

An Algorithm to Determine the Chromaticity Under Non-uniform Illuminant

Sivalogeswaran Ratnasingam and Steve Collins

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

Abstract. Colour based object recognition is a difficult problem because of the effect of scene illuminant and geometry on the captured image. In this paper the ability of an algorithm proposed by Finlayson and Drew [1] to separate similar colours is assessed. A new variant of this algorithm is then proposed that results in a slight improvement in performance. A significant performance improvement is achieved by optimising the characteristics of the sensors that are used to acquire the data for this algorithm. This optimisation process results in several combinations of sensors and associated data projections that have a comparable performance when required to distinguish between similar colours. Since this performance is comparable to that of the human visual system it is suggested that with the correct sensors this algorithm is capable of obtaining useful chromaticity information under varying illumination conditions.

Keywords: color based recognition, chromaticity constancy, colour indexing.

1 Introduction

Digital cameras have been an undoubted success in the consumer market in the last few years. However, compared to human visual system these cameras have two limitations, a limited dynamic range and variability in the apparent colour of objects. The first of these limitations causes saturation when the dynamic range of the scene is larger than that of the camera. The second limitation is subtler and means that, unlike human vision, colour information can not always be reliably used in artificial object recognition systems.

An early approach for recognizing objects based on colour was developed by Swain and Ballard [2] called ‘colour indexing’. This work is a significant contribution to colour based object recognition, but this method does not address the issue of variation of scene geometry and illumination. To solve this problem Funt and Finlayson[3] created colour indices using the ratio of responses from neighbour pixels. An alternative approach proposed by Berwick and Lee [4] used a diagonal transform to model the sensor response and a translation to model the effect of variations in the illuminant spectrum. This general approach was followed by Finlayson and Drew, who proposed an algorithm using the responses from four sensors to first compensate for changes in scene geometry and illuminant intensity, followed by a projection into a space in which the effect of changing the illuminant spectrum was minimized.

In this paper, Finlayson and Drew's algorithm [1] and a flexible method for removing illuminant effect are optimised to achieve better chromaticity constancy. The optimisation is done using a steepest descend algorithm and the performance of variations of both algorithms (normalizing by one of the channel response and by geometric mean of all four responses) is also investigated. Both algorithms map different colours into a 2-dimensional space, invariant to scene geometry and illuminant. The 2 dimensional spaces formed by these algorithms show a smooth spatial variation of colours across the space. Initial results suggest that in this space, colours that are separated by between 6 units in the CIELab space could be identified uniquely.

2 Algorithm

Differences between the spectra of different illuminants can cause significant changes between the receptor responses to the same surface when illuminated by the different illuminants. The aim of the algorithms is to extract an invariant representation of each colour under different illuminants.

The illuminants data used to assess these algorithms were obtained from CIE standard daylight [5] and the surfaces that were used were from the Munsell data set [6]. For each of the 1269 surfaces in this data set the reflectance spectra, measured using a spectrometer at 1 nm intervals, and the different colour co-ordinates, when viewed under a D65 illuminant, were available. For this study the sensitivity of the sensors were modeled using a Lorentzian function characterized by its full width at half the maximum response.

Finlayson and Drew. [1] started with the logarithm of each of four sensor responses. Then either one of these logarithmic sensor responses or the geometric mean of all four was used as a normalizing channel to compensate for possible changes in geometry and the brightness of the illuminant [7] [8]. When a single logarithmic response was used as the normalizing channel this removes the dependency on illuminant intensity and scene geometry while reducing the dimension of the space of possible responses by one. Alternatively, when the geometric mean is used as the normalizing channel one of the normalized channels is discarded to achieve the same effect. In either case the result is a 3 dimensional space of possible normalized response ratios. Within this space each colour follows a trajectory as the spectrum of the illuminant varies. To remove this variation, the 3 dimensional space is projected into a 2 dimensional space in which the variations due to changing illuminant spectrum are minimized.

