

PRIP-TR-77

March 6, 2003

A 3D-polar Coordinate Colour Representation Suitable for Image Analysis

Allan Hanbury and Jean Serra

Abstract

The processing and analysis of colour images has become an important area of study and application. The representation of the RGB colour space in 3D-polar coordinates (hue, saturation and brightness) can sometimes simplify this task by revealing characteristics not visible in the rectangular coordinate representation. The literature describes many such spaces (HLS, HSV, etc.), but many of them, having been developed for computer graphics applications, are unsuited to image processing and analysis tasks. We describe the flaws present in these colour spaces, and present three prerequisites for 3D-polar coordinate colour spaces well-suited to image processing and analysis. We then derive 3D-polar coordinate representations which satisfy the prerequisites, namely a space based on the L_1 norm which has efficient linear transform functions to and from the RGB space; and an improved HLS (IHLS) space. The most important property of this latter space is a “well-behaved” saturation coordinate which, in contrast to commonly used ones, always has a small numerical value for near-achromatic colours, and is completely independent of the brightness function. Three applications taking advantage of the good properties of the IHLS space are described: the calculation of a saturation-weighted hue mean and of saturation-weighted hue histograms, and feature extraction using mathematical morphology.

1 Introduction

The number of applications requiring colour image processing and analysis is growing continuously, in particular in the multimedia domain. Important problems currently under study include the accurate reproduction of colours on different output devices [15], and the development of reliable algorithms for processing colour images [24]. The development of these algorithms is made more difficult by the vectorial nature of colour coordinates, as well as by the large number of colour representation models available, allowing a certain colour to be equivalently encoded by many sets of coordinates.

Representations of the RGB colour space in terms of hue, saturation and brightness coordinates are often used. These representations suffer from some defects, such as the presence of unstable singularities and non-uniform distributions of their components, as described by Kender [16]. Nevertheless, they can be more intuitive than the RGB representation, and could reveal features of an image which are not clearly visible in this representation. They do not have all the good properties of the $L^*a^*b^*$ or $L^*u^*v^*$ spaces, but are simpler to calculate, and do not require any calibration information. Even though the transformation from RGB to hue, saturation and brightness coordinates is simply a transformation from a rectangular colour coordinate system (RGB) to a three-dimensional polar (cylindrical) coordinate system, one is faced with a bewildering array of such transformations described in the literature (HSI, HSV, HLS, etc. [25]). This results in a confusing choice between models which essentially all offer the same representation. Indeed, physicists have, as an aid to problem solving, been routinely converting between rectangular and 3D-polar coordinate systems for many decades without similar model choice problems. Is it not possible to achieve this simplicity in colour representation?

In this technical report, we first discuss the existing hue, saturation and brightness transforms and their shortcomings when used in image processing or analysis (section 2). Section 3 describes, in terms of vector independence and vector norms, the prerequisites for a useful 3D-polar coordinate colour representation, and section 4 summarises the basic properties of the RGB vector space used in the derivation of this representation. We present a geometrical derivation, in section 5, of an expression for calculating the saturation of an RGB vector. Sections 6 and 7 consider the consequences of restricting oneself to using respectively only the L_2 and L_1 vector norms in the derivation of the 3D-polar coordinates. An improved HLS (called IHLS) coordinate system is then suggested in section 8. We give a brief comparison, in section 9, of the distributions of the saturation and chroma expressions discussed. Efficient transformations between the RGB space and IHLS system are presented in section 10. Finally, three application examples using the suggested coordinates are given in section 11: the calculation of hue statistics, saturation-weighted hue histograms, and feature extraction in colour images using mathematical morphology.

2 Existing colour space transforms

In this section, we first review the standard definition of the terms used to describe colour intensity (section 2.1). An overview of the method of converting RGB coordinates to 3D-polar coordinates is then given (section 2.2). Lastly, we discuss the problems arising when using the

currently popular versions of these spaces in image analysis, and the reasons for which they occur (section 2.3).

In the RGB space, colours are specified as vectors (R, G, B) which give the amount of each red, green and blue primary stimulus in the colour. For convenience, we take $R, G, B \in [0, 1]$ so that the valid coordinates form the cube $[0, 1] \times [0, 1] \times [0, 1]$. For digital images, these coordinates are usually 8-bit integers, but it is easy to generalise from $[0, 1]$ to any range of values.

2.1 Brightness, luminance and lightness

The terms *brightness*, *luminance* and *lightness* are used to describe the intensity of a colour. They are often used interchangeably, although they have specific definitions assigned to them by the CIE (International Commission on Illumination). These standard definitions are [5, 22]:

Brightness: Attribute of a visual sensation according to which an area appears to emit more or less light. This attribute is measured subjectively and has no units of measurement.

Luminance: Luminance is the luminous intensity per unit surface area, measured in the SI units of candela per square metre (cd/m^2). *Luminous intensity* (unit: Candela) is radiant intensity (unit: watts/steradian) weighted by the spectral response of the human eye. The luminance measure therefore takes into account that for three light sources which appear red, green and blue, and have the same radiant intensity in the visible spectrum, the green one will *appear* the brightest, and the blue one the dimmest.

In the international recommendation for the high definition television standard [14], the following weights for calculating luminance from the (non gamma-corrected) red, green and blue components are given:

$$Y(\mathbf{c}) = 0.2126R + 0.7152G + 0.0722B \quad (1)$$

Lightness: A measurement which takes into account the non-linear response of the human eye to luminance. A source having a luminance of only 18% of a reference luminance appears about half as bright [21]. The CIE uses lightness in their $L^*a^*b^*$ and $L^*u^*v^*$ spaces.

To avoid repeatedly writing out all three of these terms, we assume that luminance and lightness functions are part of the set of brightness functions, and hence are included when only brightness functions are mentioned.

2.2 Overview of the transformation from RGB to 3D-polar coordinates

The basic idea behind the transformation from an RGB coordinate system to a hue, saturation and brightness coordinate system is described by Levkowitz and Herman [17]. One first places a new axis in the RGB space between $(0, 0, 0)$ and $(1, 1, 1)$. This axis passes through all the achromatic points (i.e. those with $R = G = B$), and is therefore called the *achromatic axis*. One then chooses a function $L(\mathbf{c})$ which calculates the brightness, luminance or lightness of colour $\mathbf{c} = (R, G, B)$. The form chosen for $L(\mathbf{c})$ defines the shape of the iso-brightness surfaces. The iso-brightness surface l contains all the points with a brightness of l , i.e. all the points satisfying

the relation $\{c = (R, G, B) : L(c) = l\}$. These iso-brightness surfaces are then projected onto a plane perpendicular to the achromatic axis and intersecting it at the origin, called the *chromatic plane* as it contains all the colour information. The *hue* and *saturation* or *chroma* coordinates of each point are then determined within the plane, where the hue corresponds to the angular coordinate around the achromatic axis¹, and the saturation or chroma corresponds to a distance from the achromatic axis.

To visualise the shape of the resulting space, the points of each iso-brightness surface l are projected onto a chromatic plane intersecting the achromatic axis at l . The solid corresponding to a colour space is constructed out of the sub-regions of each chromatic plane containing projected points. The form of this solid depends on the brightness function chosen, as is now demonstrated for the HSV and HLS models (based on the discussion in [17]).

2.2.1 The HSV model

The brightness function used in the HSV model is

$$L_{\text{HSV}}(c) = \max(R, G, B) \quad (2)$$

To visualise the iso-brightness surface corresponding to brightness l , begin with the cube having principal diagonal between $(0, 0, 0)$ and (l, l, l) . The iso-brightness surface consists of the three faces of the cube which contain the vertex at (l, l, l) , an example of which is shown in figure 1a. When this surface is projected onto the chromatic plane, one obtains a hexagon. It is clear that the surface areas of these hexagons are proportional to l , and hence the solid created by stacking these hexagons is a hexcone. A vertical slice along the achromatic axis through the HSV colour space is shown in figure 10a.

For completeness, we give the commonly used HSV model saturation and hue expressions

$$S_{\text{HSV}}(c) = \begin{cases} \frac{\max(R, G, B) - \min(R, G, B)}{\max(R, G, B)} & \text{if } \max(R, G, B) \neq 0 \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

$$H'_{\text{HSV}}(c) = \begin{cases} \text{undefined} & \text{if } S_{\text{HSV}} = 0 \\ \frac{G - B}{\max(R, G, B) - \min(R, G, B)} & \text{if } R = \max(R, G, B) \\ 2 + \frac{B - R}{\max(R, G, B) - \min(R, G, B)} & \text{if } G = \max(R, G, B) \\ 4 + \frac{R - G}{\max(R, G, B) - \min(R, G, B)} & \text{if } B = \max(R, G, B) \end{cases} \quad (4)$$

H'_{HSV} is multiplied by 60° to get a hue value H_{HSV} in degrees.

2.2.2 The HLS model

The brightness function used in the HLS model is

$$L_{\text{HLS}}(c) = \frac{\max(R, G, B) + \min(R, G, B)}{2} \quad (5)$$

¹The fact that hue is an angular value, and therefore has a periodicity of 360° , is often ignored in colour image analysis. One cannot simply take the minimum of the hue to be 0° and the maximum to be 360° , as these coordinates correspond to the same point on the circle! Furthermore, even though the origin is traditionally chosen to be in the red part of the hue circle, this does not imply that red is more important than the other colours. Further discussion can be found in [13, 20].

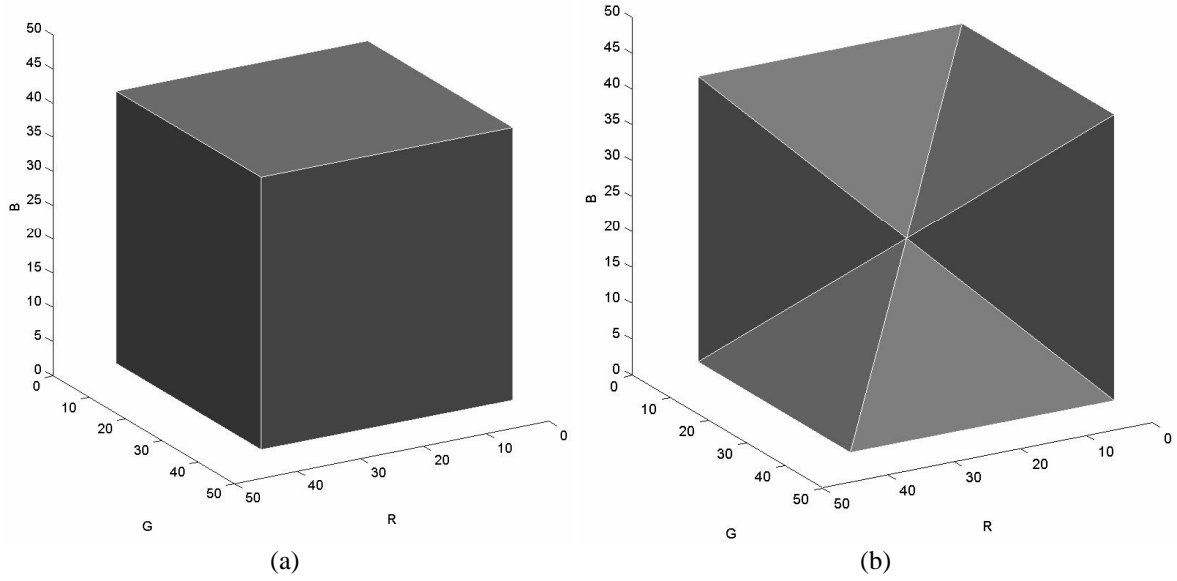


Figure 1: Example iso-brightness surfaces for two digital colour spaces in which the coordinates are encoded using 8 bits. (a) HSV for $l = 40$. (b) HLS for $l = 20$.

