Hue that is invariant to brightness and gamma

Graham Finlayson and Gerald Schaefer
School of Information Systems, University of East Anglia
Norwich NR4 7TJ, United Kingdom
{graham,gerald}@sys.uea.ac.uk

Abstract

Hue provides a useful and intuitive cue that is used in a variety of computer vision applications. Hue is an attractive feature as it captures intrinsic information about the colour of objects or surfaces in a scene. Moreover, hue is invariant to confounding factors such as illumination brightness. However hue is not stable to all of the types of confounding factors that one might reasonably encounter. Specifically, the RGBs captured in images are sometimes raised to the power gamma. This is done for two reasons. First, to make the images suitable for display (since monitors have an intrinsic non-linearity). Second, applying a gamma is the simplest way to change the contrast in images. It has also been observed that digital cameras often apply a scene dependent gamma type function (which is unknown to the user).

In this paper we show that a simple photometric ratio in log RGB space cancels both brightness and gamma. Furthermore, some simple manipulation reveals that the brightness/gamma invariant can usefully be interpreted as a hue in a log opponent colour space. We carried out indexing experiments to evaluate the usefulness of the derived hue correlate. In situations where gamma is held fixed, the new hue supports recognition equal to conventional definitions. In situations where gamma varies the new correlate supports better indexing. The new hue is also found to predict some psychophysical data quite accurately.

1 Introduction

According to the International Commission on Illumination (CIE), *hue is the attribute of a visual sensation according to which an area appears to be similar to one of the perceived colours, red, yellow, green and blue, or a combination of two of them* [5]. In more practical terms, hue is the 'name' of a colour. It is also the property of colour that people find is easiest to use.

Computer vision, in trying to mimic human's abilities, has found hue to be useful in various applications. These include [14] where a colour segmentation algorithm based on hue only is introduced. [12] presents a hue based approach to suppress the effects of cloud shadows for remote sensing applications. It has been shown that (assuming the light source of a scene is white) hue does not change in the presence of specularities [8]. Colour transformations like HSV and HLS that convert image RGB values to a hue based representation [17] allow not only for a more intuitive description of colour but can also

BMVC 2001 doi:10.5244/C.15.32

be used in applications such as object recognition [19, 8] and face tracking [16]. Colour naming, the division of colour space into regions identified by colour names, is closely linked with the concept of hue, and has been successfully used in image retrieval [13] and visual surveillance [2].

Remarkably, the human vision system can ascribe fairly constant hues to surfaces viewed in different visual contexts. Looking at the mathematical definitions of hue used in computer vision it is easy to show that CV hue is invariant to brightness. Shaded surfaces or surfaces viewed under different powers of illumination have the same intrinsic hue. But we find that other confounding factors do not cancel out.

A gamma different from 1 implies that there exists a power function relationship between scene intensities and pixel values [15]. So device responses denoted as (R, G, B) will become pixel RGBs described as $(R^\gamma, G^\gamma, B^\gamma)$. It is often wrongly assumed in computer vision that $\gamma = 1$ or that γ might be turned off and that the default resting state is $\gamma = 1$. Yet, our practical experience has shown that this is rarely the case. Indeed, it has been found that some low end digital cameras apply a gamma that depends on scene content [6].

Non unity gamma is needed because the colours that are displayed on a screen are not a linear function of the RGBs sent to the monitor. Rather, there exists a power function (or gamma) relationship between the incoming voltage and the displayed intensity. A linear image displayed on screen will look too dark and lack contrast. This is because midtones get attenuated by the gamma function in comparison to dark and light pixel values [15]. To compensate for this, images are usually stored in a way that reverses the effect of the monitor. This can be achieved by applying a gamma function with the reciprocal value of the monitor gamma as exponent. Usually this normalization takes place directly at the stage of image acquisition, i.e. in the device. It should also be noted that, as different monitors have different gammas (e.g. the ”standards” for PC and Macintosh are 2.2 and 1.8 respectively) images with different γ values are a consequence. Another reason for applying a non unity gamma is to change the contrast of an image e.g. as a preprocessing step prior to other tasks such as segmentation.