Initially the performance of the algorithm proposed by Finlayson and Drew was investigated when used with four sensors whose responses are characterized by the parameters, given in table 1, chosen to give an even spread of peak sensitivities in the wavelength range 400nm to 700nm. The data used to assess these algorithms was

Table 1. Initial sensor parameters

Sensor ID	1	2	3	4
Peak position (nm)	437.5	512.5	587.5	662.5
Width (nm)	60	60	60	60

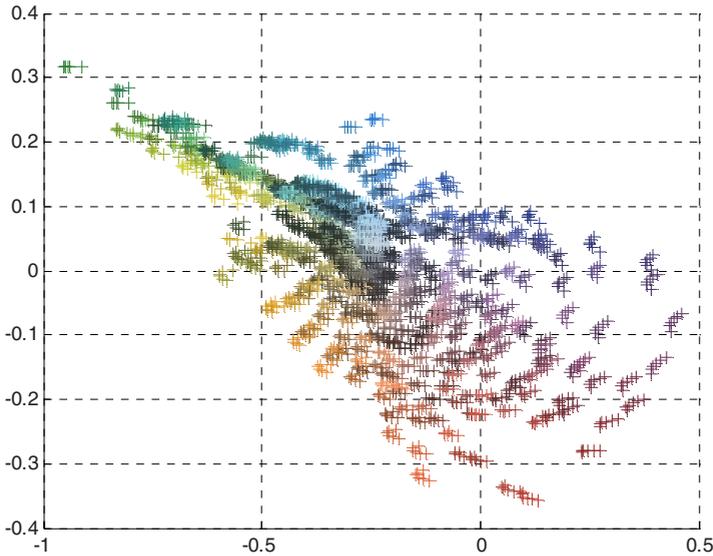


Fig. 1. 671 Munsell colours mapped on the two dimensional chromaticity space formed by the sensors given in table 1

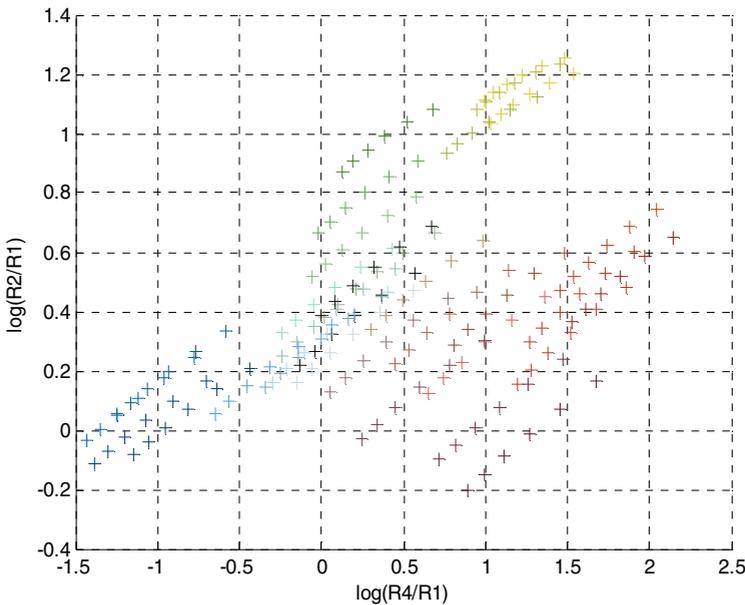
obtained by integrating the triple product of the reflectance spectra of each different colour, one of six CIE standard daylight illuminants with a colour temperature between 4000 K and 6500 K and a sensor response function over the wavelength range between 400nm and 700nm. The results in figure 1 show that the proposed algorithm creates a two dimensional space with similar colours projected into neighbouring areas. This suggests that noise in the data will not lead to a large error in the interpretation of the surface chromaticity.

This is a promising result. However, the most important assessment criterion for the algorithm is the ability to use its output to reliably distinguish colours. One colour space that has been defined to measure the perceptual differences between colours is the CIELab space. Ideally there is a just noticeable difference between two colours that are separated by a unit distance in CIELab space. However, there are various descriptions of the quality of colour matching associated with different distances in the space. For example matching colours between 3 and 6 units have been described as ‘good quality’ [9] or ‘acceptable’ [10] matching. To compare the performance of the algorithm to the human visual system, pairs of colours from the data set were identified that are separated by known distances in CIELab space.

