The Taming of the Hue, Saturation and Brightness Colour Space

Allan Hanbury
Centre de Morphologie Mathématique, Ecole des Mines de Paris
35 rue St-Honoré, 77305 Fontainebleau cedex, France
telephone: +33 1 64 69 48 05, fax: +33 1 64 69 47 07
e-mail: Hanbury@cmm.ensmp.fr

Abstract

The transformation of the RGB colour space to a hue, saturation and brightness colour space is essentially a conversion from a rectangular coordinate system to a cylindrical coordinate system. Nevertheless, a bewildering array of such conversions exist. We show that one of the main reasons for this is the dependence of the saturation values on the choice of the brightness function, and suggest a definition of saturation which is independent of the brightness. The usual ways of calculating brightness and hue are reviewed. Lastly, we examine some of the characteristics of the cylindrical colour coordinates and give a simple example in which the suggested cylindrical colour coordinates are used.

1 Introduction

The transformation of the RGB colour space to a hue, saturation and brightness colour space is essentially a conversion from a set of rectangular coordinates to a set of cylindrical coordinates. One could therefore ask how such a seemingly simple procedure could have given rise to the plethora of such transformations described in the literature, such as HSV [10], HSI [4], Triangle [10] and HLS[1]. It is shown in this article that in the definitions of these spaces, the saturation values obtained depend intimately on the expression chosen for calculating brightness, even though it is usually claimed that the saturation and brightness measures are independent. We then propose a definition of the saturation which is completely independent of the brightness function, and which therefore allows the free choice of the brightness function most suited to the task at hand. In order to complete the description of the space, we review the methods used to calculate brightness and hue, and we examine some of their characteristics.

Why, it may be asked, is such a colour representation space necessary? Surely it is better to use a standardised colour space such as the CIE L*a*b* space or its cylindrical coordinate version. The obvious objection to the use of the L*a*b* space is that one needs calibration

[1]The transformations to and from these spaces and some others are summarised by Shih [9].

information on the image which is being transformed from the RGB space, namely the colour coordinates of the source of illumination (the white point). This information is not always available for the images that are encountered in computer vision applications. If we do not have the necessary information, might it not be better to avoid making assumptions and estimations, and to use an alternate and more intuitive coordinate system for representing the information that is known, namely the coordinates of the colours in the RGB space.

We begin with a discussion of the existing cylindrical coordinate colour representations (section 2), followed by a discussion and derivation of the suggested representation (section 3). In section 4, we present some of the characteristics of this representation. Finally, a simple example of its use is given in section 5.

2 Discussion of the existing transforms

In the RGB space, colours are specified as vectors (R, G, B) which give the amount of each red, green and blue primary in the colour. For convenience, we take $R, G, B \in [0, 1]$ so that the valid coordinates form the cube $[0, 1] \times [0, 1] \times [0, 1]$. The basic idea behind the transformation to a hue, saturation and brightness coordinate system is to place a new axis between $(0, 0, 0)$ and $(1, 1, 1)$, and to specify the colours in terms of cylindrical coordinates based on this axis. The new axis passes through all the achromatic or grey points (i.e. with $R = G = B$), and will therefore be referred to as the *achromatic axis*. The *brightness* gives the coordinate of a colour on this axis, the *hue* corresponds to the angular coordinate and the *saturation* corresponds to the distance from the achromatic axis.

One of the causes of the variety of such spaces is the number of different definitions of brightness. These definitions lead to spaces which have shapes which are not simply constructed as a pile of planar cross-sections of the cube taken perpendicular to the achromatic axis. Further problems with the existing transforms are due to them originally being developed for the easy numerical specification of colours in computer graphics applications [10]. Due to the associated brightness functions, the "natural" shape of the HSV space is a cone, and of the HLS space, a double cone. A vertical slice through the achromatic axis of each of these spaces is shown in figures 1a and 1c. The problem with using these representations when specifying a colour is that there are large regions which lie outside the cones. In order to avoid complicated verification of the validity of a specified colour, these spaces were often artificially expanded into cylinders by dividing the saturation values by their maximum possible values for the corresponding brightness. Slices of the cylindrical versions of the HSV and HLS spaces are shown in figures 1b and 1d. The cylindrical versions have often been carried over into image processing and computer vision, for which they are ill-suited.

