Optimum Sensors for 'Chromaticity' Constancy in the Pixel

Sivalogeswaran Ratnasingam and Steve Collins

Department of Engineering Science,
University of Oxford, OX1 3PJ, Oxford, United Kingdom
{siva,collins}@robots.ox.ac.uk

Abstract. In machine vision systems recording the colour of an object is crucial for applications such as skin detection while it will enhance applications including colour based recognition and image retrieval. Unfortunately, almost none of the existing colour constancy algorithms have been designed to deal with the high dynamic ranges that can occur in external, naturally illuminated scenes. One algorithm that can deal with these scenes has been proposed by Finlayson and Drew. In this paper a method of assessing the performance of this algorithm, and equivalent algorithms, are proposed. The performance of this algorithm is then significantly improved by optimising the spectral response of the sensors used to obtain the data required by algorithm. Since the resulting performance is comparable to that of the human visual system it appears that this algorithm is capable of obtaining useful chromaticity information under highly varying illumination conditions.

Keyword: sensor optimisation, chromaticity constancy.

1 Introduction

The high dynamic range (HDR) of illuminant causes difficulties in capturing a scene independent of lighting conditions. This large variation in intensity (10^9 [1]) of an illuminant makes the performance of the state of the art colour constancy algorithms to degrade in HDR scenes [2]. Recent research shows that the most advanced algorithms are not able to provide stable chromaticity that can be used in colour based recognition applications [3]. In HDR scenes the effect of scene illuminant should be discounted at pixel level and this is difficult with three sensor linear responses.

Foster states that the conventional colour measurement methods concentrate not on colour constancy but on complementary aspects [4]. This suggests that to achieve better chromaticity constancy, the imaging sensors should be optimised. Ohta [5] shows that exact colour reproduction is impossible with currently available dyes. It is proved by Verhel et al. [6] that four sensors are required for accurate colour reproduction. Based on the above discussion, four sensor imaging system was chosen for solving the lighting effect in HDR scenes and optimise the sensitivity of the sensors for better chromaticity constancy. Finlayson and Drew's [7] algorithm is investigated for identification of similar colours and the sensor parameters are optimised for better

P. Foggia, C. Sansone, and M. Vento (Eds.): ICIAP 2009, LNCS 5716, pp. 748–756, 2009.
© Springer-Verlag Berlin Heidelberg 2009

performance. In this work a new method of sensor optimisation for chromaticity constancy, error metric for optimisation, and a method for comparing the ability of identifying similar chromaticities and illuminant independent chromaticity space are proposed.

In this research work, the organic dyes are to be used in implementing the optimised imaging device for achieving chromaticity constancy in HDR scenes. As most of the organic sensors have a full width at half maximum of 100 nm this work focuses on the performance of the algorithm at 100 nm width sensors. Finlayson and Drew's algorithm is described and its performance is investigated in section 2. Section 3 discusses the optimisation of sensor parameters and the ability of similar colour separability of the algorithm. In section 4, the effect of quantisation on the algorithm's performance is investigated. Finally, based on the results, conclusions have been made in section 5.

2 The Algorithm

Based upon a model in which it is assumed that Planck's blackbody equation models the illuminant spectrum and that the sensors spectral response can be modeled by the Dirac delta function Finlayson and Drew have proposed a color constancy algorithm that only uses information from 4-sensor responses [7]. The responses from each of these sensors are proportional to the logarithm of the incident photon flux in a different spectral range. The first step of the algorithm is to use either one of these logarithmic sensor responses or the geometric mean of all four responses as a normalising channel to compensate for possible changes in geometry and more importantly the brightness of the illuminant [7]. When a single logarithmic response is used as the normalising channel this automatically reduces the dimension of the space of possible responses by one. Alternatively, when the geometric mean is used as the normalising channel one of the normalised channels can be discarded. In either case the result is a three dimensional space of normalised response ratios. The results presented in this paper are for the first case. In this three dimensional space each colour follows a trajectory as the spectrum of the illuminant varies. To remove this illuminant induced variation, the 3-dimensional space is projected into a 2-dimensional space in which the variations due to changing illuminant spectrum are minimised. The resultant 2-dimensional space is approximately independent of shading, scene geometry and scene illuminant [7].