As 6-CIELab unit is acceptable quality of colour reproduction, the algorithm was tested using 135 pairs of colours with the members of each pair chosen so that they are separated by 6-CIELab units. A close inspection of the final two dimensional space shows that each colour is projected into a small area within this space. A pair of colours was deemed to be distinguishable or separable from each other if the minimum distance between the two small areas was larger than the average of the largest dimensions of the two small areas. The results in Table 2 for the algorithm,

Table 2. Seperability of colour pairs

Normalising sensor	Colour separability (%)
1	80
2	85
3	75
4	80
Geometric mean and discard 1	78
Geometric mean and discard 2	80
Geometric mean and discard 3	75
Geometric mean and discard 4	80

**Fig. 2.** Two dimensional Log channel ratio space shows the variation of different colours under six CIE standard day light illuminants

with the different options for normalization, suggest that the algorithm performs well in this task, however, there is some room for improvement.

One possible approach for improving the algorithm is to replace the direct projection from a three dimensional space after normalization to the two dimensional chromaticity space with a more flexible two step projection procedure. In this alternative approach, as before the first step is to use either one of the sensor responses or the geometric mean of all four responses to compensate for possible changes in geometry of the scene and the brightness of the scene illuminant. In the latter case one of the normalized responses is then again discarded. The resulting three ratios of responses are then used to create 2 dimensional spaces. A typical 2 dimensional space is shown

in figure 2. The variability in this 2 dimensional space is caused by the changes in the illuminant. Feature extraction based on a Karhunen-Loève (KL) transform is therefore applied to each of these spaces to extract the feature that contains the minimum illuminant induced variability. Figure 3 shows the feature extracted from the data in figure 2 for a range of Munsell colours. These results show that this transformation has resulted in a feature that is almost independent of the illuminant. However, it is clear that some very different colors, such as samples 9, 11 and 15, in figure 3 have the same feature value and would therefore be confused.

This confusion is avoided by using the features obtained from an equivalent KL transformation applied to the two additional 2-dimensional spaces created from the other pairs of response ratios (as shown in figure 4). The three features obtained by this process can be used to create a 3 dimensional space that is largely independent of the illuminant and scene geometry. Any variations in this space are therefore dominated by differences between colours. In this space a feature extraction based on KL transform is therefore applied to extract the two features that contain the most variability. These two features are then used to form the chromaticity space.

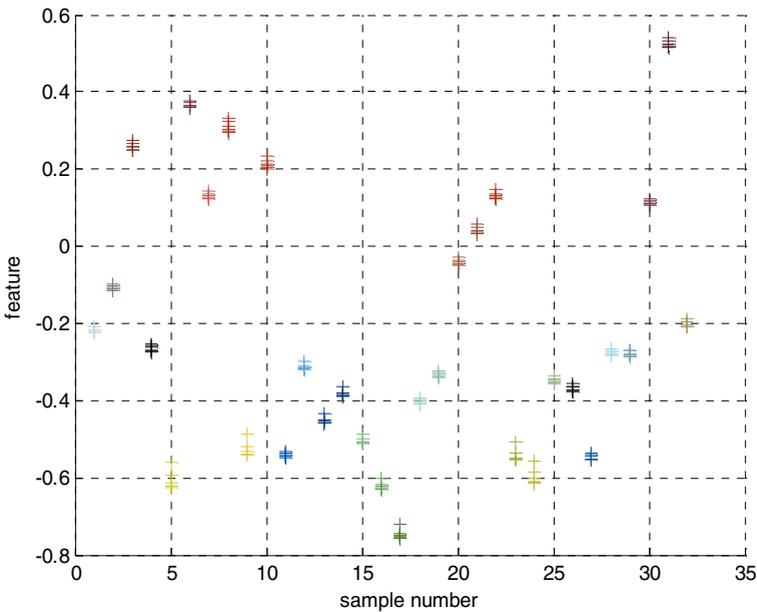


Fig. 3. Extracted chromaticity features from the two dimensional space shown in figure 2

The performance of this alternative procedure for obtaining a chromaticity space was assessed using the same sensor positions and test data as used to generate the results in table 2. Although there appears to be a slight improvement in performance this improvement is only a fraction of the possible improvement.