For brightness l , one can visualise the iso-brightness surface by starting from the cube with principal diagonal between $(0, 0, 0)$ and $(2l, 2l, 2l)$ for $l \leq 1/2$, or with principal diagonal between $(2l - 1, 2l - 1, 2l - 1)$ and $(1, 1, 1)$ for $l > 1/2$. The iso-brightness surface consists of the six triangles inside the cube with edges formed by the lines between the point (l, l, l) and the six vertices of the cube which are not on the achromatic axis. An example of this iso-brightness surface is shown in figure 1b. The projection of this surface onto the chromatic plane also results in a hexagon, except that for this model, the largest hexagon is found at $l = 1/2$. The solid produced by stacking these hexagons is therefore a double-hexcone. A vertical slice along the achromatic axis through the HLS colour space is shown in figure 10c.

For the HLS model, the hue calculated as for the HSV model (equation 4), and the commonly used saturation expression is

$$S_{\text{HLS}} = \begin{cases} 0 & \text{if } \max(R, G, B) = \min(R, G, B) \\ \frac{\max(R, G, B) - \min(R, G, B)}{\max(R, G, B) + \min(R, G, B)} & \text{if } L_{\text{HLS}} \leq \frac{1}{2} \\ \frac{\max(R, G, B) - \min(R, G, B)}{2 - [\max(R, G, B) + \min(R, G, B)]} & \text{otherwise} \end{cases} \quad (6)$$

2.3 Problems arising when using these spaces for image analysis

The HSV and HLS colour spaces were developed during the 1970's for easy numerical specification of colours in computer graphics applications [26]. In this context, the hexcone and double-hexcone shapes of the spaces are inconvenient, as it would be easy for a user to accidentally specify coordinates which lie outside the colour gamut. As computers of the time were not very speedy, additional checking to avoid this would have been unacceptable, so the solution of expanding the colour spaces into cylindrical form was adopted. This is easily done by defining

the saturation as the ratio of the actual distance of a point from the achromatic axis to the maximum distance for the corresponding brightness value. The HSV cone and HLS double-cone are thereby expanded into cylinders. Vertical slices through these cylinders are shown in figures 10b for the HSV space, and 10d for the HLS space, to be compared with the slices through the conic and bi-conic versions of the spaces in figures 10a and 10c respectively. Indeed, the commonly used saturation expressions (equations 3 and 6) describe such cylindrically shaped spaces. The dependence of these saturation measures on the corresponding brightness is easily seen. For example, given the definition of brightness in the HSV space, the first level of the HSV saturation expression (equation 3) can easily be rewritten as

$$1 - \frac{\min(R, G, B)}{L_{\text{HSV}}(\mathbf{c})} \quad (7)$$

The HLS saturation can also easily be rewritten in terms of L_{HLS} . These cylindrically shaped colour spaces have unfortunately been adopted by the image analysis community (and implemented in image analysis software²), leading to the widespread use of an unsuitable definition of saturation.

To demonstrate the unsuitability of the cylindrically shaped spaces for image processing and analysis, we use the colour image in figure 11a. This image was captured under slightly non-uniform lighting conditions, so that not all the pixels which look white have RGB coordinates of exactly (1, 1, 1). The upper part of the image was then inverted by subtracting the values in each of the R , G and B channels from the maximum possible values (i.e. 255 for this 8-bit image). The HSV saturation image calculated from this colour image is shown in figure 2b. The lower part of this saturation image, corresponding to the white region in the initial colour image, has a saturation of around zero, as expected. However, some of the black pixels in the upper part of the colour image are shown as being fully saturated. This patently contradicts the definition of saturation, which states that saturation should be low for almost-achromatic colours, and zero for greylevels. The reason is that some of the black pixels have small non-zero R , G or B components. The expansion of the HSV cone into a cylinder (demonstrated in figures 10a and b) results in these pixels getting artificially high saturation values. One therefore has the ridiculous situation where some of the black pixels are shown as being more highly saturated than the colourful regions that they surround. Because of the double-cone shape of the HLS colour space, its expansion into a cylinder produces spurious high values of saturation in both the high and low brightness regions, as shown in figure 2c. This is particularly noticeable for the orange region at the bottom of the image, on which two of the white letters are shown as having saturation values equal to the surrounding orange colour.

This demonstrates that two of the common assumptions about these models are not true when the cylindrically shaped versions are used:

1. Saturation is defined as the chromaticity of a colour, so that pixels which appear black, white or grey should have a lower saturation than colourful pixels. As was shown in the example above, pixels which appear black or white often have maximal saturation values when one of the cylindrically shaped spaces is used.

²Software already used by the author which implement cylindrically shaped colour models include: Matlab release 12.1, Aphelion 3.0, Optimas 6.1 and Paint Shop Pro 7.

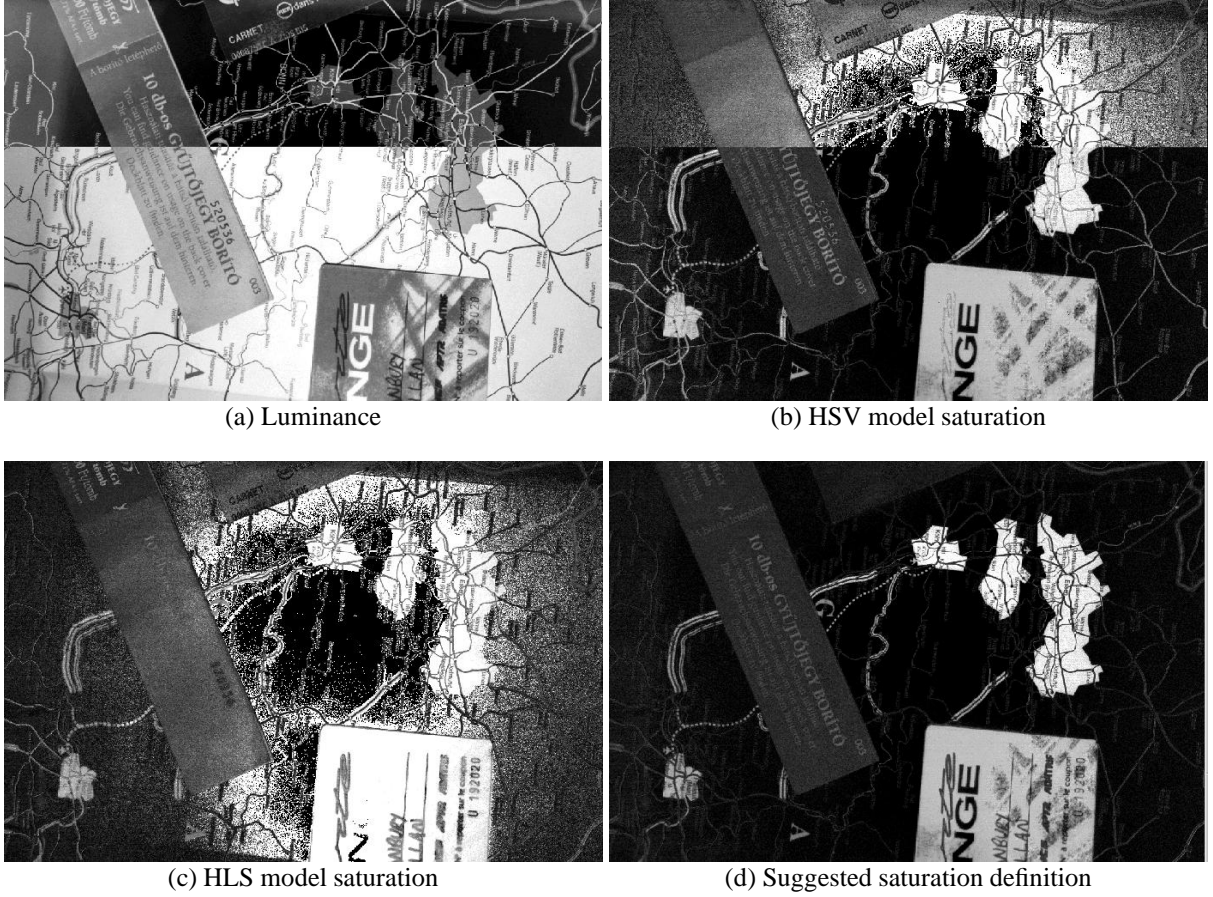


Figure 2: 3D-polar coordinate components of figure 11: (a) Luminance. (b) HSV model saturation. (c) HLS model saturation. (d) The suggested saturation measure.

2. It is often said that these spaces separate chrominance (hue and saturation) and brightness information. However, use is made of the brightness function to normalise the saturation in the cylindrically shaped spaces (as shown in equation 7). It is clear that the saturation values therefore depend critically on the brightness function chosen (demonstrated by the large differences between figures 2b and c).

We now consider two cases of the confusion that the cylindrical forms of the colour spaces can cause. Demarty and Beucher [7] applied a constant saturation threshold in the cylindrically shaped HLS space (figure 10d) 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 10d, and it is clear that this does not correspond to a constant saturation. Demarty [6] later improved the threshold by using a hyperbola in the cylindrical HSV space (figure 10b), which corresponds to a constant threshold in the conic HSV space (figure 10a). Smith [27] makes the assumption that the cylindrical HSV space is perceptually uniform when a Euclidean metric is used, but upon examining figure 10b, 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. Such an assumption is almost certainly truer in the conically shaped version of the space. This problem also affects the quantisations of the cylindrical HSV

space in which an equal number of saturation bins are used in the high and low brightness regions of the HSV space. Almost imperceptible colour changes in the low-brightness region are quantised into the same number of bins as highly-visible changes in the high-brightness region.