In this paper we look at hue in the context of changing brightness and contrast. As common hue descriptors are invariant only to brightness, we are seeking a description of hue that remains constant after a change in brightness and/or contrast (either device related or through post-processing). Given an RGB and possible brightness and gamma dependencies we show how some simple manipulation in log RGB space can cancel both factors. Moreover, the manipulation might usefully be interpreted as a hue correlate. Specifically we show that if hue is defined as the angle between two log opponent color coordinates (a red-green and a yellow-blue coordinate) then hue is brightness and gamma independent. This definition of hue naturally falls out of the algebra involved in cancelling brightness and gamma. Yet, the algebra leads to a definition that meshes well with definitions used in colour science [4].

To test the utility of the gamma invariance we scanned a dataset of 27 design images using two different gamma settings: $\gamma = 1$ and $\gamma = 2.2$. We show that corresponding hue images as obtained from the HSV colour model do indeed differ due to the different gamma settings. However, images based on our newly defined hue space look very similar. To quantitatively assess hue stability, we performed colour indexing [20] experiments. Is it possible, using hue content alone, to match images across the two different scanning settings? We found that this was the case for our new hue correlate but that HSV and HLS failed (as we might expect).

Of course this is a very simple experiment which is really only useful to show that the derived invariant is fairly stable to gamma changes. However, we wished to look to its general applicability even in situations where a varying gamma is not a problem. In a second experiment we indexed into a large set of (around 4000) images and found that the hue descriptor alone supported good performance. Moreover, our new hue correlate delivered similar performance to conventional descriptors.

In a final experiment we wondered whether our new hue descriptor has relevance to our own visual system. We found that the derived correlate can predict the constant hue lines derived in psychophysical experiment more accurately than the HSV colour space.

Before proceeding we point out to the reader that brightness and gamma invariance will not render hue appropriate to all imaging situations. If the colour of the light that illuminates the scene changes, so do the colours in the scene as recorded by a device such as a digital camera [7]. In terms of hue this manifests itself as a global shift of object hues towards the colour of the light source. This is true for our new hue definition as it is for HLS and HSV. Henceforth we will assume that the colour of the illuminant has been discounted from images.

The rest of the paper is organised as follows: Section 2 briefly explains the process of image formation, defines conventional hue based colour spaces and shows that hue in these colour models changes with a change in contrast. Section 3 introduces our log hue space. Section 4 describes the experiments we performed to demonstrate the validity and usefulness of the newly defined hue. Section 5 concludes the paper.

2 Background

2.1 Image Formation

A linear device captures colour, or R, G and B, according to:

$$\underline{p}_{lin}(x) = \underline{e}(x) \cdot \underline{n}(x) \int_{\omega} S(x,\lambda) E(\lambda) \underline{R}(\lambda) d\lambda \tag{1}$$

where λ is wavelength, $\underline{p}_{lin}$ is a 3-vector of sensor responses (RGB pixel values), S^x is the surface reflectance at location x, E the spectral power distribution of the illumination, and $\underline{R}$ is the 3-vector of sensitivity functions of the device. Integration is performed over the visible spectrum ω. The light reflected at x is proportional to $S^x(\lambda)E(\lambda)$, its magnitude is determined by the dot product $\underline{e}^x \cdot \underline{n}^x$ where $\underline{e}^x$ is the unit vector in the direction of the light source, and $\underline{n}^x$ is the unit vector corresponding to the surface normal at x. In the strictest sense Equation (1) only describes the response of Lambertian (matte) reflectances. But, in practice it is a tolerably good model for most surfaces (even those that have some highlight component).