We now consider two cases of the confusion that the cylindrical forms can cause. Demarty and Beucher [3] applied a constant saturation threshold in the cylindrical HLS space (figure 1d) to differentiate between chromatic and achromatic colours. This threshold can be represented by a vertical line on either side of the achromatic axis in figure 1d, and it is clear that this does not correspond to a constant saturation. Demarty [2] later improved the threshold by using a hyperbola in the cylindrical HSV space (figure 1b), which corresponds to a constant threshold

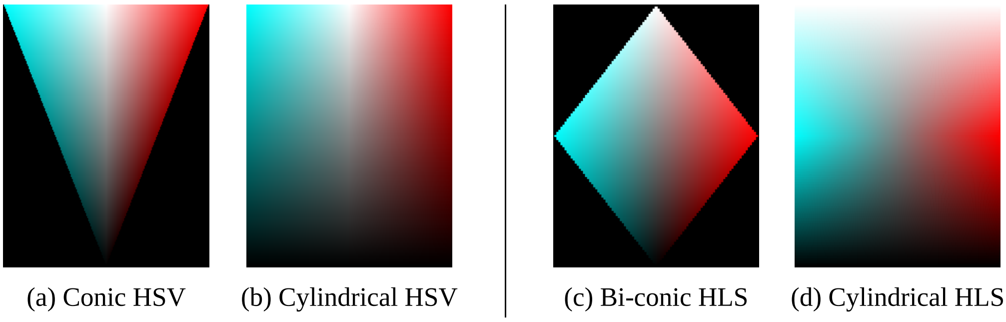

(a) Conic HSV (b) Cylindrical HSV (c) Bi-conic HLS (d) Cylindrical HLS

Figure 1: Slices through the conic and cylindrical versions of the HSV and HLS colour spaces. Colours to the right of the central achromatic axis have hues of $0°$, and colours to the left have hues of $180°$.

in the conic HSV space (figure 1a). Smith [11] makes the assumption that the cylindrical HSV space is perceptually uniform when a Euclidean metric is used, but upon examining figure 1b, one sees that a certain distance in the high brightness (top) part of the space corresponds to a far larger perceived change in colour than the same distance in the low brightness part of the space.

3 Derivation of a useful hue, saturation and brightness space

In this section, we examine a derivation of a cylindrical coordinate system in the RGB space, pointing out the choices which could (and have) lead to characteristics which are disadvantageous, and ending up with a cylindrical coordinate representation of the RGB space which is useful for computer vision. This derivation is based on the derivation of a Generalised Lightness, Hue and Saturation (GLHS) model [7] suitable for computer graphics applications.

3.1 Brightness

In order to conform to the terminology suggested by the CIE, we call a subjective measure of luminous intensity the *brightness*. The brightness function of the GLHS model is

$$L(\mathbf{c}) = w_{\mathrm{min}} \cdot \min(\mathbf{c}) + w_{\mathrm{mid}} \cdot \mathrm{mid}(\mathbf{c}) + w_{\mathrm{max}} \cdot \max(\mathbf{c}) \tag{1}$$

in which the functions $\min(\mathbf{c})$, $\mathrm{mid}(\mathbf{c})$ and $\max(\mathbf{c})$ return respectively the minimum, median and maximum component of a vector $\mathbf{c}$ in the RGB space, and w_{min}, w_{mid} and w_{max} are weights set by the user, with the constraints $w_{\mathrm{max}} > 0$ and $w_{\mathrm{min}} + w_{\mathrm{mid}} + w_{\mathrm{max}} = 1$. Specific values of the weights give the brightness functions used by the common cylindrical colour spaces: $w_{\mathrm{min}} = 0$, $w_{\mathrm{mid}} = 0$ and $w_{\mathrm{max}} = 1$ for HSV; $w_{\mathrm{min}} = \frac{1}{2}$, $w_{\mathrm{mid}} = 0$ and $w_{\mathrm{max}} = \frac{1}{2}$ for HLS; and $w_{\mathrm{min}} = \frac{1}{3}$, $w_{\mathrm{mid}} = \frac{1}{3}$ and $w_{\mathrm{max}} = \frac{1}{3}$ for HSI.

