

Aberystwyth University

On the separation of luminance from colour in images

Woodland, Alan; Labrosse, Frédéric

Published in:

Proceedings of the International Conference on Vision, Video and Graphics

Publication date:

2005

Citation for published version (APA):

Woodland, A., & Labrosse, F. (2005). On the separation of luminance from colour in images. In *Proceedings of the International Conference on Vision, Video and Graphics* (pp. 29-36). Eurographics.
<http://hdl.handle.net/2160/359>

General rights

Copyright and moral rights for the publications made accessible in the Aberystwyth Research Portal (the Institutional Repository) are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Aberystwyth Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Aberystwyth Research Portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

tel: +44 1970 62 2400
email: is@aber.ac.uk

On the separation of luminance from colour in images

Alan Woodland and Frédéric Labrosse

Computer Science Department
University of Wales, Aberystwyth
Wales, UK, SY23 3DB

e-mail: [a jw2 , f f l]@aber . ac . uk

On the separation of luminance from colour in images

Alan Woodland and Frédéric Labrosse

Department of Computer Science, University of Wales, Aberystwyth, UK
[ajw2,ffl]@aber.ac.uk

Abstract

*Many computer vision and graphics related techniques rely upon illumination invariance of images to derive meaning from images of an object under varying lighting conditions. This is all the appearance-based methods. In practice however this assumption does not hold if one is not careful with either controlling the illumination of the object when capturing its appearance or with some post-processing of the images. This paper presents results of experiments designed to analyse the usefulness for illumination invariance of two colour models, CIE $L^*a^*b^*$ and YUV, that have been designed to provide separation of the luminance information from the colour information, and compare them with more traditional colour models, RGB and HSV. This is done by evaluating the variations in each of the components of the different colour spaces in real images taken in variable illumination conditions. We also present a simple application example.*

Categories and Subject Descriptors (according to ACM CCS): I.4.7 [Image processing and computer vision]: Invariants

Please note: most images should really be viewed in colour and are available in the electronic version or should be available at <http://users.aber.ac.uk/ffl/>.

1. Introduction

Images are generally captured in the RGB colour space, mainly because of hardware constraints. Note that we refer in this paper to the RGB colour space as being the one usually used in computer graphics systems: components range from 0.0 to 1.0 (or 0 to 255), 1.0 being the maximum amount of the corresponding colour, the “colours” black and white being respectively (0.0, 0.0, 0.0) and (1.0, 1.0, 1.0). This has implied that most vision and graphics algorithms have been using RGB, a colour model that has many drawbacks. One of these is that the colour and the luminance information are not separated from each other. Although this might be fine for most algorithms, it does pose problems in some cases. For example, Swain and Ballard [SB91] proposed a method of recognising objects based upon the object’s colour alone. Further work [FF95] done on this stated that “Swain’s algorithm is very sensitive to the lighting. Simple changes in the illumination’s intensity, let alone it’s colour, radically alter the algorithm’s results”.

With the emergence of appearance-based methods, both in vision (e.g. [ML04, NL04]) and in graphics (e.g. [SWCT02]), it is important to have models that will indeed offer the separation between colour and luminance, especially in applications involving the recognition and/or fusion of images taken under different illumination conditions.

Colour constancy is part of the problem (e.g. [FDF93]). This refers to methods of transforming colours such that objects imaged under different lighting conditions look the same. In this case, the differences in lighting conditions that are addressed correspond to differences in colour, usually produced by different types of lighting such as natural lighting, possibly at different times of the day [JMW64], and different types of artificial lights. This paper is not about colour constancy but about illumination independence.

Some work towards illumination independence has been done, often involving complex methods and/or constraining assumptions. This includes explicit illumination evaluation by analysing the lighting gradient [Hor74], but assumes Mondrian-like images. More recently, implicit illumination invariance (in the context of recognition) has been achieved by learning probabilistic illumination gradient distributions using a large training set [CBJ00].

Researchers and practitioners have used colour spaces that separate luminance from colour information. The colour model known as *HSV* (Hue, Saturation, Value) is such a model and is mainly used by graphics designers for its ease of colour specification. A less known colour model having a similar property is CIE $L^*a^*b^*$ [Fai98]. This colour model has another important property, although not of use in this work: it is perceptually linear. The luminance separation offered by CIE $L^*a^*b^*$ has been used in segmentation tasks (e.g. [Hen98]). Another widely used colour model is *YUV*. This model is used in television (and more generally video) broadcasting in Europe with the PAL format. This model has been designed such that the Y component carries the luminance information (and is being attributed more bandwidth given that the human eye is more sensitive to luminance variation), the colour information being contained in the U and V components. However, we are not aware of any experimental study of the effectiveness of such colour models to separate luminance information from colour information. This paper describes experiments we performed in an attempt at characterising this effectiveness.

