

Colour Image Analysis in 3D-polar Coordinates

Allan Hanbury^{1*} and Jean Serra²

¹ Pattern Recognition and Image Processing Group, Vienna University of Technology, Favoritenstraße 9/1832, A-1040 Vienna, Austria

hanbury@prip.tuwien.ac.at,

WWW home page: <http://www.prip.tuwien.ac.at/~hanbury/>

² Centre de Morphologie Mathématique, Ecole des Mines de Paris, 35 rue Saint-Honoré, F-77305 Fontainebleau cedex, France

Abstract. The use of 3D-polar coordinate representations of the RGB colour space is widespread, although many of these representations, such as HLS and HSV, have deficiencies which render them unsuitable for quantitative image analysis. Three prerequisites for 3D-polar coordinate colour spaces which do not suffer from these deficiencies are suggested, and the results of the derivation of three colour spaces based on these prerequisites are presented. An application which takes advantage of their good properties for the construction of colour image histograms is also discussed.

1 Introduction

Representations of the RGB colour space in terms of 3D-polar coordinates (hue, saturation and brightness) are often used. Even though this corresponds to a simple coordinate transform from rectangular to 3D-polar (cylindrical) coordinate systems, the literature abounds with many different ways of performing this transformation (the HLS, HSV, HSI, etc. colour spaces). Many of these systems were developed with computer graphics applications in mind [1], and have a number of shortcomings when used for quantitative image analysis.

In this paper, we discuss the deficiencies of commonly used systems (section 2), and suggest three prerequisites for 3D-polar coordinate systems to be useful for quantitative image analysis (section 3). In section 4, we present three 3D-polar coordinate representations which were derived based on these prerequisites. Finally, an application is discussed (section 5).

2 Deficiencies of the commonly used 3D-polar spaces

To represent colours in an RGB coordinate system in terms of hue, saturation and brightness, one begins by placing a new axis, called the *achromatic axis*, into the RGB space between the pure black and pure white points [2]. The

* This work was supported by the Austrian Science Foundation (FWF) under grants P14445-MAT and P14662-INF.

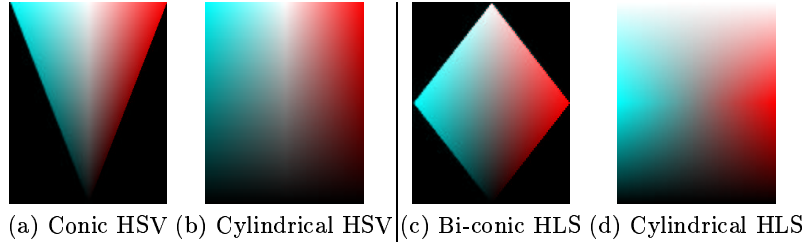


Fig. 1. Slices through the conic and cylindrical versions of the HSV and HLS colour spaces. The brightness increases from bottom to top, and the saturation increases from the centre (achromatic axis) outwards. Colours to the right of the central achromatic axis have hues of 0° , and colours to the left have hues of 180° .

choice of a function describing the brightness then gives rise to a set of iso-brightness surfaces, with each surface containing all the points having a specific brightness. These surfaces are then projected onto a plane perpendicular to the achromatic axis and intersecting it at the origin, called the *chromatic plane*, where they form hexagons. The *hue* and *saturation* or *chroma* coordinates are determined within this plane. The hue corresponds to an angular coordinate around the achromatic axis, traditionally measured with respect to pure red, and the saturation or chroma to the distance from the achromatic axis.

Two commonly used 3D-polar coordinate colour systems are HSV and HLS. The HSV system is usually described as being in the shape of a hexcone, and the HLS system in the shape of a double-hexcone, as shown in figures 1a and 1c. However, the commonly used conversion formulae, having been developed to simplify numerical colour selection, produce spaces which have been artificially expanded into cylindrical form (figures 1b and 1d).

These cylindrically shaped 3D-polar coordinate spaces are not suitable for quantitative image analysis for the reasons presented here. By definition, saturation has a low value for black, white or grey pixels, and a higher value for more colourful pixels. However, the commonly used formulae for the HSV and HLS spaces can assign a maximum saturation value to an achromatic pixel. For example, if one calculates the HSV saturation of the RGB coordinates $(0.01, 0, 0)$ using the commonly used $S_{\text{HSV}} = \frac{\max(R, G, B) - \min(R, G, B)}{\max(R, G, B)}$, one obtains a saturation of 1, even though the colour is visually indistinguishable from pure black.