2.4 Removal of the brightness dependence of the saturation expression

The simplest way of avoiding the disadvantages tied to the cylindrically shaped spaces is to remove the brightness normalisation from the saturation expressions, hence reverting to the original shapes of the spaces. Removing this brightness dependence from the saturation for the HSV model is simply done by multiplying equation 3 by the brightness L_{HSV} , giving

$$S^{\text{NC}} = \max(R, G, B) - \min(R, G, B) \quad (8)$$

where the superscript ‘NC’ indicates that this is the non-cylindrical version. For the HLS space, removing the brightness dependence is slightly more complex due to its double-cone shape. The non-cylindrical saturation is

$$S_{\text{HLS}}^{\text{NC}} = S_{\text{HLS}} \left[1 - 2 \left| \frac{1}{2} - L_{\text{HLS}} \right| \right] \quad (9)$$

which after some manipulation also reduces to equation 8. The equivalence of these non-cylindrical saturation expressions is tantalising, and we show that it is in fact derivable from the basic definition of saturation in section 5.

3 Vectors, norms and independence

We now define clearly what the notions of *vector space*, *norm* and *independence* contribute to colour image representations. In the following, the RGB space is modeled by the Euclidean space \mathbb{R}^3 , with its projections, orthogonality, etc., but we equip it successively with different norms, including amongst them the Euclidean norm.

The vector space notion associates a point $\mathbf{c} = (R, G, B)$ to the vector \vec{oc} . It defines:

- The sum of a number of vectors as being the vector made up of the sum of the vector components.
- The product of a vector and a scalar as being obtained by multiplying each component of the vector by the scalar.

Note that these operations transform vectors into vectors, and not into numbers. We know that every vector can be uniquely written in terms of its components for each system of axes. Therefore, starting from the unit cube with coordinates $0 \leq R \leq 1$, $0 \leq G \leq 1$ and $0 \leq B \leq 1$, we define the diagonal between $(0, 0, 0)$ and $(1, 1, 1)$ as the achromatic axis, and the plane intersecting the origin and perpendicular to this axis as the chromatic plane, which contains all information on the colour. Hence, the point \mathbf{c} can be written as $\vec{oc} = \vec{or} + \vec{og} + \vec{ob}$, or equivalently as $\vec{oc} = \vec{oc}_d + \vec{oc}_p$, where \mathbf{c}_d and \mathbf{c}_p are the projections of \mathbf{c} onto respectively the achromatic axis and the chromatic plane.

Can we therefore say that the vectors \mathbf{c}_d and \mathbf{c}_p , which are orthogonal, are also independent? The response depends on the meaning which we attribute to the adjective “independent”. If we refer to a possible link between the two projections \mathbf{c}_d and \mathbf{c}_p , they are obviously not independent: the points of low brightness always have low colour saturation. But the independence can also signify something else, for example that the parameters that we associate with \mathbf{c}_p (saturation, hue) don’t affect those associated with \mathbf{c}_d . In this case, if two differently coloured points \mathbf{c} and \mathbf{c}' have the same projection \mathbf{c}_p , they have the same saturation and the same hue. In order to have a colour representation adapted to image analysis, we therefore propose the following prerequisite:

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

We could go further and require that two points which have the same projection onto the achromatic axis have the same intensity. However, this would limit one to symmetric functions of R , G and B , excluding notably weighted expressions such as the luminance (equation 1).

Another useful concept on which we now base our discussion is that of the norm. It associates a parameter, which we call ξ , with every vector. This parameter is zero or positive, and its magnitude becomes larger as point \mathbf{c} moves further away from the origin, i.e. $\xi(\lambda\mathbf{c}) = \lambda\xi(\mathbf{c})$, in which $\lambda \geq 0$ is a weighting factor. Furthermore, the norm links the addition of vectors to that of numbers by the classic triangular inequality

$$\xi(\mathbf{c} + \mathbf{c}') \leq \xi(\mathbf{c}) + \xi(\mathbf{c}') \quad (10)$$

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 inequality of equation 10 should be satisfied. Lastly, it is equivalent to say that the vector \mathbf{c} is zero or that its norm is zero

$$\mathbf{c} = \mathbf{0} \Leftrightarrow \xi(\mathbf{c}) = 0 \quad (11)$$

When this last condition is not satisfied, we refer to a *semi-norm*. We note that in the triangular inequality 10 the two ‘+’ symbols do not have the same meaning: the first is with respect to vectors, and the second with respect to numbers. The same is true for the two zeros in equation 11.

We will consider in more detail the norms L_1 and L_2 , and the semi-norm $\max - \min$ (proof that it is a semi-norm is given in appendix B). Use of the L_2 norm leads to conversion formulae which are quadratic and rather difficult to invert. Conversely, the distance associated with this norm is the Euclidean distance, which is well-known and convenient to work with.

The L_1 norm has already made its appearance in colour space conversions, but without announcing itself as such. We see it for example in [3] and [9] for the achromatic axis, and in the standard triangle colour model [17]. Its associated distance is less intuitive than the Euclidean distance, but faster to implement and usually just as precise. We note lastly 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 leads to the second prerequisite:

Second prerequisite: The brightness parameters associated with colour vector \mathbf{c} and with its projection \mathbf{c}_p must be norms.

In addition to the two prerequisites already mentioned, it is convenient to introduce a third constraint, less fundamental and suggested by practical experience. It is extremely convenient if one is able to return to an RGB space image representation at the end of an image processing task, which leads us to propose the third prerequisite:

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

If we examine the HLS system in the light of the first two prerequisites, the basic reasons for the criticisms presented in section 2.3 become clear. In the HLS system, neither the saturation nor the brightness are norms, and in addition, there is no independence between the achromatic axis and the chromatic plane: it's almost impossible to develop a worse colour space.

One can show the lack of independence by considering the points $\mathbf{c} = (1/4, 1/4, 0)$ and $\mathbf{c}' = (1/2, 1/2, 1/4)$, which both project onto the same point \mathbf{c}_p on the chromatic plane. Their HLS saturations are given by $\frac{\max - \min}{\max + \min}$ as their brightness values are $\leq 1/2$. The first has an HLS saturation of 1 and the second of $1/3$ (the latter point has a smaller saturation as its brightness value is higher than that of the other, we once again come across a problem in the commonly used cylindrical form). Not only does this representation create artificial differences between points, but it fails to discriminate between points which are different: all points with brightness $\leq 1/2$ and with $\min = 0$ have the same saturation.

To show that the brightness $L = \frac{\max + \min}{2}$ does not satisfy the triangular inequality, we can use the points $\mathbf{c} = (1/2, 1/2, 0)$ and $\mathbf{c}' = (0, 1/2, 1/2)$, both having HLS brightness values equal to $1/4$, while the brightness of $\mathbf{c} + \mathbf{c}'$ is equal to $3/4$. Finally, the HLS saturation is not a norm either, as the points $\mathbf{c} = (1/3, 2/3, 1/3)$ and $\mathbf{c}' = (2/3, 1/3, 1/3)$ both have saturation of $1/3$ while their sum has a value of 1 (the term $\max - \min$, on the other hand, stays the same).

4 Properties of the space under consideration

In the RGB unit cube, the achromatic axis is placed between the points $(0, 0, 0)$ and $(1, 1, 1)$, and contains the colours for which $R = G = B$. The chromatic information is entirely encoded in the chromatic plane, perpendicular to the achromatic axis and intersecting it at the origin. Every vector \mathbf{c} of the RGB unit cube is decomposed into the vectorial sum of its projections \mathbf{c}_d onto the achromatic axis and \mathbf{c}_p onto the chromatic plane

$$\mathbf{c} = \mathbf{c}_d + \mathbf{c}_p \quad (12)$$

in which all chromatic information is encoded in the vector \mathbf{c}_p . We adopt the notation in which all the vectors projected onto the chromatic plane take a subscript p . Hence \mathbf{r}_p and \mathbf{g}_p represent respectively the projections onto the chromatic plane of the pure red vector \mathbf{r} and pure green vector \mathbf{g} .

The chromatic plane is shown in figure 3a, and we proceed to draw the reader's attention to some of the important features shown in this figure. The hexagon surrounds the regions into which points in the RGB cube are projected, and the circle circumscribing the hexagon has

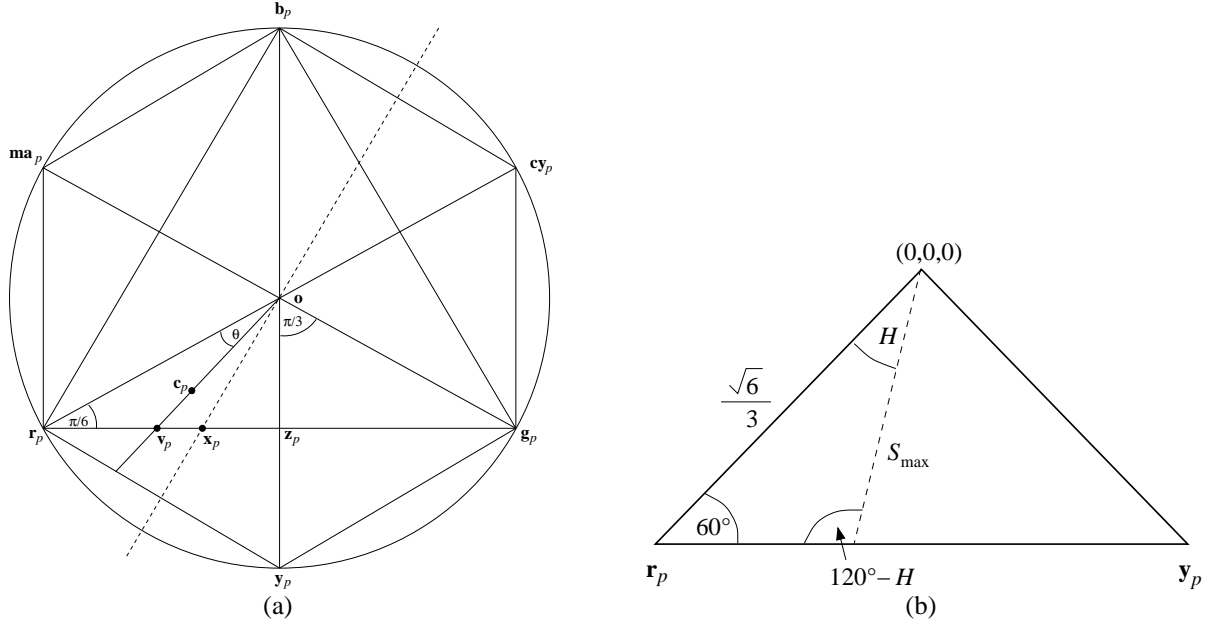


Figure 3: (a) The chromatic plane. (b) The red-yellow sector of the hexagon on the chromatic plane. The lower vertices correspond to the colours red (at the left) and yellow. The angle H takes values between 0° and 60° .