The paper is organised as follows. Section 2 describes the transformations used to convert from *RGB* to *HSV*, *YUV* and CIE $L^*a^*b^*$. Section 3 describes the experimental setup while Section 4 gives the results we obtained. Section 5 gives a simple example of a possible application of the results. Finally, Section 6 concludes and proposes future experiments.

2. Transformations

The *RGB* to *HSV* and *YUV* transformations are fairly standard. The *RGB* to *HSV* transformation is as follows:

$$\begin{aligned} \min &= \min(R, G, B), \\ \max &= \max(R, G, B), \\ H &= 60 \times \begin{cases} \text{undefined} & \text{if } \max = 0; \\ \frac{G - B}{\max - \min} & \text{if } R = \max; \\ 2.0 + \frac{B - R}{\max - \min} & \text{if } G = \max; \\ 4.0 + \frac{R - G}{\max - \min} & \text{if } B = \max, \end{cases} \\ S &= \begin{cases} 0.0 & \text{if } \max = 0; \\ \frac{\max - \min}{\max} & \text{otherwise,} \end{cases} \\ V &= \max, \end{aligned}$$

where $0.0 \leq R, G, B \leq 1.0$, H is in degrees ($0.0 \leq H < 360$) and $0.0 \leq S, V \leq 1.0$. The *RGB* to *YUV* transformation is as follows (PAL version):

$$\begin{aligned} Y &= 0.299R + 0.587G + 0.114B, \\ U &= 0.492(B - Y), \\ V &= 0.877(R - Y), \end{aligned}$$

where $0.0 \leq R, G, B \leq 1.0$, $0.0 \leq Y \leq 1.0$, $-0.492 \leq U \leq 0.492$ and $-0.877 \leq V \leq 0.877$.

The *RGB* to CIE $L^*a^*b^*$ transformation is less common, and in fact involves two transformations. First *RGB* colours are transformed into CIE *XYZ*:

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} 2.7690 & 1.7518 & 1.1300 \\ 1.0000 & 4.5907 & 0.0601 \\ 0.0000 & 0.0565 & 5.5943 \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}.$$

Note that this transformation is theoretically a function of the viewing conditions (intensity and colour of the lighting) [MG87]. The given transformation is one that is often used (e.g. [Hen98, LW01]).

The CIE *XYZ* to CIE $L^*a^*b^*$ transformation is given below, and covered in depth by Fairchild [Fai98]:

$$\begin{aligned} L^* &= 116 \left(\frac{Y}{Y_n} \right)^{\frac{1}{3}} - 16, \\ a^* &= 500 \left[\left(\frac{X}{X_n} \right)^{\frac{1}{3}} - \left(\frac{Y}{Y_n} \right)^{\frac{1}{3}} \right], \\ b^* &= 200 \left[\left(\frac{Y}{Y_n} \right)^{\frac{1}{3}} - \left(\frac{Z}{Z_n} \right)^{\frac{1}{3}} \right], \end{aligned}$$

where X_n, Y_n, Z_n are the X, Y , and Z values of the reference white, $R = G = B = 1.0$. It is to be noted that this transformation is simplified in that it strictly should not be applied to low luminance levels. The complete transformation can be found in [MG87, Fai98]. The component L^* gives the luminance and we have $-16.0 \leq L^* \leq 100.0$, from black to white. The components a^* and b^* roughly corresponds to the axes respectively from green to red and from blue to yellow. A graphical depiction of the mapping from *RGB* to CIE $L^*a^*b^*$ can be found in [Fai98]. These components have no theoretical limits (see Section 4 for a discussion).

3. Experimental setup

The setup used to capture images consisted of a statically mounted camera overlooking areas with a variety of interesting sub-areas to consider. The camera was used to take images at semi-regular intervals over a period of time long enough for there to be a noticeable change in illumination across the whole image set, but nothing else, in particular no significant changes in illumination colour. We largely avoided sun-rise and sun-set times as these do change significantly the illumination colour. Individual sets were captured during single days to avoid illumination colour changes as much as possible. These changes however were not totally removed, which is a cause for the observed variations in colour information, see Section 4.

We used a Picasso 104-2SQ framegrabber from Arvo and a WAT-202B camera from Watec. The camera provides three different modes of white balance: 3200K, 6300K and automatic. The precise nature of the white balance functions of the camera are undocumented, and for this reason the sub-areas shown in Figure 2 actually has two sets of images, the