It is often said that these spaces separate chrominance (hue and saturation) and brightness information. However, the normalisation included to convert the conically-shaped spaces into cylindrically-shaped spaces introduces an interdependence between these coordinates. For example, if one converts the RGB coordinates $\mathbf{c} = (0.5, 0.5, 0)$ into HLS coordinates with the commonly used conversion formula, one obtains $(H, L, S) = (60^\circ, 0.25, 1)$. If $\Delta\mathbf{c} = (0.25, 0.25, 0.25)$ is then added to the initial RGB vector \mathbf{c} , a modification corresponding uniquely to a change in the brightness, the HLS coordinates become $(60^\circ, 0.5, 0.5)$. In other words, the saturation value has been diminished because of an increase in the brightness value, implying that these values are certainly not independent.

3 Prerequisites for a useful 3D-polar representation

We suggest three prerequisites whose adoption leads to the development of 3D-polar coordinate colour spaces which do not suffer from the shortcomings listed in the previous section, and hence are suitable for quantitative image analysis. We model the RGB space by a Euclidean vector space over \mathbb{R}^3 , allowing us to make use of its projections, orthogonality, etc., but we equip it successively with different norms. The axes of this space are labelled R , G and B , and the space of valid coordinates is limited to $0 \leq R \leq 1$, $0 \leq G \leq 1$ and $0 \leq B \leq 1$.

The vector space notion associates a point $\mathbf{c} = (R, G, B)$ to the vector \overrightarrow{oc} . This point can be written in terms of vectors parallel to the R , G and B axes, or $\overrightarrow{oc} = \overrightarrow{or} + \overrightarrow{og} + \overrightarrow{ob}$. Equivalently, it can be written as $\overrightarrow{oc} = \overrightarrow{oc_d} + \overrightarrow{oc_p}$, where $\mathbf{c_d}$ and $\mathbf{c_p}$ are the projections of \mathbf{c} onto respectively the achromatic axis and the chromatic plane. We say that $\mathbf{c_d}$ and $\mathbf{c_p}$ are independent if the parameters associated with $\mathbf{c_p}$ (saturation, hue) do not affect those associated with $\mathbf{c_d}$. This is equivalently stated in the following prerequisite:

First prerequisite: Two distinct colours which have the same projection onto the chromatic plane, have the same chromatic parameters.

Use of a vector space also allows us to make use of *norms* and the associated triangular inequality, which says that the norm of the mean vector between \mathbf{c} and $\mathbf{c'}$ cannot be larger than the average of the norms of \mathbf{c} and of $\mathbf{c'}$. For example, two projections onto the chromatic plane which are far from the achromatic axis, but opposite each other, represent colours which are highly saturated. The vector mean of these two colours is, however, achromatic. It therefore makes sense that its norm should not be larger than the norms of the original colours, and hence that the triangular inequality should be satisfied. This leads to the following prerequisite:

Second prerequisite: The intensity parameters associated with colour vector \mathbf{c} (brightness) and with its projection $\mathbf{c_p}$ (saturation) must be norms.

For example, the HLS brightness and saturation parameters are not norms, which leads to some of the undesirable properties discussed in the previous section. Finally, motivated by practical experience, we suggest:

Third prerequisite: Every system for the representation of colour images must be reversible with respect to the RGB standard.

The colour spaces which satisfy the prerequisites presented have the following main advantages over the commonly used 3D-polar coordinate spaces:

- Achromatic or near-achromatic colours always have a low saturation value.
- The saturation and brightness coordinates are independent, as there is no normalisation of the saturation by the brightness.
- Comparisons between saturation values are meaningful, also due to the saturation normalisation having been removed.

4 Three 3D-polar coordinate colour representations

We present the conversion formulae which result when one limits oneself to using only the L_2 or L_1 norms. The full derivations of these formulae are in [3]. Lastly, we demonstrate the use of the semi-norm $\max - \min$ as a measure of saturation. When describing the length of the vector \mathbf{c}_p , we make use of two terms, *saturation* and *chroma*. We define chroma as the norm of \mathbf{c}_p , as done by Carron [4] (who uses the L_2 norm). It assumes its maximum value at the six corners of the hexagon projected onto the chromatic plane. For the saturation, the hexagon projected onto the chromatic plane is slightly deformed into a circle by a normalisation factor, so that the saturation assumes its maximum value for all points with projections on the edges of the hexagon. Poor choice of this normalisation factor has led to some of the less than useful saturation definitions currently in use.