a radius of $\frac{\sqrt{6}}{3}$. If we limit the points projected onto the plane to only those with a specific brightness, then the hexagon has an area smaller than or equal to the one shown. A point c projected onto the chromatic plane has coordinates

$$\mathbf{c}_p = \left[\frac{(2R - G - B)}{3}, \frac{(2G - B - R)}{3}, \frac{(2B - R - G)}{3} \right] \quad (13)$$

The projections onto the chromatic plane of pure red $\mathbf{r} = (1, 0, 0)$, yellow $\mathbf{y} = (1, 1, 0)$ and green $\mathbf{g} = (0, 1, 0)$ have coordinates

$$\mathbf{r}_p = \left(\frac{2}{3}, -\frac{1}{3}, -\frac{1}{3} \right), \mathbf{y}_p = \left(\frac{1}{3}, \frac{1}{3}, -\frac{2}{3} \right), \mathbf{g}_p = \left(-\frac{1}{3}, \frac{2}{3}, -\frac{1}{3} \right) \quad (14)$$

The RGB unit cube is shown in figure 4. We point out some useful correspondences between regions of the cube and their projections onto the chromatic plane. Points for which $R \geq G$ form the half-space limited by the plane $\mathbf{ow}\mathbf{y}$ which contains \mathbf{r} in figure 4. Its points are projected onto the half-plane limited by $\mathbf{y}_p\mathbf{b}_p$ and containing \mathbf{r}_p in figure 3a. Similarly, the points for which $G \geq B$ form the half-space limited by $\mathbf{ow}(\mathbf{cy})$ which contains \mathbf{g} (figure 4), with points projected onto the half-plane limited by $\mathbf{r}_p(\mathbf{cy})_p$ and containing \mathbf{g}_p . Lastly, the points such that $R + B - 2G = 0$ form the plane passing through the achromatic axis and the line $R + B = 0$ in the plane \mathbf{orb} . This plane cuts the chromatic plane along the line parallel to $\mathbf{r}_p\mathbf{b}_p$ passing through \mathbf{o} .

With respect to colour, the RGB cube is divided into six sectors delimited by the six planes each containing the achromatic axis and one of the three R , G and B axes or one of the diagonals

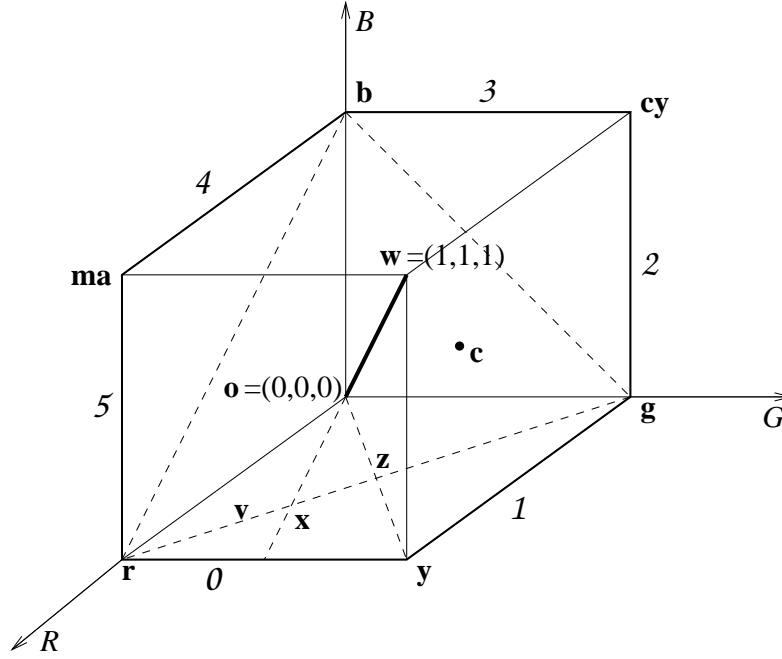


Figure 4: The RGB unit cube. The italic numbers indicate the edges corresponding to the six sectors into which the RGB cube is divided.

in the squares in $B = 0$, $R = 0$ or $G = 0$. The following equation gives the sector of a colour based on the order of magnitudes of the RGB coordinates

$$\lambda(\mathbf{c}) = \begin{cases} 0 & \text{if } R > G \geq B \\ 1 & \text{if } G \geq R > B \\ 2 & \text{if } G > B \geq R \\ 3 & \text{if } B \geq G > R \\ 4 & \text{if } B > R \geq G \\ 5 & \text{if } R \geq B > G \end{cases} \quad (15)$$

The cube edge corresponding to each sector is indicated by the italic numbers in figure 4.

We now assign a polar coordinate system to the chromatic plane, taking the vector \mathbf{r}_p as the origin of the angles. The angular values increase as one moves in an anti-clockwise direction. Every point (i.e. every vector) \mathbf{c} of the RGB unit cube can be equivalently written in terms of RGB Cartesian coordinates, or 3D-polar coordinates ($\|\mathbf{c}_d\|$, $\|\mathbf{c}_p\|$, θ).

One nevertheless has different equations for converting from one system to the other depending on whether one uses the L_2 or L_1 norm. For the L_2 norm

$$\|\mathbf{c}\|^2 = \|R\|^2 + \|G\|^2 + \|B\|^2 \quad (16)$$

and for the L_1 norm, one has

$$|\mathbf{c}| = |R| + |G| + |B| \quad (17)$$

The differences are large enough that we study more precisely, in sections 6 and 7, the advantages and disadvantages of the two approaches.

Finally, for the semi-norm $\delta = \max - \min$, we find for vector $\mathbf{c} = (R, G, B)$ and its chromatic projection \mathbf{c}_p of equation 13,

$$\delta(\mathbf{c}) = \delta(\mathbf{c}_p) = \max(R, G, B) - \min(R, G, B) \quad (18)$$

which means that this semi-norm is exclusively chromatic (it does not see variations in the brightness).

5 Geometric derivation of a saturation term

The *saturation* and *chroma* measurements are associated with the length of the vector \mathbf{c}_p . Their definitions are:

Chroma: The norm of \mathbf{c}_p is used, as done by Carron [2] (who uses the L_2 norm). It assumes its maximum value at the six corners of the hexagon projected onto the chromatic plane. The shape of the resultant space obtained by piling up the hexagons is a hexcone or double-hexcone.

Saturation: 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. The shape of the resultant space is therefore a cone or double-cone. Poor choice of this normalisation factor has led to some of the less than useful saturation definitions currently in use.

We now geometrically derive a saturation coordinate which does not suffer from the disadvantages enumerated in section 2.3 and is therefore much more useful in image processing and analysis. This derivation is based on the one for the Levkowitz and Herman GLHS model [17].

5.1 Basic saturation formulation

To calculate the saturation of a colour represented by a vector \mathbf{c} in the RGB space, we begin by considering the triangle which contains all the colours which have the same hue as \mathbf{c} (iso-hue triangle), shown in figure 5a. The achromatic axis always forms one of the sides of this triangle. The vector $\mathbf{L}(\mathbf{c}) = [l(\mathbf{c}), l(\mathbf{c}), l(\mathbf{c})]$ gives the position on the achromatic axis in RGB coordinates of the brightness value associated with \mathbf{c} . The iso-brightness line associated with \mathbf{c} is the intersection of the iso-hue triangle and $l(\mathbf{c})$ iso-brightness surface, and hence passes through $\mathbf{L}(\mathbf{c})$ and \mathbf{c} . By definition, all the iso-brightness lines in the triangle are parallel. The point with the same hue as \mathbf{c} lying furthest away from the achromatic axis is labeled $\mathbf{q}(\mathbf{c})$. This point necessarily lies on one of the edges of the RGB cube.

Traditionally, the saturation is defined as the fraction given by 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 produces a space in the form of a cylinder. We call this type of saturation a *cylindrical saturation*. The problems inherent in the use of this saturation definition have already been described in section 2.3.

In order to keep the conical form of the space, it is necessary to change the definition of the saturation. In figure 5a, instead of dividing the length of the vector from $\mathbf{L}(\mathbf{c})$ to \mathbf{c} by

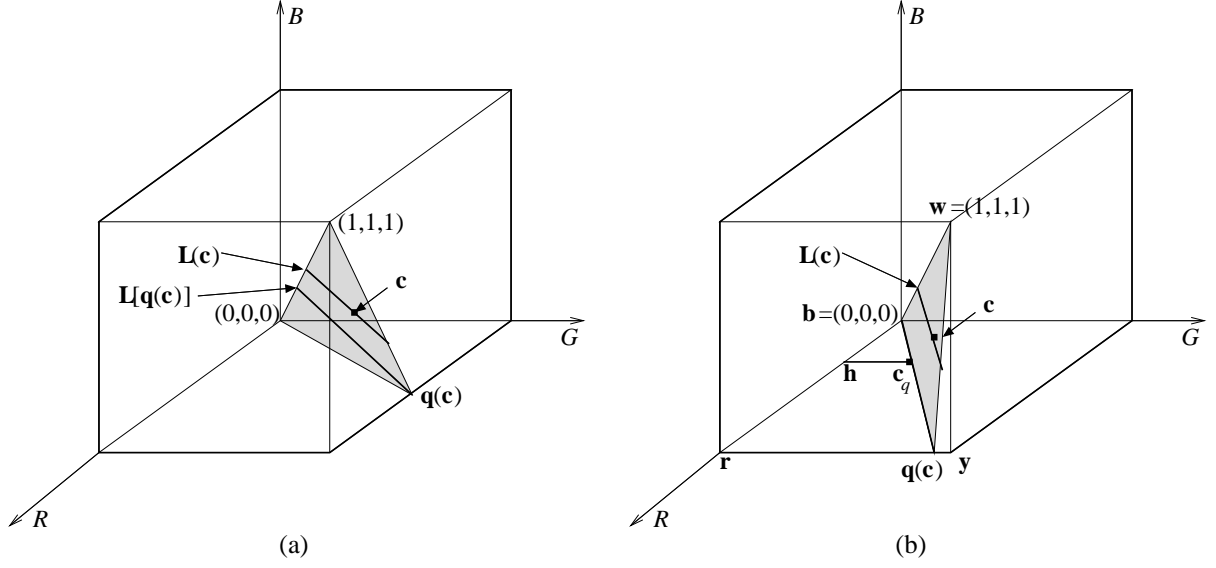


Figure 5: (a) Diagram used in the derivation of a general saturation expression. (b) Diagram used in the derivation of the simpler saturation expression. Both diagrams show the triangle which contains all the points with the same hue as c .

the length of its extension to the edge of the cube, we divide it by the length of the vector between $L[q(c)]$ and $q(c)$. This is the longest vector parallel to the iso-brightness lines, which necessarily intersects the third vertex $q(c)$ of the triangle. We therefore have the following general definition of saturation

$$S = \frac{\|L(c) - c\|}{\|L[q(c)] - q(c)\|} \quad (19)$$

which gives the natural conic or bi-conic form to the 3D-polar coordinate colour space, and which additionally is independent of the choice of the brightness function. A proof of this independence is presented in appendix A. This saturation calculated for figure 11a is shown in figure 2d. It is clear that the defects associated with the cylindrically shaped HSV and HLS models are not present. The colourful regions always have saturation values higher than the surrounding monochromatic background. Furthermore, one would obtain the same saturation values irrespective of the brightness function used.