The *luminance* is the radiant intensity per unit projected area weighted by the spectral sensitivity associated with the brightness sensation of human vision [8]. This objective measure takes into account the fact that if one looks at red, green and blue light sources of the same radiant intensity in the visible spectrum, the green will appear the brightest and the blue the darkest. The luminance function which corresponds to contemporary video displays is [8]

$$Y(\mathbf{c}) = 0.2125R + 0.7154G + 0.0721B \tag{2}$$

In the RGB space, we can visualise surfaces of iso-brightness (or iso-luminance). The surfaces of iso-brightness l contain all the points such that $L(\mathbf{c}) = l$ and intersect the achromatic axis at l. For the HSV and HLS spaces, these surfaces have a complicated shape (see [7] for details), and for the HSI space these surfaces are planes perpendicular to the achromatic axis. The surfaces of iso-luminance (equation 2) are planes oblique to the achromatic axis.

3.2 Hue

The hue angle is traditionally measured starting at the direction corresponding to pure red. The simplest way to derive an expression for this angle is to project the vector $(1, 0, 0)$ corresponding to red in the RGB space and an arbitrary vector $\mathbf{c}$ onto a plane perpendicular to the achromatic axis, and to calculate the angle between them. This gives the expression

$$H' = \arccos\left[\frac{R - \frac{1}{2}G - \frac{1}{2}B}{(R^2 + G^2 + B^2 - RG - RB - BG)^{\frac{1}{2}}}\right] \tag{3}$$

after which, in order to give a value of $H \in [0°, 360°]$, we apply

$$H = \begin{cases} 360° - H' & \text{if } B > G \\ H' & \text{otherwise} \end{cases} \tag{4}$$

An approximation to this trigonometric expression is often used, and it is shown in [7] that the approximated value differs from the trigonometric value by at most $1.12°$. A further comparison between the trigonometric hue and the approximated hue is given in section 4.

3.3 Saturation

For the derivation of an expression for the saturation of an arbitrary colour $\mathbf{c}$, we begin by looking at the triangle which contains all the points with the same hue as $\mathbf{c}$, as shown in figure 2. The intersection of this triangle and the iso-brightness surfaces are lines parallel to the line between $\mathbf{c}$ and its brightness value on the achromatic axis $\mathbf{L}(\mathbf{c}) = [L(\mathbf{c}), L(\mathbf{c}), L(\mathbf{c})]$.

Traditionally, the saturation is calculated as the length of the vector from $\mathbf{L}(\mathbf{c})$ to $\mathbf{c}$ divided by the length of the extension of this vector to the surface of the RGB cube. This definition, however, results in colour spaces in the form of cylinders discussed in section 2. Moreover, it is clear that this definition of the saturation depends intimately on the form of the brightness function chosen (i.e. on the slopes of the iso-brightness lines). An example of this dependence is shown in figure 3, in which the saturation of figure 3a is shown in figure 3b for the HSV space and figure 3c for the HLS space. In the original image, not all the pixels which appear white have RGB coordinates of exactly $(1, 1, 1)$. The slight variations in these RGB values are

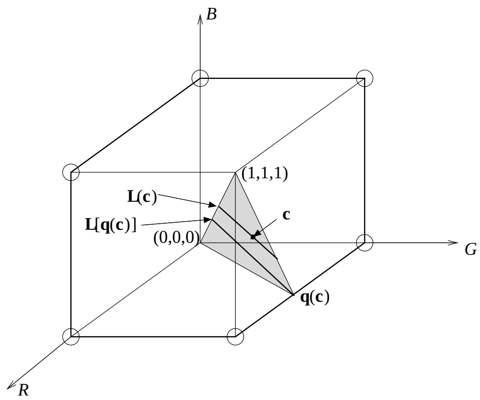

Figure 2: The triangle which contains all the points with the same hue as $\mathbf{c}$. The circled corners mark the edges of the cube containing the points furthest away from the achromatic axis.

amplified by the artificial expansion of the cones into cylinders, leading to the noisy regions in the saturation images and clearly demonstrating the dependence of the saturation on the brightness function.