4.1 L_2 norm

The conversion equations from the RGB system for the L_2 norm are easy to determine. We call the brightness, chroma and hue determined using this norm M_2 , C_2 and H_2 respectively. The M_2 and C_2 norms are scaled to the range $[0, 1]$.

$$M_2 = \frac{1}{\sqrt{3}} [R^2 + G^2 + B^2]^{1/2} \quad (1)$$

$$C_2 = \sqrt{\frac{3}{2}} \|\mathbf{c}_p\| = (R^2 + G^2 + B^2 - RG - RB - BG)^{\frac{1}{2}} \quad (2)$$

It is possible to convert the chroma measurement C_2 into a saturation measurement S_2 by dividing $\|\mathbf{c}_p\|$ by the distance from the origin to the edge of the hexagon for a given hue H , that is, the maximum value that can be taken by the norm of a projected vector $\|\mathbf{c}_p\|$ with hue H [3]. The hue is calculated as

$$H_2 = \begin{cases} 360^\circ - \arccos \left[\frac{\mathbf{r}_p \cdot \mathbf{c}_p}{\|\mathbf{r}_p\| \|\mathbf{c}_p\|} \right] & \text{if } B > G \\ \arccos \left[\frac{\mathbf{r}_p \cdot \mathbf{c}_p}{\|\mathbf{r}_p\| \|\mathbf{c}_p\|} \right] & \text{otherwise} \end{cases} \quad (3)$$

where \mathbf{r}_p is the projection of the vector representing pure red onto the chromatic plane, and $\mathbf{r}_p \cdot \mathbf{c}_p$ indicates the scalar product of the two vectors.

The advantages of using the L_2 norm are that the Euclidean distance is used and that an accurate value for the hue can be determined, which can also be used in the other representations. The biggest disadvantage is that the inverse transformation back to RGB coordinates is not simple.

4.2 L_1 norm

As the values of R , G and B are positive, the L_1 norm brightness (M_1) is simply the arithmetic mean of these three components, which can also be written in

terms of the maximum, median and minimum component of the RGB vector, notated as max, mid and min. The chroma C_1 is also written in terms of these three values. If one requires an accurate hue value, then the L_2 norm value H_2 should be used. An approximate value H_1 , requiring no trigonometric function evaluations can be derived [3]. The full conversion from RGB coordinates is:

$$M_1 = \frac{1}{3} (\max + \text{mid} + \min) \quad (4)$$

$$C_1 = \begin{cases} \frac{3}{2} (\max - M_1) & \text{if } \max + \min \geq 2\text{mid} \\ \frac{3}{2} (M_1 - \min) & \text{if } \max + \min \leq 2\text{mid} \end{cases} \quad (5)$$

$$H_1 = k \left[\lambda(\mathbf{c}) + \frac{1}{2} - \frac{\max + \min - 2\text{mid}}{2C_1} \right] \quad (6)$$

where $\lambda(\mathbf{c})$ gives the RGB cube sector number (an integer from 0 to 5) in which the vector \mathbf{c} lies [2].

Note that, given variables $R, G, B \geq 0$, every quantity $\alpha R + \beta G + \gamma B$, with weights $\alpha, \beta, \gamma \geq 0$ is still an L_1 norm on the achromatic axis. This is advantageous as it permits the use of psycho-visual *luminance* functions. A further advantage is that the system is easy to invert. However, one should beware of an inbuilt quantisation pitfall. When R, G and B are integer-valued, then C_1 is always a multiple of $1/2$. The rounding of a floating point value of C_1 to the nearest integer therefore behaves extremely erratically, as the 0.5's are sometimes rounded up and sometimes down.

4.3 max – min semi-norm

The semi-norm $S_0 = \max - \min$ (this fact is proved in [3]) obviously satisfies the first prerequisite as its value does not change when one shifts an arbitrary RGB vector parallel to the achromatic axis by adding the same amount to each component. In particular, a vector \mathbf{c} in the RGB space and its projection \mathbf{c}_p have the same value for S_0 .

Because of this independence to changes in brightness, it is impossible to build a 3D-polar colour representation based only on this semi-norm. However, it is ideally suited to describing saturation, and we can take advantage of the independence of the achromatic and chromatic components to choose either the L_1 or L_2 norm on the achromatic axis. It can be shown the S_0 corresponds exactly to the saturation measure derived for the L_2 representation, and it can be derived from the HLS and HSV spaces by removing the saturation normalisation factors [3]. Finally, for the hue, we can use the trigonometric expression (equation 3), replacing it by the approximation (equation 6) if calculation speed is important. We have named this representation the improved HLS (IHLS) space. MATLAB code implementing this conversion is available on the author's home page.