5.2 A simple expression for the saturation

In this section, we use equation 19 to derive a very simple saturation expression. For this derivation, we choose the iso-brightness surfaces to be parallel to the nearest side of the RGB cube which intersects the origin (which we are free to do due to the independence of the brightness and saturation). This is the $B = 0$ plane for sectors 0 and 1, the $R = 0$ plane for sectors 2 and 3, and the $G = 0$ plane for sectors 4 and 5. The brightness function producing such iso-brightness surfaces is

$$L(c) = \min(R, G, B) \quad (20)$$

At first, we consider only sector 0, which contains point $\mathbf{c} = (R, G, B)$ as shown in figure 5b. The brightness vector of \mathbf{c} is $\mathbf{L}(\mathbf{c}) = (B, B, B)$, due to \mathbf{c} being in sector 0. We project \mathbf{c} onto the $B = 0$ plane resulting in point \mathbf{c}_q which is the same distance from the origin as $\mathbf{L}(\mathbf{c})$ is from \mathbf{c} . The coordinates of \mathbf{c}_q are therefore $(R - B, G - B, 0)$. We then construct the line between \mathbf{h} and \mathbf{c}_q in the $B = 0$ plane parallel to the cube edge between \mathbf{r} and \mathbf{y} , forming two similar triangles with vertices \mathbf{o}, \mathbf{h} and \mathbf{c}_q , and \mathbf{o}, \mathbf{r} and $\mathbf{q}(\mathbf{c})$. The following relation is therefore valid:

$$\frac{\|\mathbf{o} - \mathbf{c}_q\|}{\|\mathbf{o} - \mathbf{q}(\mathbf{c})\|} = \frac{\|\mathbf{o} - \mathbf{h}\|}{\|\mathbf{o} - \mathbf{r}\|} \quad (21)$$

The term on the left is simply the definition of saturation given by equation 19. On the right, $\|\mathbf{o} - \mathbf{r}\| = 1$, and as the coordinates of \mathbf{h} are $(R - B, 0, 0)$, $\|\mathbf{o} - \mathbf{h}\| = R - B$. Hence the saturation $S = R - B$, which in sector 0, is equivalent to

$$S_0 = \max(R, G, B) - \min(R, G, B) \quad (22)$$

The derivation is easily done for the other five sectors to show that equation 22 is valid for them all. The simple saturation expression obtained for the non-cylindrically shaped HSV and HLS spaces (section 2.4), is therefore obtained from the general saturation definition. The $\max - \min$ expression is in fact a semi-norm, as proved in appendix B. We return to this saturation expression in section 8, in which we suggest an improved version of the HLS space. We first consider the consequences of strictly imposing either the L_1 or L_2 norm on the RGB space.

6 The framework of the L_2 norm

We expect this norm to be the best adapted to the problem, as it is based on the Pythagorean theorem, which interprets the norm in terms of vector lengths. In addition, the scalar product which accompanies it is an indispensable tool for calculating angles.

The conversion equations from the RGB coordinate system are easy to determine. We call the norms of the vectors projected onto the achromatic axis and the chromatic plane respectively M_2 and C_2 , both these norms being scaled to the range $[0, 1]$. The angle θ is called H_2 . By using relation 16 and figure 3a, the following can be derived

$$M_2 = \frac{1}{\sqrt{3}} [R^2 + G^2 + B^2]^{1/2} \quad (23)$$

$$C_2 = \sqrt{\frac{3}{2}} \|\mathbf{c}_p\| \quad (24)$$

$$= (R^2 + G^2 + B^2 - RG - RB - BG)^{\frac{1}{2}} \quad (25)$$

Note that C_2 is a measurement of the chroma, as it is simply the norm of \mathbf{c}_p multiplied by a constant. We can convert this chroma into a saturation by applying equation 19 in the chromatic plane, where it is equivalent to

$$S_2 = \frac{\|\mathbf{c}_p\|}{S_{\max}} \quad (26)$$

in which S_{\max} is 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 . The red-yellow sector of this hexagon is reproduced in figure 3b, in which the upper vertex is at the origin $(0, 0, 0)$, the lower left vertex is the projected red vector \mathbf{r}_p , and the third vertex the projected yellow vector \mathbf{y}_p . It is simple to show using figure 3b that

$$S_{\max} = \frac{\sqrt{2}}{2 \sin(120^\circ - H)} \quad (27)$$

for $0^\circ \leq H < 60^\circ$. To make this equation valid for the values of $H \in [0^\circ, 360^\circ)$, it is sufficient to replace the H in the equation by

$$H^* = H - k \times 60^\circ \text{ where } k \in \{0, 1, 2, 3, 4, 5\} \text{ so that } 0^\circ \leq H^* \leq 60^\circ \quad (28)$$

Note that this saturation expression (equation 26) gives exactly the same values as the max – min expression, as they are both derived from the same definition.

The calculation of the angle θ in figure 3a is done in terms of the scalar product between the vectors \mathbf{c}_p and \mathbf{r}_p as

$$\theta = \arccos \left[\frac{\mathbf{r}_p \cdot \mathbf{c}_p}{\|\mathbf{r}_p\| \|\mathbf{c}_p\|} \right] \quad (29)$$

$$= \arccos \left[\frac{R - \frac{1}{2}G - \frac{1}{2}B}{(R^2 + G^2 + B^2 - RG - RB - BG)^{\frac{1}{2}}} \right] \quad (30)$$

in which $\mathbf{r}_p \cdot \mathbf{c}_p$ indicates the scalar product of the two vectors. The possible values for θ are between 0° and 180° , and it is therefore necessary to expand this range of values by using

$$H_2 = \begin{cases} 360^\circ - \theta & \text{if } B > G \\ \theta & \text{otherwise} \end{cases} \quad (31)$$

Formally, the problem is solved. The variables M_2 , C_2 , S_2 and H_2 are expressed in terms of R , G and B . Nevertheless, irrespective of the theoretical equivalence of the two systems, the inverse transformation is not simple. It is desirable to simplify this pure L_2 norm system either by using another norm, which is considered in the next section, or by using a mixture of different norms, discussed in section 8.

7 The framework of the L_1 norm

In this section, we continue to use the same vector space, with the decomposition $\mathbf{c} = \mathbf{c}_p + \mathbf{c}_d$, but we assign the L_1 norm to the vectors of the space.

7.1 Brightness and chroma

Because the R , G and B coordinates are greater than or equal to zero, the L_1 norm (equation 17) of the vector \mathbf{c}_d is simply the sum of the R , G and B components. As we wish the value of the brightness to be in the range $[0, 1]$, we take it to be the arithmetic mean of the components of \mathbf{c}

$$M_1 = \frac{1}{3} (R + G + B) \quad (32)$$

We note that if two points \mathbf{c} and \mathbf{c}' have the same projection \mathbf{c}_d on the achromatic axis, we have $|\mathbf{c}\mathbf{c}'| = 0$, and by application of the triangular inequality one finds that $M_1(\mathbf{c}) = M_1(\mathbf{c}')$.

The chroma C_1 is defined as being proportional to the L_1 norm of the vector \mathbf{c}_p , that is

$$C_1 = \frac{1}{4} [|2R - G - B| + |2G - B - R| + |2B - R - G|] \quad (33)$$

The constant ensures that the chroma values lie in the range $[0, 1]$, as the expression within the parentheses has a maximum value of 4 obtained when two of the components have an extremal value, and the third has the opposite extremal value. The saturation is zero when $R = G = B$, i.e. when the point \mathbf{c} lies on the achromatic axis.

In order to remove the absolute values, linked to the choice of the L_1 norm, from equation 33, it is necessary to find the maximum, median and minimum of (R, G, B) , which we denote as max, mid and min. As equation 33 is symmetric in terms of R , G and B , it is sufficient to adopt a convention, for example

$$1 \geq R \geq G \geq B \geq 0 \quad (34)$$

and, in the calculation, to replace R by max, G by mid and B by min. When the component order in relation 34 is true, the first term of equation 33 is positive and the third is negative. The second term, which has a variable sign, distinguishes between two cases:

1. if $B + R \geq 2G$, or equivalently $M_1 \geq G$, then

$$C_1 = \frac{1}{2} [(R - G) + (R - B)] = \frac{3}{2} (R - M_1) \quad (35)$$

2. if $B + R \leq 2G$, or equivalently $M_1 \leq G$, then

$$C_1 = \frac{1}{2} [(R - B) + (G - B)] = \frac{3}{2} (M_1 - B) \quad (36)$$

For $B + R = 2G$, we find for both forms:

$$\max - \min = R - B = \frac{4}{3} C_1 \quad (37)$$

The hue, being an angle, is calculated in the same way as for the L_2 norm, using equations 30 and 31.

In summary, despite the presence of two cases for the saturation, the transformation equations 32, 35 and 36 make up a linear system much simpler than in the case of the L_2 norm. The hue calculation is the most complex, and we now develop an approximation of the trigonometric hue.

7.2 Simplified calculation of the hue in the L_1 space

The following approximation is largely based on the approach presented in [17]. Firstly, we limit ourselves to vectors having $R > G \geq B$ and non-zero chroma. Their projections onto the chromatic plane form the triangle $\mathbf{o}_p \mathbf{y}_p$ in figure 3a. Because the hue origin is the vector \mathbf{r}_p , the

angle θ in figure 3a varies from 0 to $\pi/3$ in radians, or conventionally from 0 to 1. To approximate θ , we begin with the hue-fraction equation in the HLS system, that is $H = \frac{G-B}{\max - \min}$, and we transpose it *mutatis mutandis*, that is to say by replacing $\max - \min$, which corresponds to the saturation in the HLS system, by our corresponding L_1 norm chroma expression (differing only by a factor) and taking into account the duality of the two chroma expressions. More precisely, because the line $\mathbf{r}_p \mathbf{g}_p$ is parallel to the line $\mathbf{r} \mathbf{g}$ in the $B = 0$ plane, when θ varies from 0 to 1, the point \mathbf{v} of the $B = 0$ plane describes the half-diagonal $\mathbf{r} \mathbf{z}$, and its projection \mathbf{v}_p describes the segment $\mathbf{r}_p \mathbf{z}_p$ of the chromatic plane. The point \mathbf{x} , the intersection of the plane $R + B + 2G = 0$ (which contains the achromatic axis) and the line $\mathbf{r} \mathbf{g}$, divides the two zones having different chroma definitions. In the projection, the point $\mathbf{x} = (2/3, 1/3, 0)$ gives the point \mathbf{x}_p , which corresponds to the value $\theta = 1/2$.