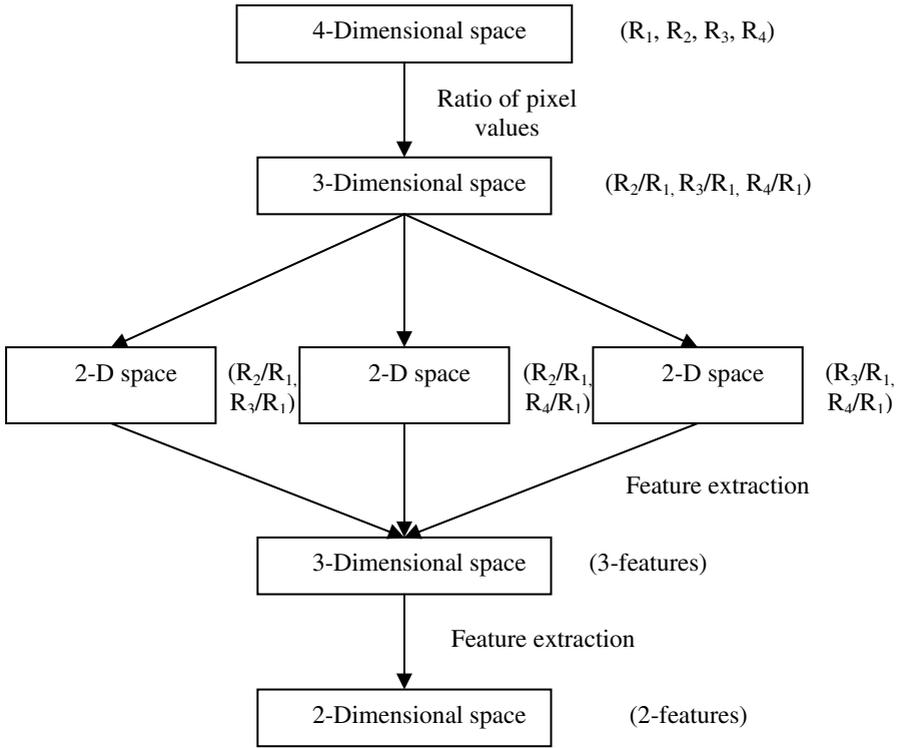


Fig. 4. The proposed method of obtaining a chromaticity space when one of the sensor responses is used to remove the effects of the lighting conditions

Table 3. Test results of the initial space using 6-CIELab test pairs

Normalising sensor	Colour separability (%)
1	85
2	80
3	80
4	80
Geometric mean and discard 1	80
Geometric mean and discard 2	80
Geometric mean and discard 3	77
Geometric mean and discard 4	73

3 Sensor Optimisation

The algorithms described in section 2 have several parameters, including the sensor characteristics and the projections into the 2 dimensional space. A closer inspection of

results such as those in figure 1 shows that although ideally each colour should be projected to a unique point in the space, in fact each colour is projected to a small area. This suggests that it might be possible to improve the separability of colours by optimising the sensor characteristics to minimise the projected area of each colour. To investigate the possibility of optimising the sensor characteristics, a steepest descent algorithm has been used to vary the position and width of each of the four sensors in a way that improves the separability of colours in the chromaticity space.

For this optimisation, 57 pairs of colours with the members of each pair chosen so that they were separated in CIELab space by between 2.95 and 3.05 units was used as the training data. There is a possibility that if the parameters are chosen to minimize the average size of the cluster of points from each colour it would merge the different colours. The error measure that was minimized by the steepest descent algorithm was therefore the average ratio of largest dimension of the small area representing each colour to the minimum distance between each pair of colours. With each of the four sensors characterized by a wavelength of peak sensitivity and a sensitivity width at half maximum, there are eight parameters that can be varied independently.