Equation (1) is the starting point for most colour based algorithms used in computer vision. Yet, in practice (outside the domain of computer vision) linear camera response is not the norm. Rather, camera response is the linear RGB response raised to some γ (gamma) power:

$$\underline{p}(x) = (\alpha \, \underline{p}_{lin}(x))^{\gamma} \tag{2}$$

where the scalar α models the interaction of surface and illumination normals and the intensity of the light source. To simplify notation still further, let us denote the linear RGB responses (Equation (1)) with $(R,\ G,\ B)$. Equation (2) can be then written as

$$\underline{p}(x) = ((\alpha R)^\gamma,\ (\alpha G)^\gamma,\ (\alpha B)^\gamma) \tag{3}$$

There are two reasons for non linear camera responses. First, colour monitors have a non linear transfer function: PC monitors apply a power of 2.2 to the signals (RGBs) driving the display. It follows that in order to achieve a true (physically accurate) colour signal, RGBs must be raised to the power of $\gamma = 1/2.2$. Monitors tied to Apple computers apply a power of 1.8 prior to display. The implication of this is that cameras calibrated to PCs and Apples require different gamma settings.

The second reason for a non-unity gamma is to change the contrast of an image. Applying a gamma larger than 1 tends to compress the signal range in the bright area of images but to bring out detail in the darker regions. Conversely a gamma of less than 1 brings out detail in bright areas but compresses the signal in darker image regions. To a first approximation most images from unknown sources (e.g. images downloaded from the Internet) can be considered to be linear after an appropriate (but unknown) gamma correction. Contrast adjustments are also made as a simple form of dynamic range compression (mapping the larger physical range (16 to 20 bit) of intensities to the 8-bit range of typical cameras). Experiments have shown that some cameras will adjust contrast in a scene dependent way without user intervention [6].

While appropriate brightness and gamma adjustments might be made (to achieve a linear image) in a calibrated lab environment, this is not in general possible. It is however reasonable to ask whether some property of colour (which is defined by the three numbers R, G and B) might be independent of the two confounding factors (α, γ).

2.2 Hue based colour models

The simplest single number used to define colour is 'hue'. Hue correlates to the colour name we might use to classify a surface (red, green, pink etc). Hue might be calculated using any of the colour spaces: HSV, HLS and IHS [17]. Indeed, all are used in image processing and computer vision. Based, on perceptual studies of how we see colour, each of these spaces codes RGB by three 'perceptual' correlates: hue, saturation and brightness. Brightness correlates to magnitude: white is brighter than grey. Saturation measures the purity of colour: a whitish pink is more desaturated than a saturated red (yet both these may have the same brightness and hue).

Though HSV, HLS and IHS differ in the way they define saturation and intensity, their definition of hue is the same. Hue is defined as [17, 9]

$$H = \cos^{-1} \frac{0.5[(R-G)+(R-B)]}{\sqrt{(R-G)(R-G)+(R-B)(G-B)}} \tag{4}$$

To increase efficiency several simpler definitions were developed [17]. However, they are also based on Equation (4) and give numerical results that differ only slightly.

We now want to inspect what effect changes in brightness and contrast, as modelled by Equation (3) have on hue as defined above. For that we substitute the RGBs from

Equation (3) into Equation (4)

$$H = \cos^{-1} \frac{0.5[((\alpha R)^\gamma - (\alpha G)^\gamma) + ((\alpha R)^\gamma - (\alpha B)^\gamma)]}{\sqrt{((\alpha R)^\gamma - (\alpha G)^\gamma)((\alpha R)^\gamma - (\alpha G)^\gamma) + ((\alpha R)^\gamma - (\alpha B)^\gamma)((\alpha G)^\gamma - (\alpha B)^\gamma)}} \tag{5}$$

We see immediately that it is possible to cancel the α terms leading to

$$H = \cos^{-1} \frac{0.5[(R^\gamma - G^\gamma) + (R^\gamma - B^\gamma)]}{\sqrt{(R^\gamma - G^\gamma)(R^\gamma - G^\gamma) + (R^\gamma - B^\gamma)(G^\gamma - B^\gamma)}} \tag{6}$$

That is, hue does not change with a change in brightness. This was what we might expect given the idea that HSV, HLS and IHS separate out brightness. However, we see that the γ exponent does not cancel. Thus, hue depends on the image gamma and will change when γ is altered.

3 Brightness and gamma invariant hue

For images from an unknown source (like those found on the web) the image gamma is also an unknown. Furthermore it is quite possible that two images of the same scene will be captured with a different gamma. This might be due to either the images being captured for different target systems (e.g. the gamma for a PC is 2.2, for Macintosh 1.8), or application of a different gamma to enhance the contrast of the image. This latter enhancement may be automatically applied by the camera. As shown in the last section, while hue is independent of brightness it depends on gamma.