As the value $\max - \min$ of the HLS system corresponds to the saturation, we replace it here by the L_1 norm chroma C_1

$$\varphi = \frac{\ell}{C_1} (G - B) \quad (\text{with } R + B \geq 2G) \quad (38)$$

We have $\varphi = 0$ for $\mathbf{r} = (1, 0, 0)$. The factor ℓ is determined by the condition of having $\varphi = 1/2$ at $\mathbf{x} = (2/3, 1/3, 0)$, which gives $\ell = 3/4$. When $R + B \leq 2G$, the duality suggests the replacement of $G - B$ by $R - G$, and φ by $1 - \varphi$, or

$$\varphi = 1 - \frac{3}{4} \frac{R - G}{C_1} \quad (\text{with } R + B \leq 2G) \quad (39)$$

In fact, at point \mathbf{x}_p , the chroma C_1 takes the same value of $1/2$ in the two modes, and we find $\varphi = 1/2$. Finally, at the extremity of the range of θ , we see by using point \mathbf{y} or point \mathbf{z} , that $\varphi = 1$.

We can reduce equations 38 and 39 which define φ to a single equation by making use of the critical element $R + B - 2G$. We find

$$B + R - 2G = C_1 (1 - 2\varphi) \quad (40)$$

which demonstrates the equivalence relations

$$B + R - 2G \geq 0 \iff 0 \leq \varphi \leq \frac{1}{2} \quad (41)$$

$$B + R - 2G \leq 0 \iff \frac{1}{2} \leq \varphi \leq 1 \quad (42)$$

A simple numerical experiment has shown that the maximum difference between this hue approximation φ and the trigonometric hue is the same as for the approximation suggested in [17], that is 1.12° .

7.3 Colour space conversions in the L_1 space

The preceding results lead to the following conversion formulae for the conversion $(R, G, B) \rightarrow (M_1, C_1, \varphi)$

$$\begin{cases} M_1 = \frac{1}{3} (R + G + B) \\ C_1 = \frac{1}{2} (2R - G - B) = \frac{3}{2} (R - M_1) & \text{if } B + R \geq 2G \\ C_1 = \frac{1}{2} (R + G - 2B) = \frac{3}{2} (M_1 - B) & \text{if } B + R \leq 2G \\ \varphi = \frac{1}{2} - \frac{B + R - 2G}{2C_1} \end{cases} \quad (43)$$

When $B + R \geq 2G$, or equivalently, when $0 \leq \varphi \leq 1/2$, the transformation is inverted as

$$\begin{cases} R = M_1 + \frac{2}{3}C_1 \\ G = M_1 - \frac{1}{3}C_1 + \frac{2}{3}C_1\varphi \\ B = M_1 - \frac{1}{3}C_1 - \frac{2}{3}C_1\varphi \end{cases} \quad (44)$$

and for $B + R \leq 2G$ (or $1/2 \leq \varphi \leq 1$) as

$$\begin{cases} R = M_1 + C_1 - \frac{2}{3}C_1\varphi \\ G = M_1 - \frac{1}{3}C_1 + \frac{2}{3}C_1\varphi \\ B = M_1 - \frac{2}{3}C_1 \end{cases} \quad (45)$$

The domain on which the equations are defined $1 \geq R \geq G \geq B \geq 0$ (without the achromatic axis) corresponds to the tetrahedron **oryw** in the figure 4. The coefficients of system 43 were chosen so that M_1 , C_1 and φ vary between 0 and 1, which does not necessarily mean that they always correspond to points inside the tetrahedron. The following equivalence relation is nevertheless easily verifiable

$$1 \geq R \geq G \geq B \geq 0 \iff \frac{2}{3}C_1 \leq M_1 \leq 1 - \frac{2}{3}C_1 \quad (46)$$

In practice, this condition is not too limiting, as it is simple to prevent operators applied to vector **c**, expressed as (M_1, C_1, φ) from giving a result outside the RGB cube.

The last case to study is that in which the vector **c** lies on the achromatic axis, i.e. the case for which $R = G = B$. The system 43 is no longer valid, as we introduce a division by zero, and must be replaced by

$$\begin{cases} M_1 &= \frac{1}{3}(R + G + B) \\ C_1 &= \frac{1}{2}(2R - G - B) = \frac{1}{2}(R + G - 2B) \\ C_1(1 - 2\varphi) &= B + R - 2G \end{cases} \quad (47)$$

which shows that φ is indeterminate. This does not mean that it is impossible to find the colour **c**, but that the chromatic intensity $C_1(1 - 2\varphi)$ is zero.

7.4 Conversions to the complete digital cube

We move from one sector of the RGB cube to another by adding to φ the sector number given by $\lambda(\mathbf{c})$ of equation 15. The hue is therefore approximated by

$$H_1 = (\lambda(\mathbf{c}) + \varphi)k \quad (48)$$

of the same structure as in the HLS system. The coefficient k determines the working units: $k = 60$ for degrees, and $k = 42$ to get resulting values between 0 and 252 which fit into 8-bits.

In parallel to the conversion from φ to H_1 , it is convenient to rewrite the brightness and saturation in terms of max, mid and min functions, which leads to the replacement of system 43 by

$$M_1 = \frac{1}{3}(\max + \text{mid} + \min) \quad (49)$$

$$C_1 = \begin{cases} \frac{3}{2}(\max - M_1) & \text{if } \max + \min \geq 2\text{mid} \\ \frac{3}{2}(M_1 - \min) & \text{if } \max + \min \leq 2\text{mid} \end{cases} \quad (50)$$

$$H_1 = k \left[\lambda(\mathbf{c}) + \frac{1}{2} - \frac{\max + \min - 2\text{mid}}{2C_1} \right] \quad (51)$$

This system of equations is also valid for 8-bit RGB values, and the resulting coordinates are encodable on 8-bits if $k = 42$ is used. It is important to notice that for 256 discrete R , G and B input levels, C_1 takes values which are multiples of $1/2$ (i.e. 512 discrete levels between 0 and 256). Care should therefore be taken when rounding off these values for an 8-bit representation.

For the inverse transformation, the value of H_1 gives, via λ , the order of magnitude of R , G and B and the value of $k\varphi$. Based on whether $k\varphi$ is $\leq k/2$ or $\geq k/2$, we use system 44 or 45 replacing R , G , and B by the components ordered according to the value of λ , and replacing φ by $k\varphi$.

8 The max – min semi-norm and an improved HLS space

We now return to the HLS system and suggest an improvement which overcomes the disadvantages linked to the classic saturation definition. We will refer to this space as the improved HLS or IHLS space. In fact, it is not necessary to modify the HLS space much in order that it be compatible with the three prerequisites. It is sufficient to pass from its cylindrical version to the conic version, which is done by replacing the HLS saturation by the function max – min. A proof that this quantity is a semi-norm is given in appendix B. We now briefly consider the three components of the IHLS system.

8.1 Saturation

The semi-norm $S_0 = \max - \min$ obviously satisfies the first prerequisite on the independence of the projection onto the chromatic plane. Adding to point \mathbf{c} a vector parallel to the achromatic axis simply reduces to adding the same constant to every component of the (R, G, B) vector, which does not modify the value of max – min. In particular, in the RGB space, the vector \mathbf{c} and its projection \mathbf{c}_p on the chromatic plane have the same value for max – min.

On the other hand, this semi-norm is not invariant by projection onto the achromatic axis, in contrast to L_1 and L_2 . As the point \mathbf{c} approaches the achromatic axis, the value of S_0 gets smaller, becoming zero when \mathbf{c} is on the achromatic axis. The projection \mathbf{c}_d of the vector \mathbf{c} onto this axis always has a value of zero for the semi-norm S_0 . It is therefore impossible to build a representation based only on this semi-norm, which is blind to brightness. However, we see also that for all the vectors in the chromatic plane (and only for these vectors), the quantity max – min becomes a norm. This is why we use it for the saturation, complementing it by taking the L_1 or L_2 norm on the achromatic axis.

8.2 Brightness

The L_1 and L_2 norms, and max – min semi-norm which we have studied all guarantee the first prerequisite on independence: two different points \mathbf{c} and \mathbf{c}' having the same projection \mathbf{c}_p onto the chromatic plane have the same saturation and the same hue. As this property remains valid independent of the norm chosen on the achromatic axis, one can easily replace M_2 in equation 23, or M_1 in equation 32 by some weighted mean M (which is still a norm)

$$M(R, G, B) = \alpha R + \beta G + \gamma B \quad \text{where } \alpha + \beta + \gamma = 1 \quad (52)$$

or even by a non-linear estimator, provided that it is a norm. The hue circle remains the same, giving the same weighting to the three fundamental colours, as does the saturation, even if this is not the case for the brightness which replaces M_1 or M_2 .

8.3 Hue

For the hue, we have developed an exact trigonometric expression, given by equations 30 and 31, and a simpler approximate form given by equation 51. Both these hue formulations still show evenly spaced spurious high values in the hue histogram when one converts from an RGB space containing discretely spaced values [16]. However, with the trigonometric hue formulation, it is easy to remove or reduce the height of these histogram spikes by calculating a hue histogram only for pixels having a saturation above a chosen threshold [10]. With the approximation, removal of the spikes is more difficult. Due to the high speed of modern computers, it is highly recommended that the exact trigonometric form be used.

9 Comparison of the saturation and chroma formulations

We compare the distributions of three of the saturation and chroma formulations discussed in this report: the max – min saturation expression (equation 22), the L_2 norm chroma (equation 24), and the L_1 norm chroma (equation 50). These distributions are shown in figure 6.

To calculate the distributions, we start with a $256 \times 256 \times 256$ RGB cube with a point at each set of integer-valued coordinates. For the max – min and L_2 norms, the saturation (and chroma) values of each point are calculated (as floating point values), and then rounded to the nearest integer. Histograms showing the distribution of these integer values (256 levels) are shown in figure 6.

The L_1 chroma measure has an inbuilt quantisation pitfall. When R , G and B are integer-valued, then the L_1 norm M_1 is always a multiple of $\frac{1}{3}$, and therefore max – M_1 and M_1 – min in equation 50 are also multiples of $\frac{1}{3}$. As both of these expressions are multiplied by $\frac{3}{2}$ when calculating the chroma, C_1 is always a multiple of $\frac{1}{2}$. In other words, R , G and B values quantised into 256 levels produce values of C_1 quantised into 512 levels. The rounding of a floating point value of C_1 to the nearest integer therefore behaves extremely erratically, as the $\frac{1}{2}$'s are sometimes rounded up and sometimes down, depending on whether their floating point values are just above or just below 0.5, thereby producing many spurious peaks and valleys in the histogram. For the L_1 chroma distribution in figure 6, 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 the 256 bin histogram shown.