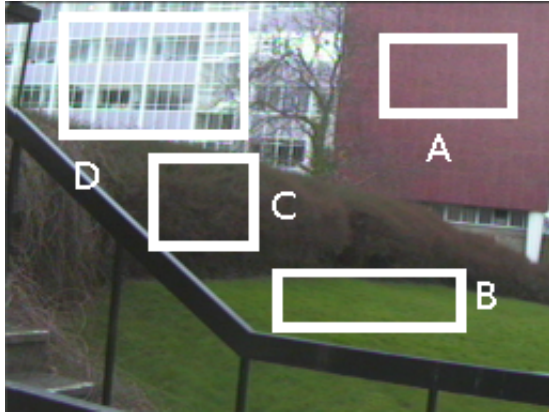


Figure 1: The first set of images

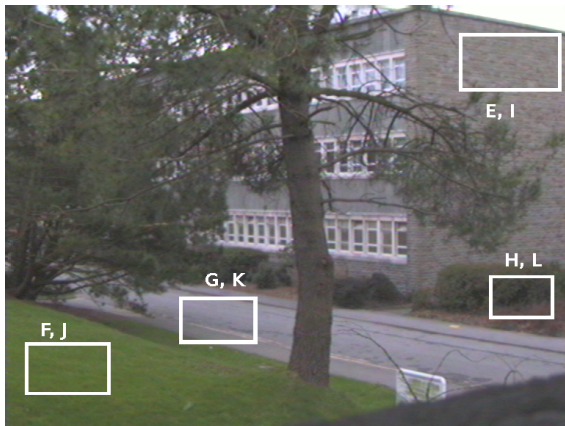


Figure 2: The second set of images. The sub-areas corresponding to the first letter of each pair are with automatic white balance, the others are fixed to "natural light".

first (E, F, G, H) with automatic white balance and the second set (I, J, K, L) with the "natural light" white balance setting. This enabled us to consider the effects of white balancing upon the transformations discussed previously. Moreover, the automatic lighting adjustment was turned off but the aperture of the lens was adjusted between image sets to ensure that the images would not be saturated (some days were much brighter than others). Despite this, a few sub-areas of a few images were saturated, as discussed later.

The sub-areas selected for consideration are shown in Figures 1 and 2. These were selected primarily because they are less subjected to changes for reasons other than illumination. None of the sub-areas selected include for instance any sky, since the colour of sky is liable to change with the weather and in particular passing clouds. Another element of the sets largely avoided corresponded to the trees owing to the potential for movement between images caused by wind.

With any image capture process a certain amount of noise within the image is inevitable. This unavoidable problem

Table 1: Properties of the different sub-areas: number of images and size (in pixel)

Sub-area	A	B	C	D	E	F
Number	11	11	11	11	24	24
Size	1650	1600	1548	1980	9100	6572
Sub-area	G	H	I	J	K	L
Number	23	24	34	34	34	32
Size	5088	3950	9100	6572	5088	3950

Table 2: Some correspondances between RGB and CIE $L^*a^*b^*$

RGB	CIE $L^*a^*b^*$			colour
(1.0 0.0 0.0)	49.126	113.48	112.29	red
(0.0 1.0 0.0)	92.238	-128.15	143.53	green
(1.0 1.0 0.0)	99.587	-34.057	156.20	yellow
(0.0 0.0 1.0)	9.5101	182.43	-155.35	blue

makes performing direct pixel by pixel comparisons of any sub-area difficult and unreliable. There are however ways of mitigating the effects of this noise. Assuming that the noise function is likely to be approximately Gaussian, it is possible to use the fact that the mean noise of a large enough set of pixels should be 0 to our advantage. To this end the mean values of each channel of each sub-area have been used in analysis of the captured images, thus the pixels in each sub-area produce one triplet for each input image.

Having carefully selected suitable sub-areas as described above it was still necessary to discard several frames from the sets due to pedestrians and cars. These images were identified and removed manually.

Figure 3 shows some of the sub-areas of some of the image sets. They are indicative of the range of illumination conditions present across the image sets. The nature of the selected areas results in the sub-areas passing through patches of shade, and the time period of capture resulted in significant perceivable changes across the image sets. Table 1 gives the number of images for each sub-area as well as their size.

4. Results

In order to facilitate the comparison, each axis in each colour space was normalised to the interval $[0;1]$. In the case of the a^* and b^* components of CIE $L^*a^*b^*$, which can in theory range from $-\infty$ to $+\infty$, the lowest and highest CIE $L^*a^*b^*$ values that could be produced by transforming valid RGB colours were used as the range, Table 2. However, impossible colours in CIE $L^*a^*b^*$ can be transformed into RGB colours and provide a much wider range for the a^* and b^* components. For example black-red in CIE $L^*a^*b^*$ $(-16.0, 375.0, 0.0)$ can be transformed into $(0.996, 0.0, 0.004)$ in RGB. This means that the performed

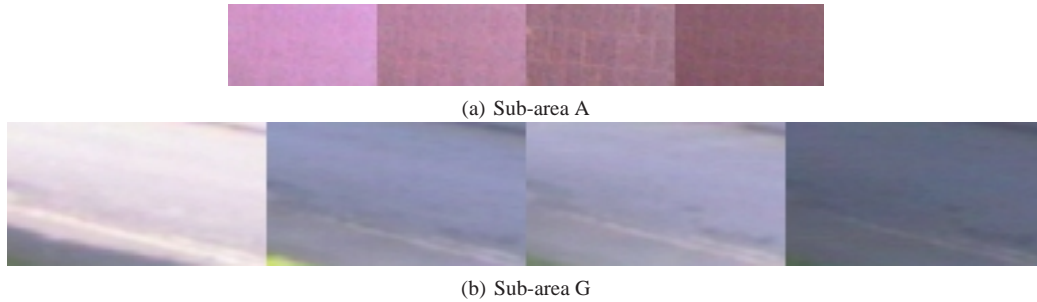


Figure 3: Typical sub-areas of the different image sets

normalisation is conservative and could have reduced even more the obtained variances (see later).