To test this algorithm's ability to extract data that can be used to distinguish perceptually similar colours the algorithm has been implemented in MATLAB. The algorithm was then applied to the simulated responses of sensors whose inputs were obtained from the reflectance spectra of Munsell colours[8] illuminated by 6 different CIE standard daylight illuminants (D40, D45, D50, D55, D60 and D65) [9], Both the reflectance spectra of the Munsell colours and the standard daylights are sampled at 1nm intervals in the range between 400nm and 700nm and the spectral sensitivity of the sensors was modeled using a Lorentzian function.

Initial investigations of the algorithm were performed with sensors that have a full width half maximum of 100nm and peak spectral responses at 437.5nm, 512.5nm, 587.5nm and 662.5nm. Figure 1 shows the two dimensional chromaticity space formed by the algorithm when the sensor with the lowest peak wavelength response was used to normalise the other responses. In this figure the colour of each cross matches the Munsell colour used to generate the corresponding data point. This representation means that the results clearly show a smooth variation in chromaticity across the space with white at its centre. The other noticeable feature of these results is that dark colours are spread across the space however, a closer inspection of these results shows that they have the similar reflectance spectra (chromaticity) to their neighbours but with a low reflectivity.

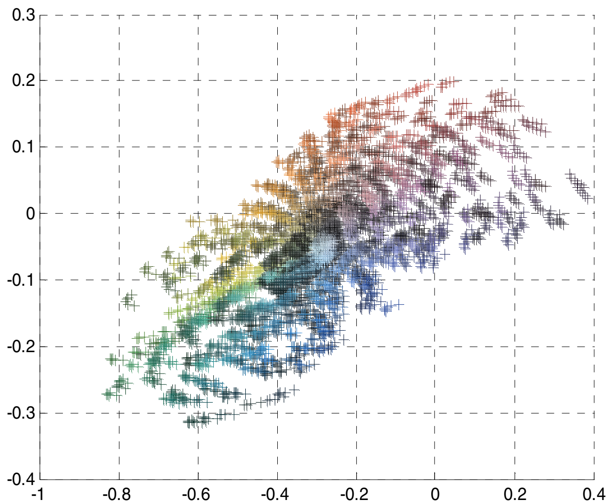

Fig. 1. The chromaticity space with 1269 munsell colours projected when applying 6-CIE standard daylight illuminants. Assuming the white point as the centre of the space, colour red is projected on the positive y-axis, orange is in the second quadrant, yellow is on the negative x-axis, green is in the third quadrant, blue is projected on the negative y-axis, purple is on the positive x-axis and brown is in the first quadrant.

A key feature of the algorithm is that it normalises the sensor responses in order to accommodate variations in the brightness of the illuminant. This also normalises the brightness component of a colour preserving only the chromaticity. This feature of the algorithm has been confirmed by scaling the reflectance representing each Munsell colours by different factors to create families of chromaticities, such as the one shown in Figure 2. As expected each of these families are projected to the same point in the two dimensional space created by the algorithm, for example the scaled reflectances in Figure 2 are all mapped to the same point which is circled in Figure 2.

A closer inspection of the space (figure 1) shows that each colour sample creates a small cluster of six points in the chromaticity space, one for each of the six day light illuminants. This space therefore has the basic characteristics required to separate different chromaticities. It is important to use a perceptually relevant scale to quantify the similarity between neighbouring chromaticities. For this the separability of the projected areas different colour samples has been tested.