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 also decreases very rapidly as one approaches higher chroma values because it is calculated in the hexagonally shaped sub-region of the chromatic plane. The L_1 norm chroma approximates the L_2 chroma well (if the quantisation effects are taken into account), and the histogram is more regular.

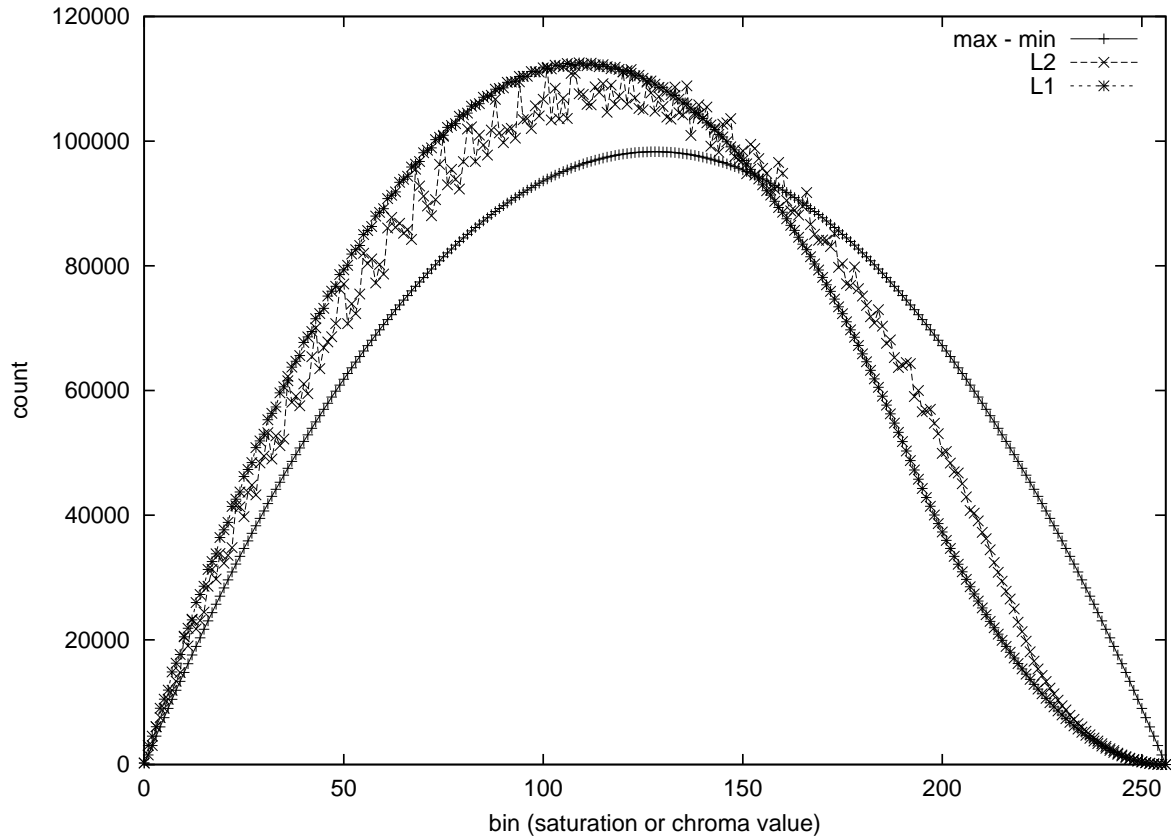


Figure 6: The saturation and chroma histograms.

10 Transformations to and from the IHLS space

Two efficient algorithms to calculate the luminance, trigonometric hue, chroma and saturation from RGB coordinates are given here. We emphasize that identical 3D-polar coordinate values are produced by both algorithms. The inverse transformation, from IHLS to RGB coordinates, is also derived.

We have chosen to use luminance because of its psycho-visual properties. Due to the independence of the brightness and the saturation, one is free to replace the luminance Y with another amplitude measure (brightness, luminance or lightness). The given inverse transformation algorithm, however, only works for amplitude measures which are linear combinations of the R , G and B coordinates. For other brightness definitions, such as those involving $\min(R, G, B)$ or $\max(R, G, B)$ functions, the inverse transformations are more complicated, and will have to be derived.

MATLAB routines implementing the following transformations are available at <http://www.prip.tuwien.ac.at/~hanbury>.

10.1 RGB to IHLS

Two algorithms for producing exactly the same IHLS coordinates are presented, the second algorithm being easier to invert than the first.

10.1.1 The simplest implementation

For the simplest implementation, one calculates an amplitude measure (equation 1, 2, or 5), the saturation using equation 22, and the hue using equations 30 and 31. We have therefore used an L_1 norm for the luminance, the max – min norm for the saturation, and the scalar product of the L_2 norm for the hue.

$$Y(c) = 0.2126R + 0.7152G + 0.0722B \quad (53)$$

$$S(c) = \max(R, G, B) - \min(R, G, B) \quad (54)$$

$$H'(c) = \arccos \left[\frac{R - \frac{1}{2}G - \frac{1}{2}B}{(R^2 + G^2 + B^2 - RG - RB - BG)^{\frac{1}{2}}} \right] \quad (55)$$

$$H(c) = \begin{cases} 360^\circ - H' & \text{if } B > G \\ H' & \text{otherwise} \end{cases} \quad (56)$$

10.1.2 An alternative

An alternative way of arriving at exactly the same IHLS values is now presented. It is based on the algorithm suggested by Carron [2]. The changes with respect to Carron's version are the extension to calculate the saturation from the chroma, and the use of luminance instead of brightness. It is also similar to the IHS system described by Pratt [23], except for a change in the hue origin and a different saturation definition. It still uses the L_1 norm luminance expression, but an L_2 norm for the saturation. This algorithm allows a more straightforward inverse transformation to be derived as it contains no max or min functions.

The first step is to calculate the luminance Y and two chrominance coordinates

$$\begin{bmatrix} Y \\ C_1 \\ C_2 \end{bmatrix} = \begin{bmatrix} 0.2125 & 0.7154 & 0.0721 \\ 1 & -\frac{1}{2} & -\frac{1}{2} \\ 0 & -\frac{\sqrt{3}}{2} & \frac{\sqrt{3}}{2} \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \quad (57)$$

followed by the calculation of the chroma $C \in [0, 1]$ (this chroma is equal to the chroma C_2 derived in the framework of the L_2 norm in section 6)

$$C = \sqrt{C_1^2 + C_2^2} \quad (58)$$

and the hue $H \in [0^\circ, 360^\circ]$

$$H = \begin{cases} \text{undefined} & \text{if } C = 0 \\ \arccos\left(\frac{C_1}{C}\right) & \text{if } C \neq 0 \text{ and } C_2 \leq 0 \\ 360^\circ - \arccos\left(\frac{C_1}{C}\right) & \text{if } C \neq 0 \text{ and } C_2 > 0 \end{cases} \quad (59)$$

We derive from equations 24, 26 and 27 the value of the saturation $S \in [0, 1]$

$$S = \frac{2C \sin(120^\circ - H^*)}{\sqrt{3}} \quad (60)$$

in which

$$H^* = H - k \times 60^\circ \text{ where } k \in \{0, 1, 2, 3, 4, 5\} \text{ so that } 0^\circ \leq H^* \leq 60^\circ \quad (61)$$

10.2 Inverse transformation from IHLS to RGB

To transform colours represented in the IHLS coordinate system obtained using either of the algorithms of section 10.1 to RGB coordinates, one first calculates the chroma values from the saturation values (using equation 60)

$$C = \frac{\sqrt{3}S}{2 \sin(120^\circ - H^*)} \quad (62)$$

where H^* is given by equation 61. From the chroma, one calculates

$$C_1 = C \cos(H) \quad (63)$$

$$C_2 = -C \sin(H) \quad (64)$$

For the case where the hue is undefined: $C_1 = C_2 = 0$. Finally, the inverse of the matrix used in equation 57 is used to obtain R , G and B

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 1.0000 & 0.7875 & 0.3714 \\ 1.0000 & -0.2125 & -0.2059 \\ 1.0000 & -0.2125 & 0.9488 \end{bmatrix} \begin{bmatrix} Y \\ C_1 \\ C_2 \end{bmatrix} \quad (65)$$

11 Application examples

Three applications in which the use of the suggested IHLS model is advantageous are described. The first is the calculation of hue statistics, for which better results can be obtained by utilising a weighting by saturation values. This saturation weighting is then applied in the calculation of hue histograms. We finally show an example in which the saturation plays a dominant role in a mathematical morphology operator.

11.1 Colour statistics

In a 3D-polar coordinate colour space, standard statistical formulae can be used to calculate statistical descriptors for the brightness and saturation coordinates. The hue, as has been pointed out, is an angular value, so circular statistical descriptors [8] should be calculated for it.

11.1.1 Hue statistics

We initially summarise some of the standard circular statistics formulae. Given n hue values $H_i, i = 1, \dots, n$, the *mean direction* \overline{H} is the direction of the resultant vector of the sum of n unit vectors having directions H_i . This direction is given by

$$\overline{H} = \arctan\left(\frac{B}{A}\right) \quad (66)$$

where

$$A = \sum_i \cos H_i, \quad B = \sum_i \sin H_i \quad (67)$$

and the necessary care to taken to expand the output of the arctan function to the range $[0, 360^\circ]$. The *mean length* of the resultant vector is

$$\overline{R} = \frac{\sqrt{A^2 + B^2}}{n} \quad (68)$$

The value of the mean length is in the range $[0, 1]$ and can be used as an indicator of the dispersion of the data (similar to the variance). If $\overline{R} = 1$, all the H_i are coincident. Conversely, a value of 0 does not necessarily indicate a homogeneous data distribution, as certain non-homogeneous distributions can also result in this value.

11.1.2 Saturation-weighted hue statistics

The calculation of statistics based only on the hue, described above, has the disadvantage of ignoring the close relationship between the chrominance coordinates (hue and saturation). For weakly saturated colours (greylevels), the hue value is unimportant. Indeed, for zero-saturated colours, the hue value is meaningless. We can take these different levels of importance into account in the statistics by weighting the hues by their corresponding saturations.