Tables 3, 4, 5 and 6 give the mean and variance over all images of the corresponding sets for the different sub-areas and colour spaces. Direct comparison of the variance of components of the CIE $L^*a^*b^*$ and RGB colour models are not statistically valid comparisons because the transformation from RGB to CIE $L^*a^*b^*$ is not linear, which means they have different distributions. However, for our purposes, simple comparison, such as the order of magnitude, do present meaningful evidence.

From these tables, it is clear that the different colour spaces provide different degrees of separation of luminance from colour. R , G and B variances are all similar for given sub-areas, showing that the illumination being the only varying factor, colour and luminance are not separated. In HSV , the variance of the hue is often, **but not always**, lower than the variance of the saturation and value (which are supposed to convey most of the luminance information). Some of these unexpected results are discussed below. In both the YUV and CIE $L^*a^*b^*$ colour spaces, the variance of the luminance component (respectively Y and L^*) is **always** at least one (if not two or three) orders of magnitude higher than the variance of the colour components (respectively UV and a^*b^*). This shows that in both models the colour components carry significantly less luminance information than the luminance component. Moreover, the two models perform similarly.

It is important to remember that some of the variation in the colour components is due to changes in the lighting colour which is unavoidable in outdoors images over long periods of time.

Some details do require closer analysis. Figures 4 and 5 show the graphs of the components in RGB , HSV , YUV and CIE $L^*a^*b^*$ for two of the sub-areas. Sub-area F shows points within the hue component of HSV that are slightly out of range. Figure 6 shows the sub-area F in different images: 2 and 3 are the anomalous images, 4 and 5 are the subsequent “normal” images. At first sight the cause of the reason for images 2 and 3 might appear obvious (cast shadows), however closer inspection of the whole set revealed that images 8, 9, 10 and 11 (Figure 6 (bottom)), which appear to fit the straight line trend, have visual characteristics similar to

Table 7: Mean hue values for sub-area F for some of the images of the set

Image	2	3	4	5
Mean H value	66.86	65.86	84.72	80.81
Image	8	9	10	11
Mean H value	77.26	75.50	74.48	73.92

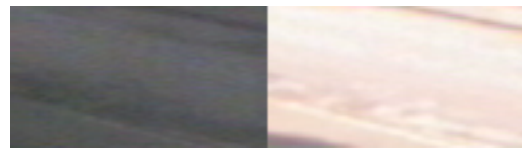


Figure 7: Two extremes of sub-area K (images 0 and 14), one showing saturation

the points which don’t fit in. Table 7 shows the actual hue values of the images on Figure 6. Visually the two out-of-line images have a more yellow tint, which is confirmed by the slightly higher values of the component b^* and the hue value close to yellow (60) for these images. These points are however not very out-of-line, especially considering the very different appearances of the different images due to shadowy and non-shadowy images.

The graph of the HSV values of sub-area K (Figure 5(b)) however do show out-of-line values. This is evident when comparing the variance in that sub-area with other sub-areas: it is an order of magnitude higher (this also happens to a lesser extent for sub-areas H and I, although for the former the reason is probably due to changes of the scene — motion in the bushes). Upon inspection of the image set we discovered that the five points which seem completely anomalous are around midday on a sunny day. Figure 7 shows images 0 and 14 showing two extremes of luminance. The RGB values of the image are close to 1.0, showing saturation of the camera, which caused problems determining the hue accurately. It is interesting to note however that the U and V components of YUV and a^* and b^* components of the CIE $L^*a^*b^*$ images (Figures 5(c) and 5(d)) at the same time fit the straight line we would expect to see if they were invariant.

Also of interest with sub-area K is its comparison with the same sub-area with automatic white balance enabled (sub-

Table 3: Comparison of channels of RGB images

Sub-area	R		G		B	
	mean	var	mean	var	mean	var
A	0.621406	0.00939583	0.449919	0.00499542	0.531759	0.013032
B	0.465028	0.0270994	0.582289	0.0303846	0.310973	0.0343862
C	0.355393	0.00972204	0.333366	0.00816249	0.303951	0.00781691
D	0.74851	0.00671774	0.814118	0.00353575	0.859049	0.0015707
E	0.582263	0.0026201	0.580631	0.00251443	0.616985	0.00488161
F	0.524702	0.0137953	0.598824	0.0172351	0.310418	0.00203099
G	0.557145	0.0234549	0.556644	0.0186704	0.632994	0.0157839
H	0.307693	0.00222363	0.303001	0.001782	0.312923	0.00185453
I	0.571447	0.00568363	0.537771	0.00424631	0.514765	0.0041639
J	0.466192	0.0122142	0.503468	0.0116693	0.264127	0.00141135
K	0.576415	0.0293905	0.543128	0.0222896	0.559224	0.0175174
L	0.305581	0.00133488	0.290489	0.000986616	0.267991	0.00106279