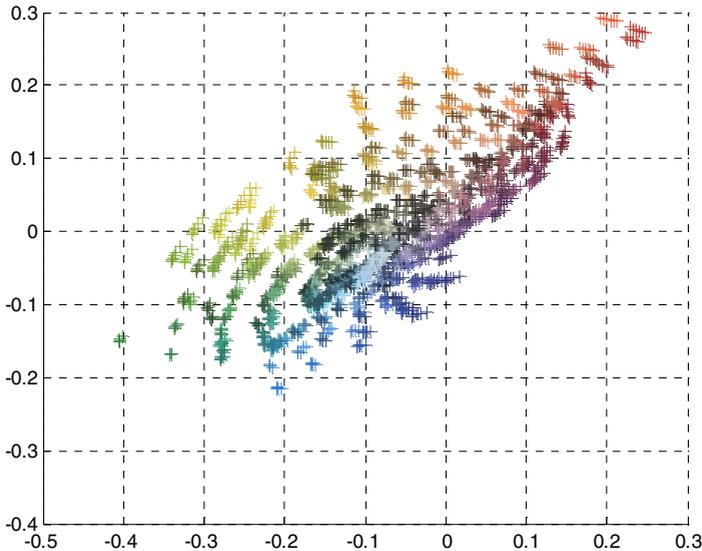


Fig. 5. The two dimensional space formed with a set of optimised sensors obtained when starting from the parameters given in table 1

The typical result in figure 5 is the optimised space corresponding to the initial space shown in figure 1. Again similar colours are neighbours in this space. However, a comparison of the two figures show that the result of this optimisation process is to shrink the area in the two dimensional space covered by the Munsell colours. The results in table 4 show that for all methods of normalisation and projection, when starting with the sensor parameters from table 1, the optimisation lead to a reduction in the area covered by the Munsell colours.

Table 4. Comparison of initial and final chromaticity spaces

Normalising channel	Area of the initial space	Area of the final space	Initial capacity	Final capacity
Results for direct projection				
1	1.1	0.40	0.97e4	3.0 e4
2	2.1	0.18	1.1 e4	1.6 e4
3	1.5	0.26	0.82e4	1.2 e4
4	1.2	0.37	0.93e4	1.1 e4
Geometric mean (discard 1)	0.6	0.14	0.86e4	3.3 e4
Geometric mean (discard 2)	0.4	0.15	0.86e4	0.86e4
Geometric mean (discard 3)	0.6	0.53	1.1 e4	1.1 e4
Geometric mean (discard 4)	0.7	0.34	1.0 e4	1.3 e4
Results for indirect projection				
1	1.7	0.36	1.0 e4	4.0 e4
2	3.2	0.85	1.1 e4	2.4 e4
3	2.1	0.31	0.78e4	0.88e4
4	1.6	0.09	0.87e4	3.9 e4
Geometric mean (discard 1)	1.0	0.47	0.94e4	1.4 e4
Geometric mean (discard 2)	0.6	0.33	0.83e4	3.2 e4
Geometric mean (discard 3)	0.8	0.35	1.1 e4	1.8 e4
Geometric mean (discard 4)	0.9	0.10	1.0 e4	0.80e4

Table 5. Comparison of colour separability of colours that differ by 6 CIE Lab units

Normalizing channel response	Colour separability of initial space (%)	Colour separability of optimised space (%)
Results for direct projection		
1	80	94
2	85	90
3	75	90
4	80	80
Geometric mean (discard 1)	78	90
Geometric mean (discard 2)	80	88
Geometric mean (discard 3)	75	81
Geometric mean (discard 4)	80	85
Results for indirect projection		
1	85	95
2	80	92
3	80	85
4	80	92
Geometric mean (discard 1)	80	93
Geometric mean (discard 2)	80	94
Geometric mean (discard 3)	77	90
Geometric mean (discard 4)	73	88