It turns out that it is quite straightforward to derive a single scalar value from an R, G and B measurement that cancels both brightness and gamma. From observing Equation (3) we see that applying a log transform to RGBs removes the γs from the exponent and turns them into multiplicative scalars. At the same time the brightness α becomes an additive rather than a multiplicative term:

$$\begin{array}{ccc} (\log((\alpha R)^\gamma), & \log((\alpha G)^\gamma), & \log((\alpha B)^\gamma)) = \\ (\gamma \log(\alpha) + \gamma \log(R), & \gamma \log(\alpha) + \gamma \log(G), & \gamma \log(\alpha) + \gamma \log(B)) \end{array} \tag{7}$$

Taking differences of colour channels allows us to remove the brightness terms

$$\begin{pmatrix} ((\gamma \log(\alpha) + \gamma \log(R)) - (\gamma \log(\alpha) + \gamma \log(G)) \\ (\gamma \log(\alpha) + \gamma \log(R)) + (\gamma \log(\alpha) + \gamma \log(G)) - 2(\gamma \log(\alpha) + \gamma \log(B)) \end{pmatrix} =$$

$$\begin{pmatrix} \gamma \log(R) - \gamma \log(G) \\ \gamma \log(R) + \gamma \log(G) - 2\gamma \log(B) \end{pmatrix} \tag{8}$$

We note that the way we define the above differences describes coordinates in an opponent colour representation [1]. They are similar to the opponent colour axes used by the human visual system [11] and so have perceptual relevance.

Finally, ratios of the opponent colour coordinates are formed to cancel gamma

$$\frac{\gamma \log(R) - \gamma \log(G)}{\gamma \log(R) + \gamma \log(G) - 2\gamma \log(B)} = \frac{\log(R) - \log(G)}{\log(R) + \log(G) - 2\log(B)} \tag{9}$$

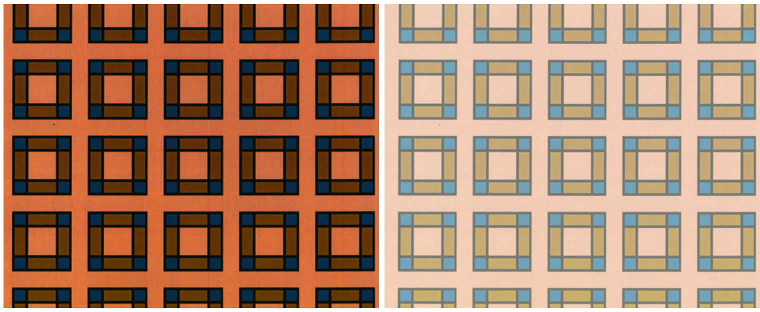

Figure 1: Example of a design image scanned using two different gamma settings: $\gamma = 1.0$ on the left, and $\gamma = 2.2$ on the right.

So, the above single scalar is independent of brightness and gamma and might be an appropriate function for use in computer vision. However, compared to the idea of hue (red, green, blue etc), this log-difference ratio does not seem so intuitive.

An alternative to removing the gamma term through ratios is to calculate the angle of the vector from Equation (8) with respect to the x-axis, i.e. we calculate the inverse tangent of the ratio in Equation (9)

$$H = \tan^{-1} \frac{\log(R) - \log(G)}{\log(R) + \log(G) - 2\log(B)} \tag{10}$$

In color science (e.g. in the CIELab space[4]), hue is defined as an angle in a red-green and blue-yellow coordinate space. Here we have shown that the simplest strategy for removing brightness and gamma dependency from RGB measurements results in an analogous hue correlate. Hue is that part of an image signal which is invariant to changes in brightness and gamma.