Table 4: Comparison of channels of HSV images

Sub-area	H		S		V	
	mean	var	mean	var	mean	var
A	0.923205	0.000716362	0.275952	0.000210198	0.621449	0.00941345
B	0.244905	0.000365237	0.477069	0.0153586	0.582307	0.0303785
C	0.170476	0.00147369	0.150291	0.00204692	0.356686	0.00966549
D	0.540148	0.0013136	0.14216	0.00295646	0.86559	0.00162167
E	0.624378	0.00350893	0.0749948	0.0010004	0.62101	0.00448823
F	0.213306	0.000195616	0.466489	0.00404809	0.598971	0.0172324
G	0.672254	0.00166297	0.147646	0.00310416	0.635455	0.0165073
H	0.543851	0.0107882	0.0851954	0.000257043	0.323045	0.00207663
I	0.224351	0.0209946	0.102828	0.00248228	0.572856	0.00553873
J	0.200177	0.000766661	0.462375	0.0100005	0.504041	0.0116285
K	0.581517	0.0373274	0.0889594	0.000405788	0.590098	0.0264975
L	0.226809	0.00781674	0.134189	0.00184084	0.308348	0.00115915

area G). Although not immediately comparable, the images having been grabbed at different times, the means and variances in *RGB* for the sub-areas E and I (and similarly for F-J, G-K and H-L) are very similar. Given that the automatic white balance is performed in the *RGB* space, this tends to show that the sub-areas are comparable. Looking at the *YUV* and CIE $L^*a^*b^*$ values, again the means and variances are also similar for the three components. However, the means and variances of the hue and saturation in the *HSV* colour space are very different (up to several orders in magnitude for the variance and a factor of 2 for the mean). The value however is comparable. This shows that the properties of the automatic white balance, performed in the *RGB* colour space, are preserved by the transformations to *YUV* and CIE $L^*a^*b^*$ but not by the transformation to *HSV*, at least the colour part of it.

These results show that the *YUV* and CIE $L^*a^*b^*$ colour spaces do separate efficiently luminance from colour, even in presence of saturation. In the next section, we briefly present an application for this result.

5. Application: appearance-based matching

In this section, we outline a typical use of the results. Many appearance-based methods (e.g. [NL04]) involve computing the Euclidean distance between images in the image space (the $(n \times m)$ -dimensional space of all n by m pixels images). As mentioned in Section 1, the main problem behind this is that if one is not careful, two views of the same object under different illumination situations might appear very different.

We show here that computing the Euclidean distance between images in *RGB*, *YUV* (just considering components U and V) and CIE $L^*a^*b^*$ (just considering components a^* and b^*) can lead to very different outcomes. Figure 8 shows the Euclidean distance between the first image of sets 2 (with natural light white balance) and 1 in the full *RGB* space and the *YUV* and CIE $L^*a^*b^*$ spaces without considering the luminance information (and normalising the components). This shows that recognition in *RGB* would fail while it would succeed even with a simple threshold-based method: the distance between the first image of sets 1 and 2 are respectively 339.169, 51.8382 and 75.6715 in *RGB*, *UV* and a^*b^* , which makes the former less discriminant than

Table 5: Comparison of channels of YUV images

Sub-area	Y		U		V	
	mean	var	mean	var	mean	var
A	0.510524	0.00693963	0.510618	0.000259126	0.555441	5.45659e-05
B	0.516298	0.0288678	0.397338	0.00135369	0.474365	7.0001e-05
C	0.336599	0.00853715	0.483676	5.57209e-05	0.509397	1.90681e-05
D	0.799623	0.00405117	0.529713	0.000197122	0.474443	9.11961e-05
E	0.585263	0.00270849	0.515861	0.000158297	0.4985	2.51616e-05
F	0.543783	0.0134644	0.383317	0.00167873	0.49046	6.69464e-05
G	0.565498	0.0196081	0.533748	0.000177854	0.495824	6.01723e-05
H	0.305535	0.00188744	0.503694	4.16806e-05	0.501079	1.20192e-05
I	0.545217	0.00454084	0.484774	0.000110528	0.513115	5.57189e-05
J	0.465038	0.0098817	0.399545	0.00144044	0.500577	4.80288e-05
K	0.554915	0.0236337	0.502154	0.000190761	0.51075	0.000106009
L	0.292437	0.0010653	0.487777	3.31153e-05	0.506572	1.37353e-05