In order to keep the conic or bi-conic forms of the spaces, it is necessary to change the definition of the saturation. Instead of the definition given above, we divide the length of the vector from $\mathbf{L}(\mathbf{c})$ to $\mathbf{c}$ (in figure 2) by the length of the vector between $\mathbf{L}[\mathbf{q}(\mathbf{c})]$ and $\mathbf{q}(\mathbf{c})$, that is, the longest vector parallel to $[\mathbf{L}(\mathbf{c}), \mathbf{c}]$ included in the iso-hue triangle, the vector which necessarily intersects the third corner $\mathbf{q}(\mathbf{c})$ of the triangle. We then end up with the following expression for the saturation

$$S = \frac{\|\mathbf{L}(\mathbf{c}) - \mathbf{c}\|}{\|\mathbf{L}[\mathbf{q}(\mathbf{c})] - \mathbf{q}(\mathbf{c})\|} \tag{5}$$

which is independent of the choice of the brightness function. This independence can be shown by using similar triangles [6]. An example of this saturation measurement is shown in figure 3d, where it should be compared to the corresponding HSV and HLS examples. The most visible improvement resulting from this definition is that both the white and black regions of the colour image are assigned a low saturation value.

The points the furthest away from the achromatic axis are those on the edges of the RGB cube between the circled corners in figure 2. These points correspond to the most highly saturated colours, and if we project them onto a plane perpendicular to the achromatic axis, they form the edges of a hexagon, which correspond to the maximum distance a point can be from the achromatic axis for a given hue. A simpler expression for the saturation of point $\mathbf{c}$ can be obtained by projecting it onto this hexagon, and dividing the distance of the projected point from the centre of the hexagon by the distance from the centre to the hexagon edge at the same hue value.

3.4 Chroma

Carron [1] suggests the use of the distance of a point from the achromatic axis without the maximum distance normalisation as an approximation to the saturation, which he calls *chroma*.

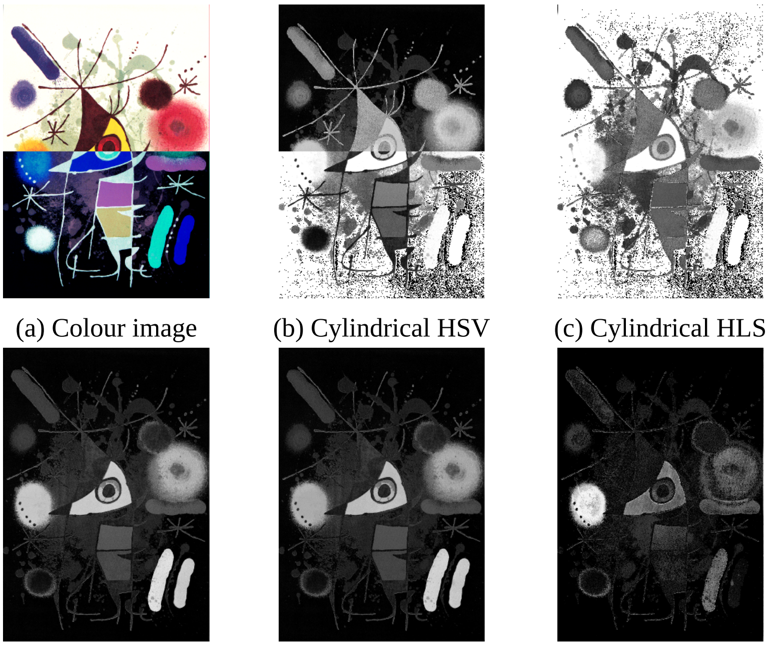

(a) Colour image (b) Cylindrical HSV (c) Cylindrical HLS

(d) Saturation S (e) Chroma C (f) $S - C$

Figure 3: (a) "Le Chanteur" by Joan Miró, with the bottom half inverted (by subtracting the values in the three colour channels from 255). The cylindrical saturation is shown in (b) for the HSV space and (c) for the HLS space. The brightness-independent (d) saturation and (e) chroma are shown, as well as (f) the difference between images d and e (contrast-enhanced).

This distance is multiplied by a constant so that for the six vertices of the projected hexagon, the chroma has a value of one. An example of the chroma is shown in figure 3e, and the difference between the chroma and the saturation images is shown in figure 3f (the contrast has been enhanced for better visibility, the maximum pixel value in the image is 0.107). The maximum possible difference between a saturation and a chroma value for a colour is 0.134.