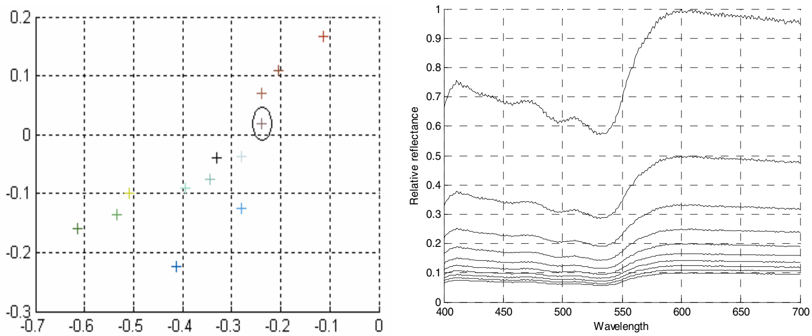

Fig. 2. Illustrates the chromaticity space, all synthetic reflectances shown in figure are projected as a single point (shown with in the circle) with some neighbor colours. These synthetic reflectance spectra were obtained by scaling the spectra with different factors.

The CIELab space is one colour space that has been designed so that distances in the space are proportional to the perceptual differences between colours. Although the method of calculating a colour co-ordinates in this space have been agreed there are several slightly different qualitative descriptions of colours separated by different distances in CIELab space. One example of a set of qualitative descriptions is that due to Abrardo et al. who 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 [10].

Since the algorithm is expected to differentiate the chromaticity between different surfaces the Munsell colours used to test the algorithm were taken for a small range of L values in CIELab space. After examining a histogram of the L values within the Munsell colours those colours with L values between 47.8 and 50.2 were chosen to obtain the largest possible set of colours that differ mainly in their chromaticity. From these colours two sets of test samples with 100 pairs of colours in each set 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. These distance ranges were chosen considering the available data.

In testing the performance of the algorithm 6-CIE standard illuminants (D41, D46, D51, D56, D61 and D64) and these two test sets (reflectance spectra) are used to generate the spectra of light incident on the simulated sensors. The algorithm is then applied to project each reflectance spectrum and illuminant pair into the chromaticity space. As expected each of the six illuminants meant that each reflectance sample created a cluster of six points in the chromaticity space. The first step in our method of determining the seperability of these spectra is to determine the smallest circle that encloses each of these clusters. If the centre distance between the two circles enclosing the members of each pair is larger than the sum of radii of the two circles then this pair of spectra is considered to be separable. Using this definition of seperability the performance of the algorithm was investigated for widths (full width at half maximum) of the sensor responses from 20 nm to 200 nm. The algorithm was tested with all four options of normalising the sensor responses and the results for the best choice of normalizing channel (see section 3) are reported in this paper.

Figure 3 shows the performance of the algorithm when used to separate perceptually similar colours with data obtained from sensors with different widths. Since the algorithm was used upon a mathematical model that assumed a very narrow spectral response it is not surprising that the algorithms performance decreases as the sensor width increases. Despite this tendency these results suggest that with this algorithm it is possible to separate colours that represent a good match to each other.

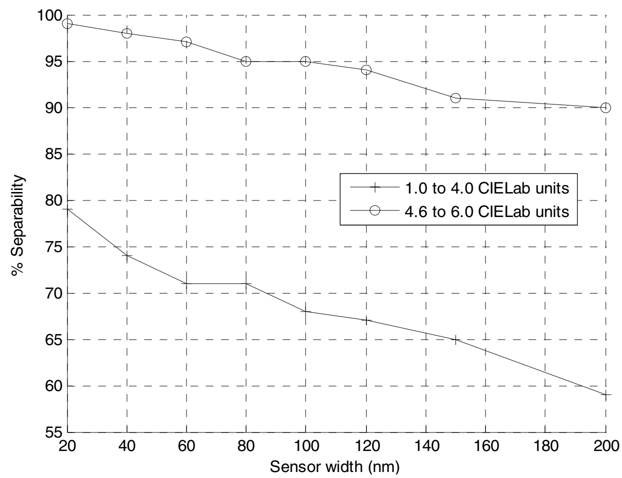

Fig. 3. Illustrates the test results of both test sets when applying the unquantised channel response to the algorithm. Width of the sensors is varied from 20 nm to 200 nm.