4 Experimental Results

To evaluate our new hue definition, we created a small image database of 27 colourful designs, each scanned in twice with different gamma settings of the scanner. In the first case the direct sensor responses were saved, i.e. no gamma was applied resulting in linear images as commonly used in computer vision applications. The designs were scanned for a second time using a gamma of 2.2. An example image pair is shown in Figure 1. The differences are obvious with the linear image appearing much darker and with less contrast. From Equation (6) we expect the hues from corresponding regions of the two images to differ due to a change in image gamma. To illustrate this we converted the original RGB images into the HSV colour space, fixed brightness (value) and saturation over the whole image, and transformed them back to RGB. The resulting images are shown in Figure 2. The difference in image colours is again quite evident. We then performed the same procedure for our new log hue space, i.e. convert the images into hue based representation, fix saturation and brightness, and convert them back to RGB. The result for the image pair from Figure 1 is given in Figure 3. Clearly, as expected, the two images look much closer to each other than is the case for the HSV based images. Here

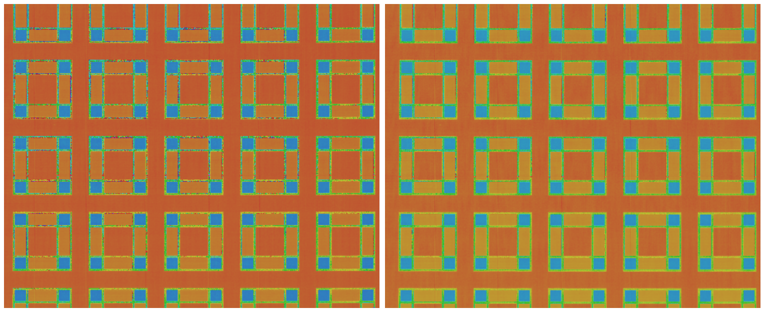

Figure 2: Hue images of the two images from Figure 1 based on the HSV model.

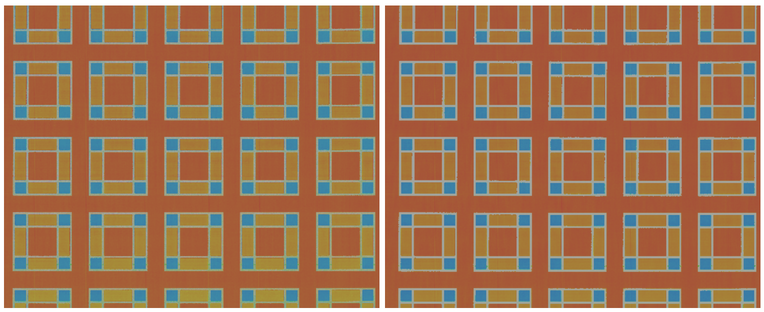

Figure 3: Hue images of the two images from Figure 1 based on the new log hue model.

we have a visual confirmation of the brightness and gamma invariance of the new hue correlate.

To better quantify hue stability, we performed colour indexing [20] experiments on the above dataset. For this we divided the 54 images into two halves, according to their gamma settings. One half was assigned model images, i.e. those images we are searching through, and the other half query images, i.e. the images that are used as input for a search. We transformed all the images into hue based representations, both the conventional one from Equation (4) and the log based hue from Equation 10. We take only the resulting hue angles and generate a 16-bin histogram by quantising the possible hue range into discrete intervals. For each query image a matching score to each of the models is calculated as the intersection of the corresponding two histograms. (Histogram intersection establishes the overlap of two histograms [20].) The retrieved images are then sorted in order of their matching score.

The results of this experiment are listed in Table 1. They are given in terms of average match percentile, the percentage of the correct images retrieved in 1st, 2nd, and 3rd rank, and the worst rank in which a corresponding image was retrieved. Average match percentile is a standard measure used in the colour indexing literature [20]. A match percentile of e.g. 99 informs us that the correct image was retrieved in the top 1% of all model images in the database. From Table 1 we see that image gamma indeed influences the matching performance based on conventional hue definition. The achieved match percentile is only about 95, moreover the recognition rate (i.e. the percentage of 1st rank retrievals) is only slightly over 50 (only half of the images being correctly identified), and the worst rank in which the correct image was retrieved is 8. For a small image database

hue model	MP	1st	2nd	3rd	worst
HSV hue	94.59	51.85	14.81	3.70	8
log hue	99.72	92.59	7.41	0.00	2

Table 1: Results of the colour indexing experiment on an image database of 27 colourful design scanned with two different gamma setting.

such as the one used in the experiment, this is clearly not good enough. In contrast, if we turn our attention to the results obtained from the log hue definition, we see immediately that here the performance is very good. The average match percentile is 99.72 which corresponds to all correct images retrieved in 1st place except for two that rank 2nd.