3.5 Summary of the transform

A simple method to calculate the luminance, trigonometric hue, chroma and saturation coordinates is given here, based on the one suggested by Carron [1]. The changes with respect to the version given by Carron are the extension to calculate the saturation from the chroma, and the use of luminance instead of brightness. The first step is

$$\begin{bmatrix} Y \\ C_1 \\ C_2 \end{bmatrix} = \begin{bmatrix} 0.2125 & 0.7154 & 0.0721 \\ 1 & -\frac{1}{2} & -\frac{1}{2} \\ 0 & -\frac{\sqrt{3}}{2} & \frac{\sqrt{3}}{2} \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \tag{6}$$

followed by the calculation of the chroma $C \in [0, 1]$

$$C = \sqrt{C_1^2 + C_2^2}$$

and the hue $H \in [0°, 360°]$

$$H = \begin{cases} \text{undefined} & \text{if} \quad C = 0 \\ \arccos\left(\frac{C_1}{C}\right) & \text{if} \quad C \neq 0 \text{ and } C_2 \leq 0 \\ 360° - \arccos\left(\frac{C_1}{C}\right) & \text{if} \quad C \neq 0 \text{ and } \mathrm{C}_2 > 0 \end{cases}$$

and, if required, the saturation $S \in [0, 1]$

$$S = \frac{2C \sin(120° - H^*)}{\sqrt{3}}$$

in which

$$H^* = H - k \times 60° \text{ where } k \in \{0, 1, 2, 3, 4, 5\} \text{ so that } 0° \leq H^* \leq 60° \tag{7}$$

The inverse of this transform is easily derived.

4 Characteristics of the space

In this section we examine the distributions of the hue, saturation, chroma and brightness coordinates for a transformation of a set of points equidistantly spaced in the RGB space to completely fill up the cube $[0, 1] \times [0, 1] \times [0, 1]$ (the distance between the points is 0.01).

We begin by examining the hue using the two calculation methods available, the trigonometric method and the approximate method. Histograms of 360 bins, with each bin corresponding to one degree, were calculated (the value in bin 360 is equal to the value in bin 0). The histograms for the trigonometric approach and for the approximate approach are shown in figures 4a and 4b respectively. The distributions are not smooth because we are calculating the angular coordinates of points distributed on a grid, but the most striking characteristic of these histograms are the strong peaks at each multiple of $60°$. If we ignore the peaks, it appears as if the approximate calculation gives a flatter distribution than the trigonometric calculation, for which the hexagonal structure of the planar cross-sections of the space is clearly visible.

What causes the peaks? Their distribution at multiples of $60°$ suggests that the hexagonal shape of the planar cross-sections are the cause. The fact that we are piling up many such hexagons to form the colour space could lead to a surplus of points in these directions forming the exaggerated peaks. To test this, we re-calculated the histograms using only the points with saturation values larger than 0.2, which gave the histograms shown in figures 4d (trigonometric method) and 4e (approximate method). By removing the interior part of the space, we have removed the peaks for the trigonometric hue. However, they are still present for the approximate hue, always accompanied by a one bin wide depression on either side. This demonstrates that the approximate method has a tendency to inflate the number of points assigned hues which are multiples of $60°$.

We now look at the brightness and luminance distributions, for which histograms (with 100 bins) are shown in figure 4c. The brightness measure used here is $L = \frac{1}{3}R + \frac{1}{3}G + \frac{1}{3}B$. These distributions do not have any particular features, and the choice is dependent on the preferences of the user or the requirements of the application.