3 Sensor Optimisation

The sensors used to obtain the data for the results in Figure 3 were positioned uniformily but there is no reason to believe that this is the best choice of sensors. The sensor spectral characteristics was therefore optimised using a steepest descend algorithm. In this optimization process the position of the peak spectral responses of each sensor was allowed to vary while keeping all four widths equal and constant. There is a possibility that if the parameters are chosen to minimise the average size of the cluster of points from each colour it would merge the different colours. To overcome this problem, 100 pairs of training samples were chosen in the same way as the test sets but with the members of each pairs separated by 6.0 to 7.1 units in the CIELab space. The success of the algorithm means that the separability of these colours will not improve significantly. The error measure used in this optimisation was therefore the ratio between the average of the largest dimension of both clusters of a pair and the distance between the cluster centres. This ratio is averaged over all 100 pairs of samples in the training set. This error measure was chosen to measure the spread of a colour due to variation of illuminant on the chromaticity space relative to the separation between similar colours and the sensor parameters were chosen to minimise this ratio. In this optimisation six CIE standard daylight illuminants are applied to illuminate the training samples. Figure 4 illustrates the chromaticity space formed by the optimised sensor response. The optimised sensors (peak positions: 422.1, 452.8, 523.1 and 557.2 nm) were obtained by starting the optimisation process with initial sensor positions used in section 1 and full width at half maximum of 100 nm.

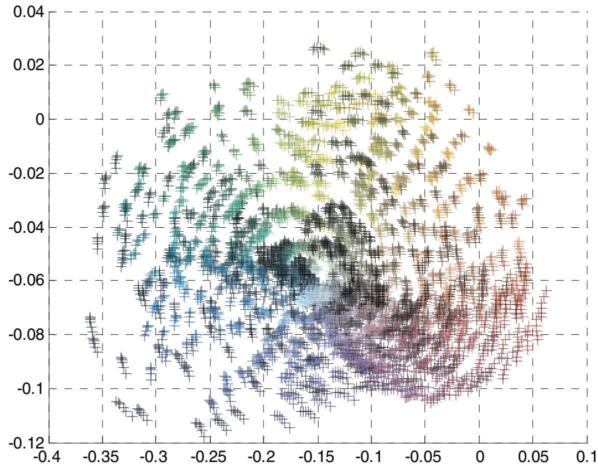

Fig. 4. Chromaticity space formed by the optimised sensor set. Six different CIE standard daylight illuminant spectra and 1269 munsell data were applied in generating this space.

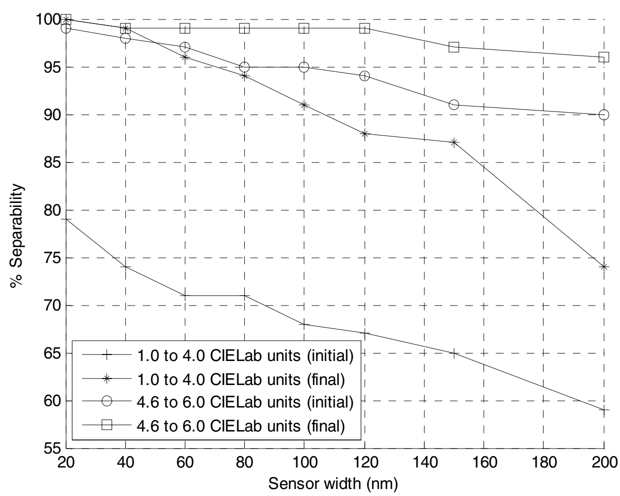

Fig. 5. Illustrates the initial and final performance of the algorithm when testing with both test sets. The channel responses applied in this test are unquantised sensor responses.

The chromaticity space formed by the optimised sensors (figure 4) shows a smooth variation of colour across the space and the average cluster spread compared to the total area of the space is 40 % smaller than in the initial space. This simple measure suggests that samples should be more easily separated.