Table 6: Comparison of channels of CIE $L^*a^*b^*$ images

Sub-area	L^*		a^*		b^*	
	mean	var	mean	var	mean	var
A	0.781374	0.00160439	0.469862	6.19425e-05	0.483003	0.000103588
B	0.813247	0.00628527	0.337735	0.000378381	0.59613	0.000808511
C	0.689234	0.00335536	0.414029	6.66229e-06	0.513768	4.00805e-05
D	0.925527	0.000631121	0.404477	8.10288e-06	0.485153	4.56153e-05
E	0.833681	0.000558209	0.418187	7.89346e-06	0.488326	5.54222e-05
F	0.829317	0.00371697	0.350987	0.000115615	0.596781	0.000571085
G	0.817236	0.00426067	0.424017	6.298e-06	0.475042	8.572e-05
H	0.669865	0.000936666	0.416159	2.6361e-06	0.494474	3.71124e-05
I	0.814373	0.0011698	0.417491	5.70455e-06	0.507978	4.52581e-05
J	0.786547	0.00298242	0.362291	0.000129361	0.591734	0.000689276
K	0.811956	0.00493214	0.423138	1.19296e-05	0.494532	7.0182e-05
L	0.662119	0.000593182	0.412932	6.54945e-06	0.51105	3.39147e-05

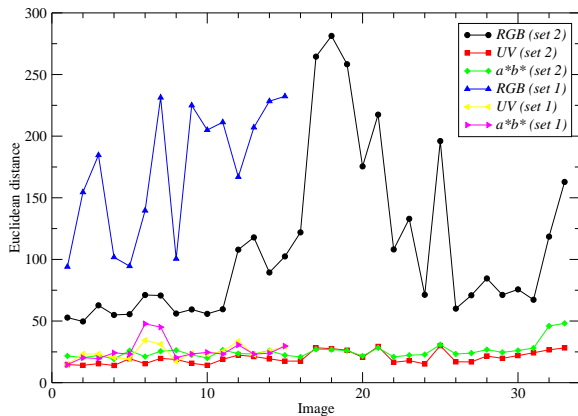


Figure 8: The Euclidean distance in RGB and CIE $L^*a^*b^*$ (without luminance information) between the first image of sets 2 (with natural light balance) and 1 and the remaining images of the same set

the others. The graph of the distance in a^*b^* for set 1 shows two points that are out-of-line. They correspond to two im-

ages showing extreme saturation or a passer-by (and were not included in the previous statistical analysis) and thus are different from the test image. The two final images of set 2 also show anomalies: the images were taken late in the afternoon and thus were becoming significantly darker (and the last but one contained a passing car). These anomalies do not appear in the YUV colour space.

6. Conclusion

We have shown that overall the a^* and b^* [resp. U and V] components of the CIE $L^*a^*b^*$ [resp. YUV] colour space have a variance at least one order of magnitude smaller than the corresponding L^* [resp. Y] component, even in the case of saturated images. (Remember that the normalising ranges we used were conservative and that possibly more realistic ranges would have lead to even lower variations in a^* and b^* .) In other words, both YUV and CIE $L^*a^*b^*$ colour models are good at separating luminance information from colour information.