Although the reduction in the area of the projected space is unexpected the optimisation process was designed to maximise the distinguishability of colours. It is possible that this smaller space is better than the previous space if the projected area of each colour has reduced by a larger fraction than the area covered by all the colours. In table 4, the area of the chromaticity space has been estimated using the area of a smallest ellipse which encloses all the 1269 munsell colours. As the area covered by a typical individual colour is curvilinear in shape the area of each colour cluster is calculated by fitting smallest ellipse around each colour. The average area occupied by a colour is then obtained by averaging the areas of all the 1269 colours. The number of colours that can be accommodated within the project space was then estimated by finding the ratio between the area of chromaticity space to this average area. Table 4 shows that although the areas of the two dimensional spaces have been decreased by optimising the sensor characteristics, in most cases there are potentially more distinguishable colours in the chromaticity spaces after optimisation.

The results in table 5 show that as required the optimisation process has increased the percentage of colour pairs that can be distinguished after projection onto the chromaticity space. Furthermore, the performance of the better combinations of sensors and projection methods are so good and so similar that it is impossible to determine which of them is the best. However, from these results it is clear that the performance of these algorithms has been significantly improved by optimisation of the spectral sensitivities of the sensors.

4 Conclusions

An algorithm that extracts a two dimensional representation of the chromaticity of different colours has been investigated. The algorithm uses the outputs from four sensors to discount the effect of scene illuminant and geometry and since it processes information at the pixel level it is expected to work well in non-uniformly illuminated scenes.

The ultimate performance of the algorithm will be application dependant. However, it has been suggested that its general ability to distinguish colour chromaticity under varying illuminants can be assessed using Munsell colours with known separations in CIE Lab space. With uniformly spaced sensors the performance of the algorithm as proposed by Finlayson and Drew was encouraging but not perfect. A variant of this algorithm with more parameters was therefore investigated with a marginal improvement in performance. A method to change the sensor characteristics to optimise the performance of the algorithm when required to distinguish similar colours was therefore proposed. Several variants of the algorithm could then distinguish more than 90% of the pairs of perceptually 'well-matched' colours. Since the performance of several variants of the algorithms were indistinguishable other factors, such as the availability of sensors with the appropriate characteristics will determine the final choice between the different variants of the algorithm.

References

1. Finlayson, G.D., Drew, M.S.: 4-sensor camera calibration for image representation invariant to shading, shadows, lighting, and specularities. In: ICCV 2001: IEEE International Conference on Computer Vision, pp. 473–480 (2001)
2. Swain, M.J., Ballard, D.H.: Color Indexing. *International Journal of Computer Vision* 7, 11–32 (1991)
3. Funt, B.V., Finlayson, G.D.: Color constant color indexing. *IEEE transactions on Pattern analysis and Machine Intelligence*, 522–529 (1995)
4. Berwick, D., Lee, S.W.: A Chromaticity Space for Specularity, Illumination color- and Illumination Pose-Invariant 3-D Object Recognition. In: Proc. Int'l Conf. Computer Vision, pp. 165–170 (1998)
5. Colorpro Communications, <http://www.colorpro.com/info/data/daylites.html>
6. Database – Munnsell Colours Matt, <ftp://ftp.cs.joensuu.fi/pub/color/spectra/mspec/>
7. Finlayson, G.D., Hordley, S.D.: Color constancy at a pixel. *J. Opt. Soc. Am. A* 18, 253–264 (2001)
8. Finlayson, G.D., Schiele, B., Crowley, J.L.: Comprehensive Colour Image Normalization. In: Fifth European Conference on Computer Vision, Freiburg, Germany, June 1998, pp. 475–490 (1998)
9. 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, Scottsdale, Arizona (1996)
10. Hardeberg, J.Y.: Acquisition and reproduction of color images: colorimetric and multispectral approaches. Ph.D. dissertation (Ecole Nationale Supérieure des Télécommunications, Paris (1999)