The above experiment illustrates that common values of gamma impact on conventional definitions of hue to the extent that indexing performance deteriorates. In a second experiment we wanted to evaluate the general utility of the new definition. Does the new hue correlate capture intrinsically useful information? Can it be used as an alternate to hue as defined in HSV (HLS or IHS)? We took a large image database comprising 4100 image triplets (this image set is similar to the one described in [3]). Each triplet consists of one original image, taken from the Corel Photo stock, and two cropped versions where $1/3$ of the image was removed either horizontally or vertically. Clipping the left and right side of images simulates portrait image. Clipping top and bottom simulates panoramic capture. We point out to readers that this clipping is exactly what happens in the APS photographic system (a full resolution image is always captured but panoramic and portrait prints can be made through clipping).

Again, we performed image retrieval on this dataset. The original images make up the model set, while the cropped images are the query images. As above, we quantised hue into 16 and 8 values and indexed only on these. The results that we obtained proved to be excellent. The average match percentile over the whole dataset (i.e. 8200 query images, 4100 model images) is 99.80 for the 16 bin and 99.22 for the 8 bin histograms, and is comparable to the performance achieved by indexing on conventional HSV hue (99.95 and 99.73 percentile respectively). This shows that hue as we have defined it in Equation (10) provides a powerful cue for object recognition. Not only that. Hue also allows for a compact representation of colour content. The amount of compression that was achieved here - an image described by 16 or 8 numbers only - is similar to other methods introduced in the literature [3, 18].

Though our angular definition of hue is similar in spirit to those used in color science (which are designed to model perceptual response), we wished to examine this relationship in more detail. In Figure 4 we have plotted lines of constant hue in the colour space defined by the opponent log colour axes from Equation 8. The data for this plot is taken from from [10] and was derived through psychophysical experiment. The lines thus connect points human observers judged to have the same hue. We can see that, with the exception of one hue locus in the blue region, all the lines are fairly straight and hence correspond well with the human visual system. In fact, our newly derived hue space is in better agreement with psychophysical data than conventional hue spaces. This is demonstrated in Figure 5 where the same hue lines as in Figure 4 are plotted in the hue-saturation plane of the HSV colour space. Here, clearly the lines appear more curved.

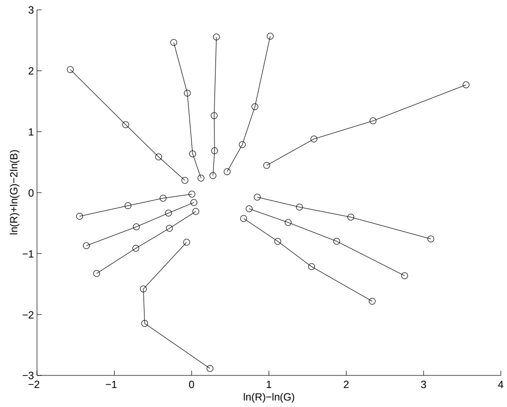

Figure 4: Lines of constant hue plotted in the new log hue space.

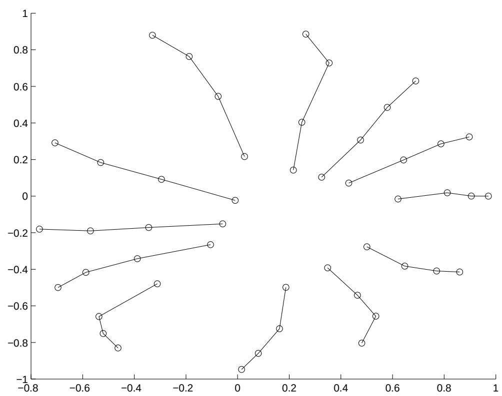

Figure 5: Lines of constant hue plotted in HSV hue-saturation plane.