4.4 Comparison of the saturation and chroma formulations

We compare the distributions of three of the saturation and chroma formulations discussed: the $S_0 = \max - \min$ saturation expression, the L_2 norm chroma C_2

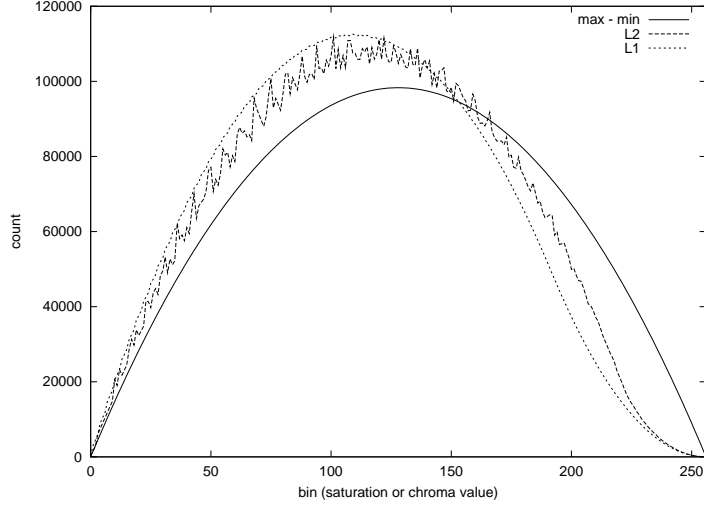


Fig. 2. The saturation and chroma histograms.

(equation 2), and the L_1 norm chroma C_1 (equation 5). We begin with a $256 \times 256 \times 256$ RGB cube with a point at each set of integer-valued coordinates, and look at the 256-level histograms resulting from calculating the saturation and chroma for each of these points. For the max – min and L_2 norms, the S_0 and C_2 values of each point are calculated (as floating point values for the latter, which are then rounded to the nearest integer). To avoid the quantisation problems outlined previously, the values of C_1 were first multiplied by 2 to get a series of integers between 0 and 512, and then adjacent pairs of histogram bins were combined to produce a 256 bin histogram. The histograms are shown in figure 2.

The max – min saturation distribution is regular and symmetric around the central histogram bin because of the normalisation coefficient which deforms the hexagonally shaped sub-region of the chromatic plane into a circle. Conversely, the L_2 chroma has a rather irregular distribution due to the discrete space in which it is calculated. It decreases rapidly as one approaches higher chroma values as it is calculated in hexagonally shaped sub-regions of the chromatic plane. The L_1 norm chroma approximates the L_2 chroma well (if the quantisation effects are taken into account), and its histogram is more regular.

5 Application

A number of applications which take advantage of the good properties of the proposed colour representation have already been suggested, including colour morphology [3] and luminance/saturation histograms [6]. Here we present an application of circular statistics to the calculation of colour image histograms.

As the hue is an angular value, circular statistics [7] should be used to calculate its mean direction. We take advantage of the correlation between colour-

fulness and saturation value in the suggested colour representation to apply a weighting to the circular mean so that achromatic colours have less influence on the result [5]. The circular mean of a set of angular data is defined as the direction of the resultant vector of the sum of unit vectors in the given directions, and the saturation-weighting is thus easily included by replacing the unit vectors by vectors with lengths proportional to the saturation. One can thus easily determine the saturation-weighted hue mean of a whole image. We next propose that this mean (and its associated variance) be calculated separately for pixels belonging to each luminance level of a colour image, leading to the construction of a colour histogram analogous to the greyscale one.

Given a colour image in the IHLS space, the luminance values are first quantised into $N + 1$ levels labeled by $\ell = \{0, 1, 2, \dots, N\}$. Then, for each value of ℓ , the following circular statistics descriptors are calculated, where $\overline{H}_{S\ell}$ is the saturation-weighted hue mean for luminance level ℓ , and $\overline{R}_{n\ell}$ is the associated mean length. The latter is essentially the inverse of the circular variance, and assumes values in the range $[0, 1]$. A value approaching one indicates that the hue values are more closely grouped. $A_{S\ell}$ and $B_{S\ell}$ are intermediate values.

$$A_{S\ell} = \sum_x S_x \cos H_x \delta_{L_x \ell}, \quad B_{S\ell} = \sum_x S_x \sin H_x \delta_{L_x \ell} \quad (7)$$

$$\overline{H}_{S\ell} = \arctan\left(\frac{B_{S\ell}}{A_{S\ell}}\right), \quad \overline{R}_{n\ell} = \frac{\sqrt{A_{S\ell}^2 + B_{S\ell}^2}}{\sum_x \delta_{L_x \ell}} \quad (8)$$

where H_x , L_x and S_x are the hue, luminance and saturation at position x in the image, and the sums are over all the pixels in the image. The symbol $\delta_{L_x \ell}$ is the Kronecker delta which limits the calculation to luminance level ℓ .