This independence from luminance obviously does not

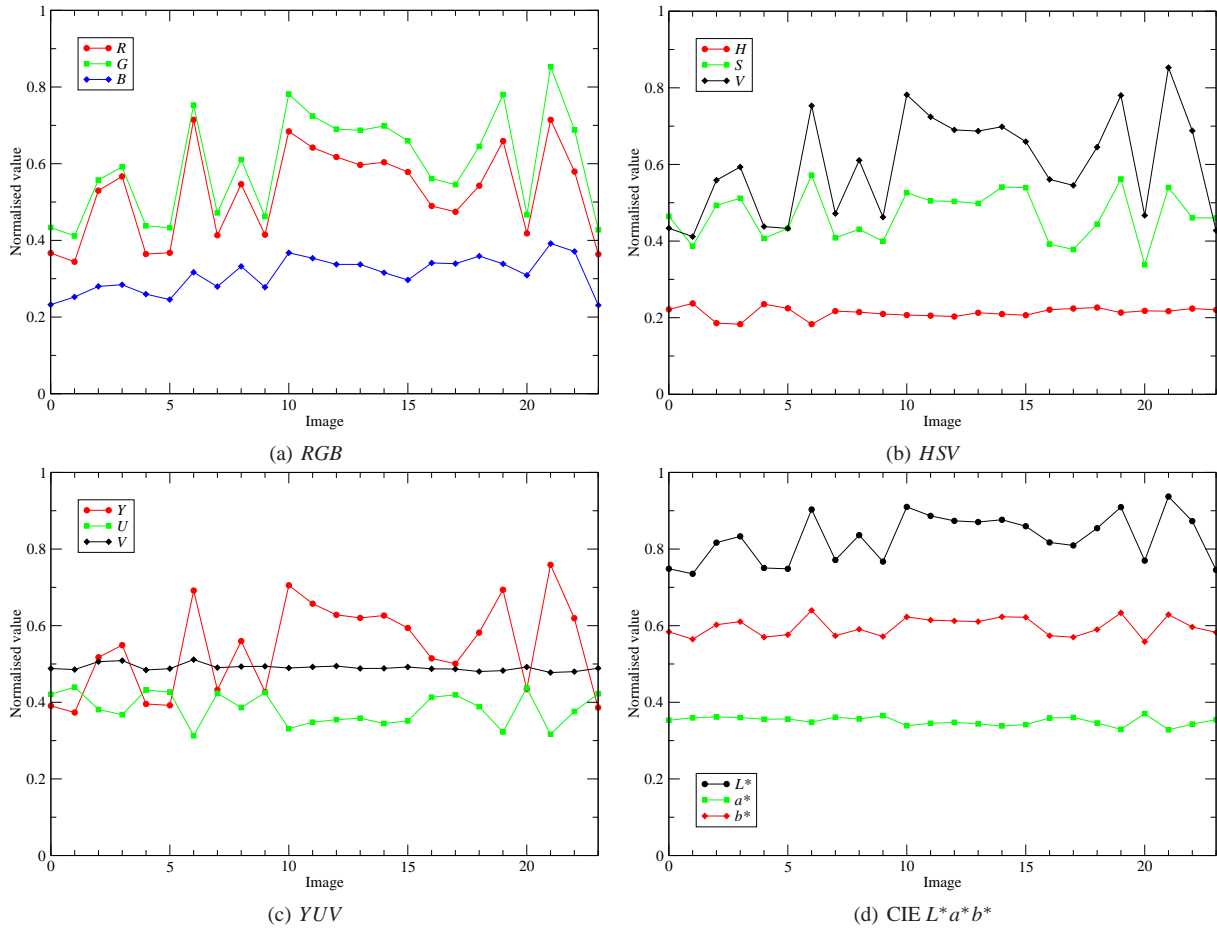


Figure 4: Results of sub-area F, Figure 2

hold in the *RGB* colour space. More surprisingly, this also does not hold in the *HSV* colour space (the variance in the hue is not consistently lower than that of the saturation and value). Moreover, we have seen that the *HSV* transformation destroys the effect of automatic white balance.

The relative simplicity of the transformation from *RGB* to CIE $L^*a^*b^*$ and even more so to *YUV* makes it worthwhile considering these models in applications needing luminance invariance. In particular, we have shown that appearance-based matching would work better if only considering the a^* and b^* components compared to using *RGB*.

Future work will include a study of the performance of appearance-based methods in robot mapping and navigation under varying lighting conditions.

References

[CBJ00] CHEN H. F., BELHUMEUR P. N., JACOBS D. W.: In search of illumination invariants. In *Proc. of the IEEE Conf. on Comp. Vis. and Pat. Rec.* (2000), pp. 254–261.

[Fai98] FAIRCHILD M. D.: *Color Appearance Models*. Addison-Wesley, 1998.

[FDF93] FINLAYSON G. D., DREW M. S., FUNT B. V.: Diagonal transforms suffice for color constancy. In *Proc. of the Int. Conf. on Comp. Vis.* (1993), pp. 164–171.

[FF95] FUNT B. V., FINLAYSON G. D.: Color constant color indexing. *IEEE Trans. on Pat. Anal. and Machine Intel.* 17, 5 (1995), 522–529.

[Hen98] HENRICSSON O.: The role of color attributes and similarity grouping in 3-D building reconstruction. *CVIU* 72, 2 (1998), 163–184.

[Hor74] HORN B. K. P.: Determining lightness from an image. *CVIP* 3 (1974), 277–299.

[JMW64] JUDD D., MACADAM D., WYSZECKI G.: Spectral distribution of typical daylight as a function of correlated color temperature. *J. of the Optical Society of America A* 54 (1964), 1031–1040.

[LW01] LABROSSE F., WILLIS P.: Towards continuous image representations. In *Proc. of the Int. Conf. in Central*

