly spaced and 10 bits to represent their output data good results are obtained for both sets of colours if the full width at half maximum of the sensors spectral response is 80nm or less. Both sets of colours can be separated more successfully if the sensor spectral responses are optimised. Furthermore, results have been obtained that suggest that the system has the expected high dynamic range. The effect of changing the data used to generate the results has been investigated. The results obtained suggest that applications that require discrimination between colours that are very good matches to each other may require application specific combinations of sensors. However, it appears that colours that are good matches to each other can be separated using a generic sensor system.

References

[1] S. Sugawa, N. Akahane, S. Adachi, K. Mori, T. Ishiuchi and K. Mizobuchi, "A 100 dB Dynamic Range CMOS Image Sensor Using a Lateral Overflow Integration Capacitor," ISSCC Dig. Tech. Papers, :352-353, (2005).

[2] G. D. Finlayson, and S. D. Hordley. "Color constancy at a pixel". J. Opt. Soc. Am. A, 18:253-264, (2001).

[3] B. Choubey, H. Cheng and S. Collins, A wide dynamic range CMOS image sensor with an adjustable logarithmic response. IS&T/SPIE's 20th Annual Symposium Electronic Imaging Conference :6816-6819, (2008).

[4] T. Morimune, H. Kajii, and Y. Ohmori. "Photoresponse Properties of a High-Speed Organic Photodetector Based on Copper–Phthalocyanine Under Red Light Illumination". IEEE Photonics Technology Letters, Vol. 18, (2006).

[5] J. Romero, J. Hernández-Andrés, J. L. Nieves and E. M. Valero. "Spectral Sensitivity of Sensors for a Color-Image Descriptor Invariant to Changes in Daylight Conditions". Colour Research and Application, 31:391-398, (2006).

[6] M. Ebner, Color constancy Wiley- IS&T Series in Imaging Science and Technology, (2007).

[7] G. D. Finlayson, and M. S. Drew. 4-sensor camera calibration for image representation invariant to shading, shadows, lighting, and specularities. IEEE International Conference on Computer Vision, :473-480, (2001).

[8] J. A. Marchant., and C. M. Onyango. "Shadow-invariant classification for scenes illuminated by daylight". J. Opt. Soc. Am. A, 17:1952–1961, (2000).

[9] Database – Munnsell Colours Matt ftp://ftp.cs.joensuu.fi/pub/color/spectra/mspec/.

[10] Munsell Color Science Laborartory, "Daylight Spectra." http://mcsl.rit.edu/

[11] A. Abrardo, V. Cappellini, M. Cappellini, and A. Mecocci. Art-works colour calibration using the VASARI scanner. In Proceedings of IS&T and SID's 4th Color Imaging Conference: Color Science,Systems and Applications, Scottsdale, Arizona, 94-97, (1996).

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

[13] J. Hernández-Andrés, J. Romero, J. L. Nieves, and R. L. Lee, Jr., "Color and spectral analysis of daylight in southern Europe," J. Opt. Soc. Am. A 18:1325-1335, (2001).

[14] S. E. J. Arnold, V. Savolainen and L. Chittka. "FReD: The floral reflectance spectra database. Nature Precedings (2008) http://dx.doi.org/10.1038/npre.2008.1846.1.

Author Biography

Sivalogeswaran Ratnasingam received the B.Sc. Eng degree in Electrical and Electronic Engineering from the University of Peradeniya Sri Lanka, in 2003 and the MSc. degree in Mobile and Satellite Communications from the University Surrey UK in 2006. He is currently working towards the Ph.D. degree at the University of Oxford, Oxford, UK, in the field of colour vision. Mr. Sivalogeswaran won the Nokia prize for best overall performance in MSc in 2006.

Steve Collins received the B.Sc. degree in theoretical physics from the University of York, York, U.K., in 1982 and the Ph.D. degree from the University of Warwick, Warwick, U.K., in 1986. From 1985 to 1997, he was with the Defence Research Agency, where he worked on various topics, including the origins of 1/f noise in MOSFETs and analogue information processing. Since 1997, he has been with the University of Oxford, Oxford, U.K, where he has continued his interest in smart imaging sensors and nonvolatile analogue memory.

Javier Hernández-Andrés received his Ph.D. degree in Physics from the University of Granada, Spain, in 1999. Since 2003 he is an associate professor in the Department of Optics at the same University. His research interests are colour-image processing, multispectral colour science, applied colorimetry, colour vision and atmospheric optics.