Again the ability to separate different reflectances at different sensor widths has been tested. The results in figure 5 show that optimisation of the position of the peak spectral response of each sensor improves the performance of the algorithm in separating similar colours. A significant improvement is achieved in separating the most similar colours (that is those separated by between 1.0 and 4.0 CIELab units). In fact these results suggest that the algorithm can show a comparable performance to the human visual system.

4 Quantisation

One aspect of a real camera that has not been included in the model used to obtain the results in Figure 5 is the quantization of the sensor responses by the analogue-to-digital converter (adc). To study the effect of quantisation on the performance of the algorithm, the algorithm was optimised with different quantisers. The first stage in representing the quantiser is to determine the maximum sensor response. The CIE standard daylight illuminant (D65) was used to determine the maximum response and since the optimisation process changes the effective position of the sensor, the maximum sensor response was found by shifting the sensor's peak response from 400nm to 700 nm in 1 nm steps. The maximum response was calculated for each of the sensor widths of interest and the maximum response for each width was divided by 2^n, n is the number of bits in the quantiser which was varied between 6 and 16. Each sensor response was then equated to the nearest of these quantized values.

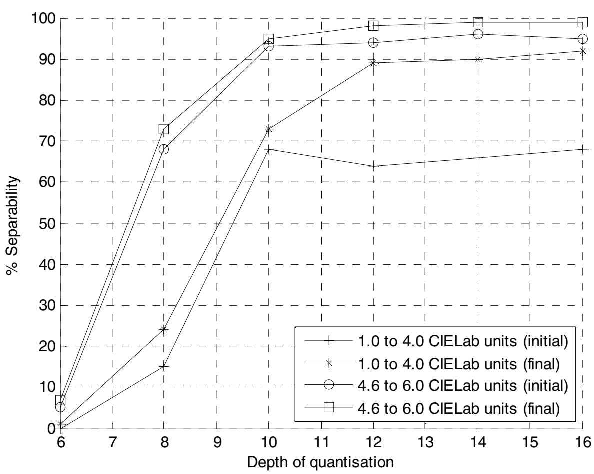

Fig. 6. Illustrates the initial and final seperability of colours as a function of quantisation depth. Full width at half maximum of all four sensors was kept at 100 nm.

Figure 6 shows the separability of similar colours with varying numbers of bits used to represent the sensor responses for two sets of data and two sets of sensors peak positions. These results show that in all four cases the performance of the algorithm increases rapidly when the number of quantisation bits increases from 6 to 10.

This suggests that the least significant bit of a 10 bits quantiser is small enough to capture the information required to separate perceptually similar colours. These results also show that optimizing the sensor positions can result in a significant increase the in separability of similar colours, especially when more than 10 bits are used to represent the output from each sensor.

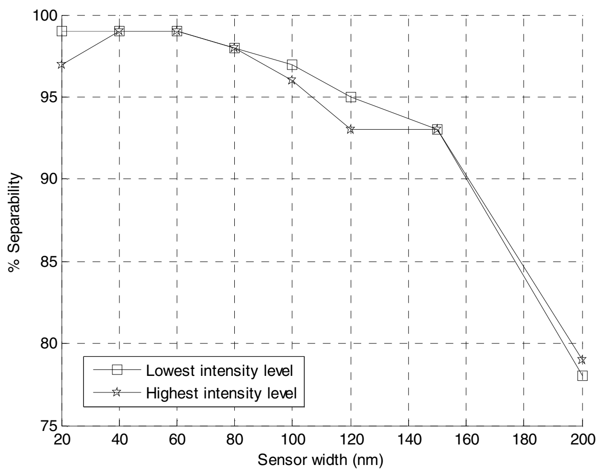

Fig. 7. Shows the test results of both test sets with initial and optimised sensor sets. The channel responses are quantised with a 10-bits quantiser.