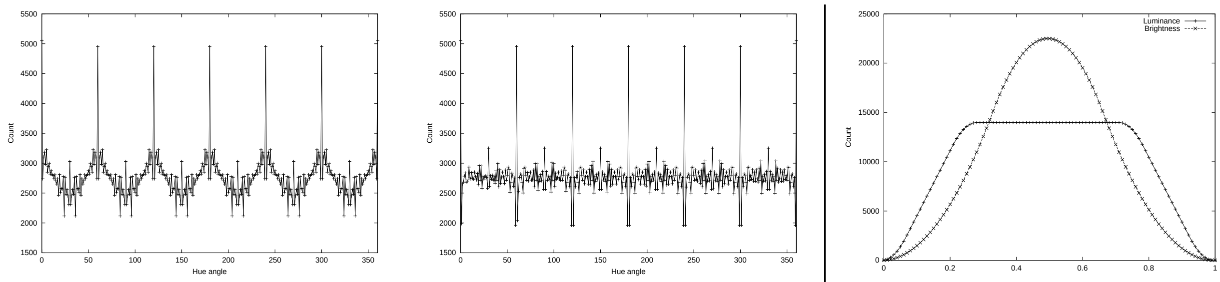

(a) Trigonometric H (b) Approximated H (c) Luminance and Brightness

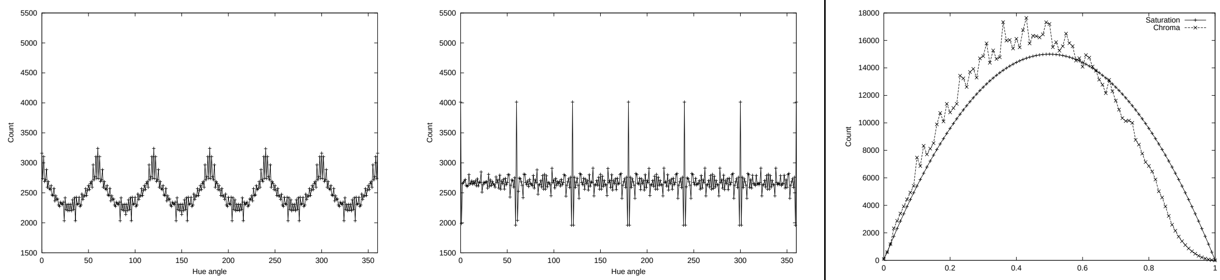

(d) Trig. H with threshold on S (e) Approx. H with threshold on S (f) Saturation and Chroma

Figure 4: The distributions of the hue, saturation, chroma, luminance and brightness coordinates after a transformation from an RGB cube containing a uniform distribution of points.

We look lastly at the saturation and chroma, for which one hundred bin histograms are shown in figure 4f. The saturation distribution is regular and symmetric around 0.5 due to its normalisation coefficient. The chroma distribution, on the other hand, is very irregular because of the distances calculated in a digital space, and descends to zero rapidly at the upper end due to the the hexagonal form of the planar cross-sections of the space.

The choice between the use of the approximate hue or trigonometric hue, and between chroma or saturation depends on the computing power available[2]. Given the computing power on our desktops, there is no excuse for not using the accurate versions in normal image analysis tasks. The only area in which one could consider using approximate hue or chroma are in very high speed industrial inspection tasks where the use of trigonometric hue and saturation might require the use of an extra DSP processor. However, a better approximation of the trigonometric hue can be obtained by the use of look-up tables for the trigonometric functions.

5 An example

We give a simple example of the use of the suggested hue, saturation and luminance coordinates. Figure 5a is a colour image in which we wish to extract the greyish lines between the mosaic tiles. The saturation of this image is shown in figure 5b, in which it is visible that the lines to be extracted have, in general, a lower saturation than the tiles. A morphological closing operation

[2]The approximate hue was introduced during the 1970's to speed up the interactive choice of colours in computer graphics programs.

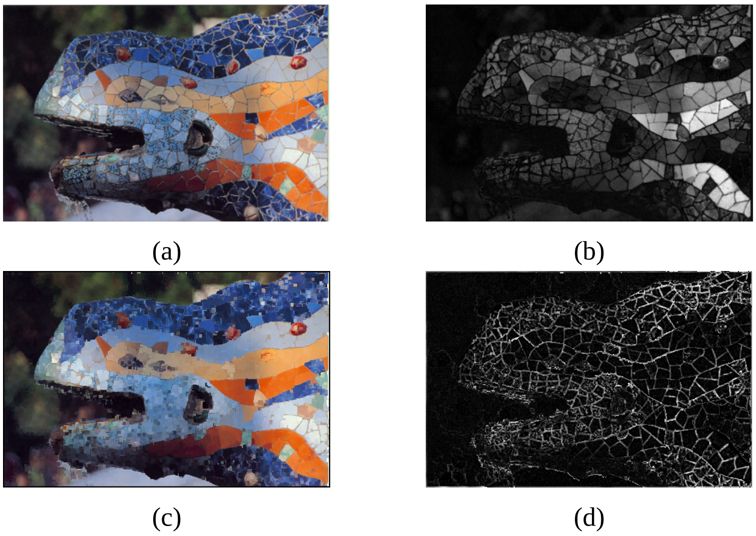

Figure 5: (a) The lizard image (size 544×360 pixels). (b) The saturation of the lizard image. (c) A morphological closing of the lizard image using a lexicographical order with saturation at the first level. (d) The top-hat — the Euclidean distance between images a and c.