5 Conclusions

We have demonstrated that hue, as conventionally defined, is not a stable cue when the gamma of images changes. To overcome this we have derived a new definition of hue that is invariant to image gamma and brightness. Invariance is achieved through a transform to a 2-dimensional colour log opponent coordinate system. With respect to polar coordinates, hue is the angle of opponent colours.

Experiments demonstrate that this hue definition indeed outperforms classical hue spaces (HLS, HSV and IHS) when the image gamma is not held fixed. Moreover, it works as well as conventional definitions when gamma does not vary. Experiments also demonstrate that the new hue correlate appears to be more perceptually relevant than conventional measures used in computer vision.

Acknowledgements

This work was funded by EPSRC grant GR/L60852.

References

[1] D.H. Ballard and C.M. Brown. *Computer Vision*. Prentice Hall, 1982.

[2] R. Benavente, G. Sánchez, R. Baldrich, M. Vanrell, and J. Lladós. Normalized colour segmentation for human appearance description. In *15th IEEE Int. Conf. on Pattern Recognition*, 2000.

[3] J. Berens and G. Finlayson. Log-opponent chromaticity coding of colour space. In *15th IEEE Int. Conf. on Pattern Recognition*, pages 206–211, 2000.

[4] CIE. *Colorimetry*. CIE Publications 15.2, Commission International de L'Eclairage, 2nd edition, 1986.

[5] CIE. *International Lighting Vocabulary*. CIE Publications 17.4, Commission International de L'Eclairage, 4th edition, 1989.

[6] G. Finlayson and G. Schaefer. Colour indexing arcoss devices and viewing conditions. In *2nd International Workshop on Content-Based Multimedia Indexing*, 2001.

[7] D.A. Forsyth. A novel algorithm for color constancy. *Int. J. Computer Vision*, 5:5–35, 1990.

[8] T. Gevers and A.W.M. Smeulders. Color based object recognition. *Pattern Recognition*, 32:453–464, 1999.

[9] R.C. Gonzales and R.E. Woods. *Digital Image Processing*. Addison Wesley, 1992.

[10] P. Hung and R. Berns. Determination of constant hue loci for a crt gamut and their predictions using color appearance spaces. *Color Research Applications*, 20:285–295, 1995.

[11] L. Hurvich and D. Jameson. An opponent-process theory of color vision. *Psychological Review*, 64:384–390, 1957.

[12] J.G. Liu and J.McM. Moore. Hue image RGB colour composition. a simple technique to suppress shadow and enhance spectral signature. *International Journal of Remote Sensing*, 11:1521–1530, 1990.

[13] B.M. Mehtre, M.S. Kankanhalli, A.D. Narasimhalu, and G.C. Man. Color matching for image retrieval. *Pattern Recognition Letters*, 16(3):325–331, 1995.

[14] F. Perez and C. Koch. Hue color segmentation determines object boundaries. *Int. J. Computer Vision*, 12, 1994.

[15] C. Poynton. The rehabilitation of gamma. In *SPIE Conf. Human Vision and Electronic Imaging III*, volume 3299, pages 232–249, 1998.

[16] Y. Raja, S. McKenna, and S. Gong. Colour model selection and adaptation in dynamic scenes. In *5th European Conference on Computer Vision*, pages 460–474, 1998.

[17] J. Sangwine and R.E.N. Horne. *The Colour Image Processing Handbook*. Chapman & Hall, 1998.

[18] G. Schaefer, G. Qiu, and G. Finlayson. Retrieval of palettised colour images. In *Proc. of SPIE, Vol. 3972 (Storage and Retrieval for Image and Video Databases VIII)*, pages 483–493, 2000.

[19] M. Stricker and M. Orengo. Similarity of color images. In *SPIE Conf. on Storage and Retrieval for Image and Video Databases III*, volume 2420, pages 381–392, 1995.

[20] M.J. Swain and D.H. Ballard. Color indexing. *Int. J. Computer Vision*, 7(11):11–32, 1991.