Although the results in Figure 6 suggest that there might be some advantage from using a 12 bit adc to represent the response of sensors most cameras are currently supplied with either an 8 bit or a 10 bit adc. The final aspect of the impact of quantization that has been investigated is to confirm colour seperability despite the high dynamic range of illuminants. For this investigation a 6-decade intensity range of the illuminant was represented by dividing the daylight illuminant spectra by 10^6. The seperability of colours for the highest and lowest illuminant intensities are shown in Figure 7. This figure shows the results obtained with data quantised to 10-bits when separating pairs of colours that differ by between 4.6 and 6.0 CIELab units. The first significant feature of these results is that the performance of the algorithm drops at larger width. This is because as the sensor width increases the correlation between the sensor responses to the two members in each pair will increase and eventually the difference between the two responses will be smaller than the difference represented by the quantized sensor responses. When this occurs the quantized data will not contain the information needed to separate the member of each pair. These results suggest that any sensors used to obtain the data for this algorithm should have a full width half maximum of 100nm or less.

Another important conclusion that can be drawn from results such as those in Figure 7 is that the results obtained are almost identical for the two simulated illumination intensities. Since these differ by a factor of 10^6 these results suggest that this combination of sensors and algorithm is capable of dealing with HDR scenes. Furthermore, these results show that using the data for sensors with a full-width half maximum of less than 100nm it is possible to separate almost all the colours that would be described as representing either a good [10] or acceptable [11] match to each other. This suggests that using data from sensors with optimised spectral responses, the performance of the algorithm in separating similar chromaticities is comparable to the human visual system.

5 Conclusion

An existing algorithm has been investigated for chromaticity constancy. The algorithm uses the ratios of the outputs from four sensors to accommodate potentially large variations in the amount of reflected light falling on a sensor. This normalization process results in three normalized sensor responses that are independent of the illuminant intensity but dependant upon the illuminant spectrum. This illuminant dependence is then removed by extracting a two dimensional illuminant independent chromaticity descriptor from these three normalized responses. A method of assessing the resulting space has been proposed based upon the ability to easily separate perceptually similar colours. Using this method of assessment it has been shown that colours are more easily separated if the spectral response of the sensors used to gather the image are optimized. Furthermore, this is particularly important when the effects of quantization to represent the effect of an analogue to digital converter in the camera are taken into account. However, using a 10 bit adc and sensors with optimized peak spectral responses it is possible to separate colours that are good matches to each other even when the illuminant brightness changes by 10^6.

References

1. Holst, G.C., Lomheim, T.S.: CMOS/CCD Sensors and camera systems. SPIE press, Washington (2007)
2. Ebner, M.: Color constancy. Wiley- IS&T Series in Imaging Science and Technology (2007)
3. Funt, B., Barnard, K., Martin, L.: Is colour constancy good enough? In: Burkhardt, H.-J., Neumann, B. (eds.) ECCV 1998. LNCS, vol. 1406, pp. 445–459. Springer, Heidelberg (1998)
4. Foster, D.: Does colour constancy exist? Trends in Cognitive Science 7, 439–443 (2003)
5. Ohta, N.: Optimisation of Spectral Sensitivities. Phot. Sci. Eng. 27, 193–201 (1983)
6. Vrhel, M.J., Trussell, J.H.: Optimal color filters in the presence of noise. IEEE Trans. Image Processing 4, 814–823 (1995)
7. Finlayson, G.D., Drew, M.S.: 4-sensor camera calibration for image representation invariant to shading, shadows, lighting, and specularities. In: ICCV, pp. 473–480 (2001)
8. Database – Munnsell Colours Matt, `ftp://ftp.cs.joensuu.fi/pub/color/spectra/mspec/`
9. Colorpro Communications, `http://www.colorpro.com/info/data/daylites.html`
10. Abrardo, A., Cappellini, V., Cappellini, M., Mecocci, A.: 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, pp. 94–97 (1996)
11. Hardeberg, J.Y.: Acquisition and reproduction of color images: colorimetric and multispectral approaches. PhD dissertation, Paris (1999)