with a square structuring element of size 5×5 pixels was applied to the initial colour image to give figure 5c. The colour vectors were completely ordered by using a lexicographical order with saturation at the first level, luminance at the second level and hue at the third level (the angular nature of the hues were taken into account). This closing succeeds in expanding the tiles to cover the grey lines. Finally, a form of top-hat was calculated by taking the Euclidean distance (in cylindrical coordinates) between figures 5a and 5c to give the greyscale image in figure 5d, in which the pixels of highest grey level correspond to the features we wish to extract.

This representation of the RGB space in cylindrical coordinates can be used in any application in which one of the HLS, HSI, etc. spaces are traditionally used, ensuring that the algorithms are not hampered by a poor representation of the data. Nevertheless, one should remember to take into account the angular nature of the hue component [5].

6 Discussion and conclusion

A critical evaluation of the hue, saturation and brightness or luminance colour spaces is presented, spaces which are essentially representations of the RGB space in cylindrical coordinates. These spaces are often used in computer vision, even though many of the suggested transformations found in the literature are optimised for the numerical specification of colours, and are badly suited to direct application to image processing. Two of the undesirable properties discussed are the artificial expansion of the conic or bi-conic spaces into cylinders, and the resulting dependence of the saturation on the brightness function used.

We have presented a formulation of the saturation which is independent of the brightness function, allowing an unconstrained choice of any brightness function (which has parallel iso-brightness surfaces), including a psycho-visual measure of the luminance. Comparisons of the distributions of the cylindrical coordinates are presented, as well as a simple example which uses the suggested cylindrical colour coordinates. This example makes use of a Euclidean distance in the suggested colour space to approximate a morphological top-hat. It would be more correct to calculate this difference in the L*a*b* space for which the Euclidean metric is defined, but in which it is less easy to calculate a saturation-based morphological closing. Further applications using the suggested colour space are given in [6], including a real-time wood colour matching application.

References

[1] T. Carron. *Segmentations d'images couleur dans la base Teinte-Luminance-Saturation: approche numérique et symbolique*. PhD thesis, Université de Savoie, 1995.

[2] C-H. Demarty. *Segmentation et Structuration d'un Document Vidéo pour la Caractérisation et l'Indexation de son Contenu Sémantique*. PhD thesis, CMM, Ecole des Mines de Paris, 2000.

[3] C-H. Demarty and S. Beucher. Color segmentation algorithm using an HLS transformation. In *Proceedings of the International Symposium on Mathematical Morphology (ISMM '98)*, pages 231–238, 1998.

[4] R. C. Gonzalez and R. E. Woods. *Digital Image Processing*. Prentice Hall, 1992.

[5] A. G. Hanbury and J. Serra. Morphological Operators on the Unit Circle. *IEEE Transactions on Image Processing*, 10(12):1842–1850, 2001.

[6] A. G. Hanbury. *Morphologie Mathématique sur le Cercle Unité: avec applications aux teintes et aux textures orientées*. PhD thesis, CMM, Ecole des Mines de Paris, 2002.

[7] H. Levkowitz and G. T. Herman. GLHS: A generalised lightness, hue and saturation color model. *CVGIP: Graphical Models and Image Processing*, 55(4):271–285, 1993.

[8] C. Poynton. Frequently asked questions about gamma. URL: http://www.inforamp.net/˜poynton/PDFs/GammaFAQ.pdf, 1999.

[9] T-Y. Shih. The reversibility of six geometric color spaces. *Photogrammetric Engineering and Remote Sensing*, 61(10):1223–1232, October 1995.

[10] A. R. Smith. Color gamut transform pairs. *Computer Graphics*, 12(3):12–19, 1978.

[11] J. R. Smith. *Integrated Spatial and Feature Image Systems: Retrieval, Compression and Analysis*. PhD thesis, Columbia University, 1997.