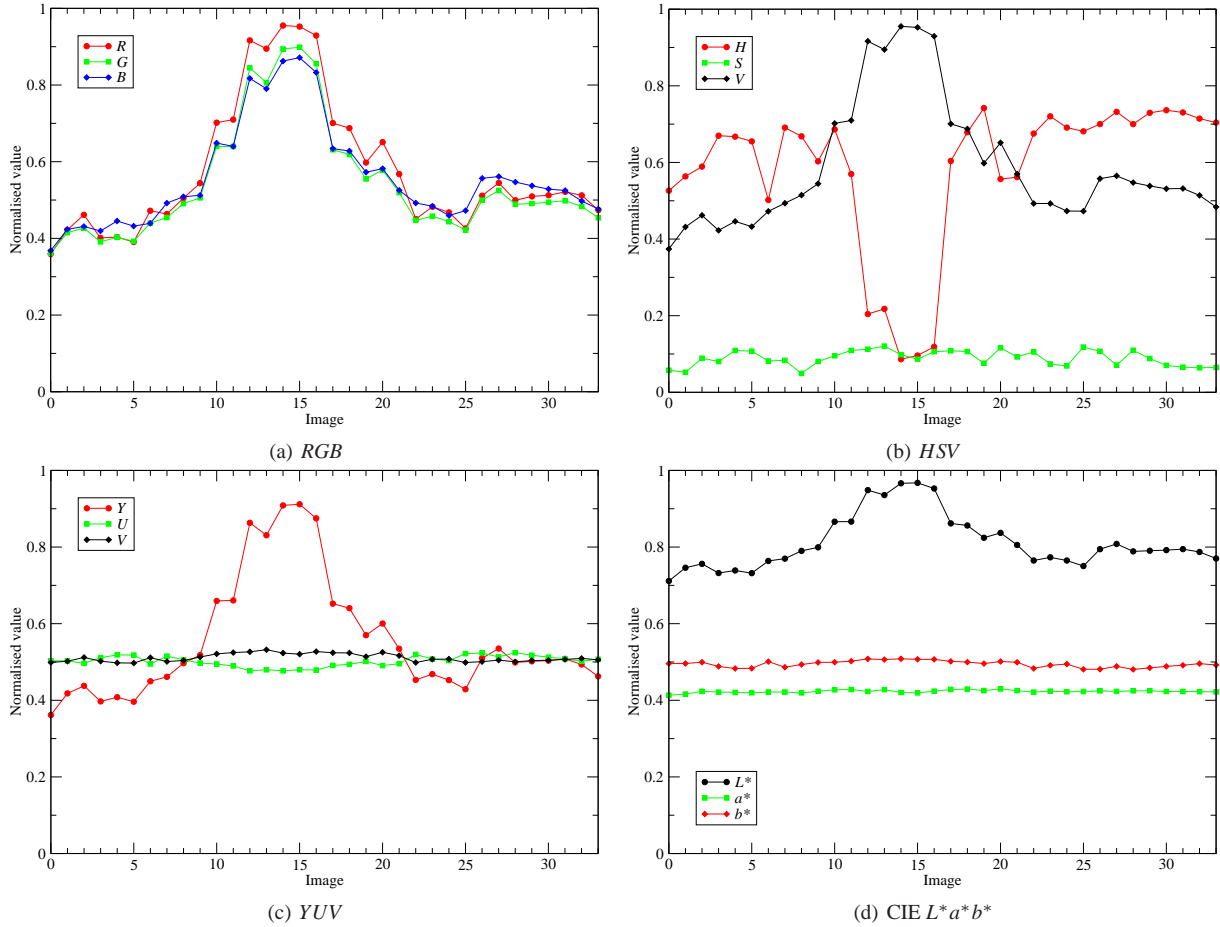


Figure 5: Results of sub-area K, Figure 2

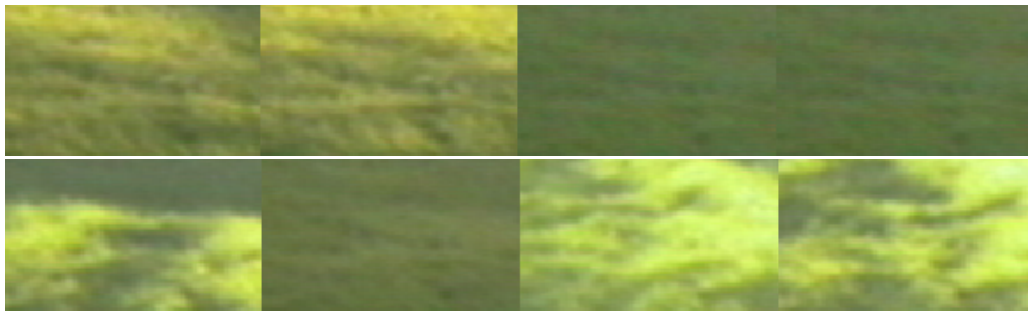


Figure 6: Images 2, 3, 4 and 5 (top) and 8, 9, 10 and 11 (bottom) from sub-area F

Europe on Comp. Graph., Visualization and Comp. Vis. (WSCG) (2001), pp. 206–213.

[MG87] MEYER G. W., GREENBERG D. P.: Perceptual color spaces for computer graphics. In *Color and the Computer*, Durrett H. J., (Ed.). 1987, pp. 83–100.

[ML04] MITCHELL T., LABROSSE F.: Visual homing: a purely appearance-based approach. In *Proc. of Towards Autonomous Robotic Systems* (2004), pp. 101–108.

[NL04] NEAL M., LABROSSE F.: Rotation-invariant appearance based maps for robot navigation using an artificial immune network algorithm. In *Proc. of the Congress on Evolutionary Computation* (2004), vol. 1, pp. 863–870.

[SB91] SWAIN M. J., BALLARD D. H.: Color indexing. *Int. J. of Comp. Vis.* 7, 1 (1991), 11–32.

[SWCT02] SHUM H.-Y., WANG L., CHAI J.-X., TONG X.: Rendering by manifold hopping. *Int. J. of Comp. Vis.* 50, 2 (2002), 185–201.