We therefore have two histograms of colour information, the mean hue and its associated mean length as a function of luminance. These could conceivably be used directly in image matching and database retrieval applications. For visualisation purposes, these two histograms can very simply be combined into a single histogram, in which the height of the bar at luminance ℓ corresponds to the mean length $\overline{R}_{n\ell}$, and its colour is given by the fully saturated colour corresponding to the mean hue $\overline{H}_{S\ell}$. As the mean hue associated with a very low mean length value does not give much information, we set the colours of the bars with a mean length below a threshold (here 0.05) to the greylevel corresponding to the associated luminance value.

Figure 3c shows the colour histogram of figure 3a, with figure 3b showing its luminance. A luminance quantisation parameter of $N = 100$ was used. One sees clearly from the histogram that the dominant colour for the low luminance parts (the vegetation) is green, the very uniform (i.e. with high mean length) blues associated with the water, and the reds of the highly luminous roofs.

6 Conclusion

We have pointed out the deficiencies of commonly used 3D-polar coordinate colour representations (specifically HLS and HSV, but also applicable to many

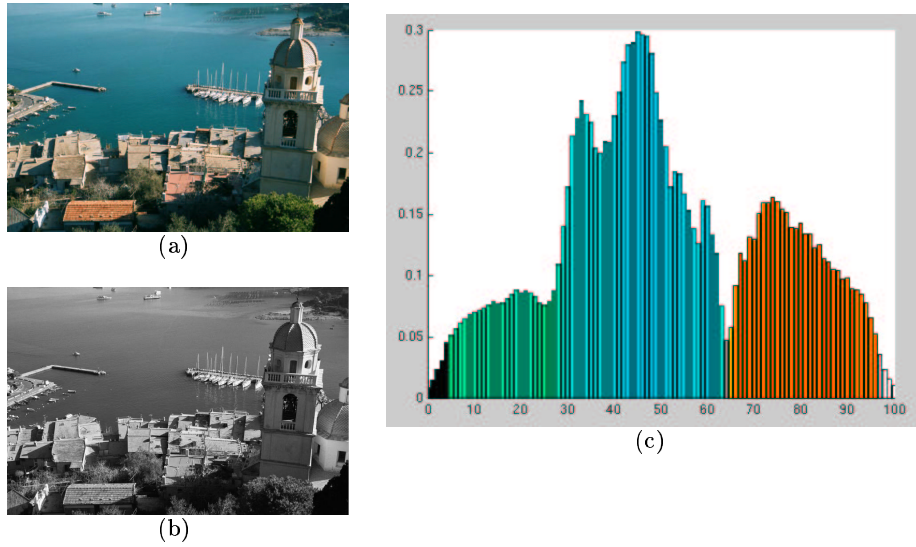


Fig. 3. (a) Colour image (from the University of Washington content-based image retrieval database), and its (b) Luminance. (c) Colour histogram of image (a).

others) rendering them unsuitable for quantitative image analysis. We then list three prerequisites for such colour representations to have useful properties, and summarise three sets of conversion equations between the RGB space and 3D-polar coordinate spaces based on these prerequisites. The good properties of the suggested representations give rise to many applications in image analysis, of which we have described one: the construction of a colour image histogram based on circular statistics. Further work is being undertaken to continue the development of these applications.

References

1. Smith, A.R.: Color gamut transform pairs. *Computer Graphics* **12** (1978) 12–19
2. Levkowitz, H., Herman, G.T.: GLHS: A generalised lightness, hue and saturation color model. *CVGIP: Graphical Models and Image Processing* **55** (1993) 271–285
3. Hanbury, A., Serra, J.: A 3D-polar coordinate colour representation suitable for image analysis. Technical Report PRIP-TR-077, TU Wien (2002)
4. Carron, T.: Segmentations d’images couleur dans la base Teinte-Luminance-Saturation. PhD thesis, Université de Savoie (1995)
5. Hanbury, A.: Circular statistics applied to colour images. In: *Proceedings of the Computer Vision Winter Workshop 2003*. (2003) 55–60
6. Angulo, J., Serra, J.: Colour feature extraction from luminance/saturation histogram in L_1 representation. Technical Report N-01/03/MM, Centre de Morphologie Mathématique, Ecole des Mines de Paris (2003)
7. Fisher, N.I.: *Statistical Analysis of Circular Data*. Cambridge University Press (1993)