Given n pairs of values, the hue H_i and its associated saturation S_i , we proceed as before, except that instead of finding the resultant of unit vectors, the vector with direction H_i has length S_i . The hues associated with small saturation values will therefore have less influence on the direction of the resultant vector. This weighting is simply done by replacing equation 67 by

$$A_S = \sum_{i=1}^n S_i \cos H_i, \quad B_S = \sum_{i=1}^n S_i \sin H_i \quad (69)$$

and replacing A and B in equation 66 by A_S and B_S . We denote by \overline{H}_S the resultant saturation-weighted hue mean. The mean length (equation 68) becomes

$$\overline{R}_S = \frac{\sqrt{A_S^2 + B_S^2}}{\sum_{i=1}^n S_i}$$

In practice, for images which contain only strongly saturated colours, there is not a significant difference between the values of weighted and unweighted hue means. Figure 11b shows an image in which this difference is important. As is visible in figure 7a, the saturation of the two brown cells is higher than the saturation of the surroundings. For this image, the unweighted hue mean is $\overline{H} = 326.9^\circ$, and the saturation-weighted hue mean is $\overline{H}_S = 19.7^\circ$. To show the difference, thresholds on the hue band image were calculated for the intervals $[\overline{H} - 20^\circ, \overline{H} + 20^\circ]$ and $[\overline{H}_S - 20^\circ, \overline{H}_S + 20^\circ]$, and these are shown in figures 7b and 7c respectively. On examining these images, it is clear the the saturation-weighted hue mean corresponds to the hue of the most highly saturated regions, the two cells, whereas the unweighted hue mean is skewed by the hues associated with the surrounding low-saturation regions.

11.2 Hue histograms

Hue histograms are often used as an image feature for retrieval of colour images from databases. In these histograms, one generally wishes to exclude achromatic and near-achromatic pixels, for

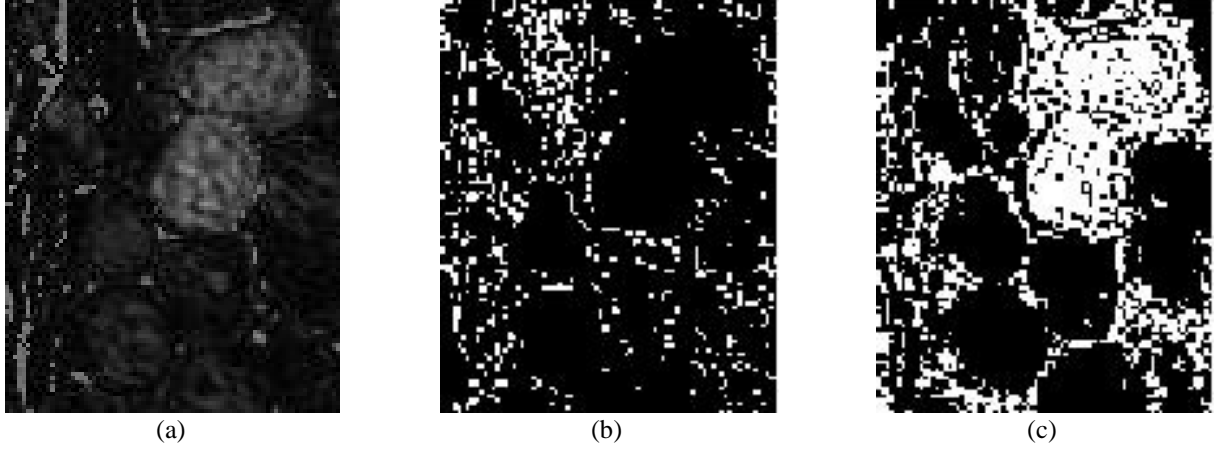


Figure 7: (a) Saturation of figure 11b. (b) Pixels of figure 11b with hue values in the interval 20° on each side of the non-weighted hue mean $[\overline{H} - 20^\circ, \overline{H} + 20^\circ]$. (c) Pixels of figure 11b with hue values in the interval $[\overline{H}_S - 20^\circ, \overline{H}_S + 20^\circ]$ around the saturation-weighted hue mean.

which the hue has little meaning. As the traditional (cylindrical) saturation is essentially useless in discriminating between achromatic and chromatic colours in the low-brightness part of the HSV space, a number of heuristics, summarised by Stokman and Gevers [29], have been used. Tico et al. [30], for example, suggest using the standard deviation of the R , G and B coordinates in conjunction with a fuzzy membership function containing two user-specified parameters, to calculate a weight differentiating between chromatic and achromatic colours, the basic idea being that the more colourful (higher saturated) pixels receive higher weighting in the hue histogram than the less colourful (lower saturated) ones. The saturation measurement suggested in this report can be directly used as such a weight. In building the saturation-weighted hue histogram for the specific case of the hues having been rounded to the nearest integer, the total in bin θ ($\theta \in [0^\circ, 1^\circ, \dots, 360^\circ]$) of the histogram is simply calculated as

$$W_\theta = \sum_x S_x \delta_{\theta H_x}$$

where the sum is over all the pixel positions x in the image, H_x and S_x are respectively the hue and saturation at point x , and δ_{ij} is the Kronecker delta function. A saturation-weighted brightness histogram can be calculated analogously by using the inverse weighting $(1 - S_x)$, thereby privileging the low saturation (achromatic) pixels. An alternative saturation weighting for the hue, which shifts the hue values around the circle, is described in [11], where it is used in the context of colour ordering for mathematical morphology.

11.3 Mathematical morphology and the lattice approach

The application of mathematical morphology to colour images is difficult due to the vectorial nature of the colour data. The most commonly adopted approach is to use one of the vector orders suggested by Barnett [1]. The *lexicographical order* is convenient as it imposes a total order on the vectors, thereby ensuring that there are no pairs of vectors for which the order is

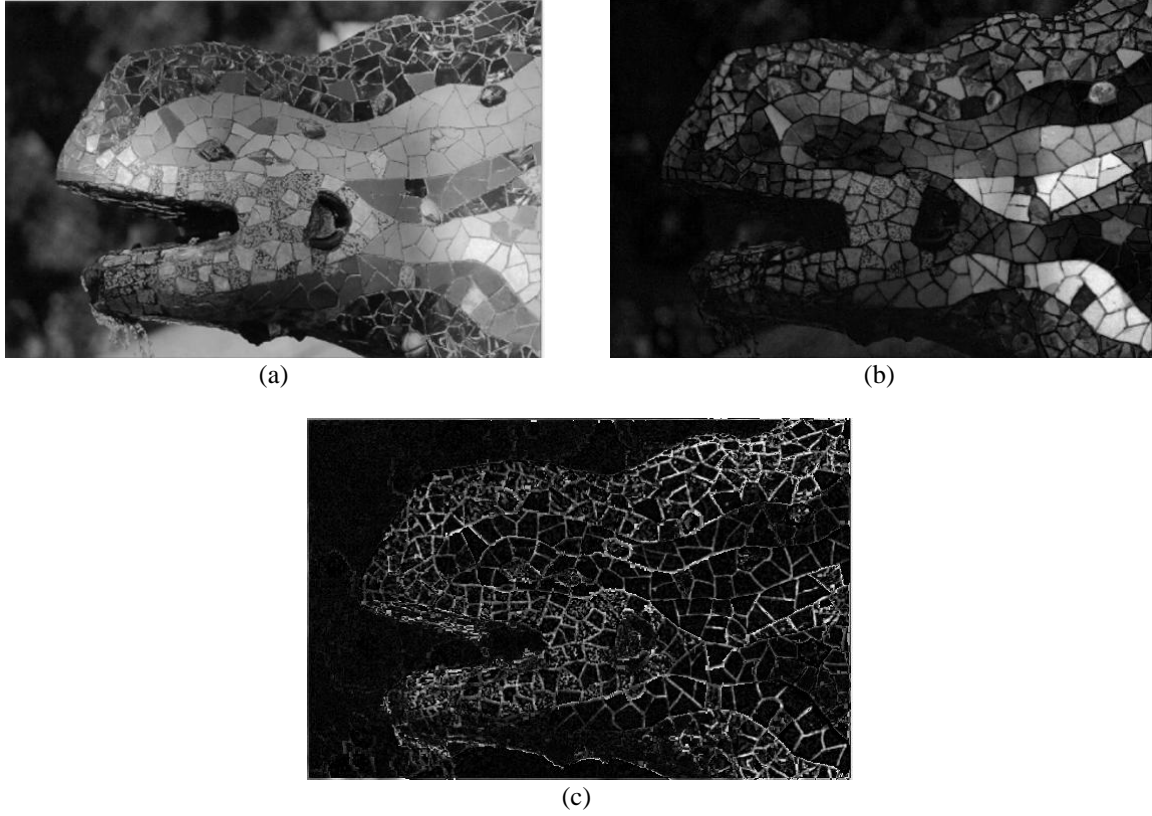


Figure 8: (a) Luminance of figure 12a. (b) Saturation of figure 12a. (c) The top-hat — the Euclidean distance between the corresponding pixels in figures 12a and b.

uncertain, contrary to the *marginal* and *reduced* orders [4]. The use of a lexicographical order directly in the RGB space requires that one of the colours red, green or blue be arbitrarily elevated to a dominant role. The more homogeneous 3D-polar coordinate representation can avoid this disadvantage. It is often suggested that using a lexicographical order with a brightness measure at the top level gives the “best” results [18, 19], although certain problems (one of which is shown in this section) can be solved by the use of a lexicographical order with saturation at the top level.

Consider the colour image in figure 12a, in which we have given ourselves the task of extracting the grey lines between the mosaic tiles. The luminance image in figure 8a shows that some of the tiles have a lower luminance than the grey lines, whereas others have a higher luminance. In the saturation image of figure 8b, one sees that the saturation of the grey lines is almost always lower than that of the more colourful tiles. We therefore choose the following lexicographical order between any two vectors in 3D-polar coordinates $\mathbf{v}_i = (Y_i, S_i, H_i)$

$$\mathbf{v}_i > \mathbf{v}_j \quad \text{if} \quad \begin{cases} S_i > S_j \\ \text{or } S_i = S_j \text{ and } Y_i > Y_j \\ \text{or } S_i = S_j \text{ and } Y_i = Y_j \text{ and } H_i \div H_0 < H_j \div H_0 \end{cases} \quad (70)$$

where the symbol \div on the last level represents the acute angle between two angular values, i.e.

$$\theta_i \div \theta_j = \begin{cases} |\theta_i - \theta_j| & \text{if } |\theta_i - \theta_j| \leq 180^\circ \\ 360^\circ - |\theta_i - \theta_j| & \text{if } |\theta_i - \theta_j| \geq 180^\circ \end{cases} \quad (71)$$

The variable H_0 on the third level of the lexicographical order is the position of the origin of the hue circle. As the third level of a lexicographical order is hardly ever used in practice [12], its value does not have much effect on the results. Applying a morphological closing using this order with a square structuring element of size 5×5 pixels to figure 12a produces figure 12b, in which one sees that the tiles have been expanded to cover the grey lines. Finally, a form of top-hat [28] is calculated by determining the Euclidean distance (in 3D-polar coordinates) between the corresponding pixels of figures 12a and 12b to give the greyscale image in figure 8c, in which the pixels with highest greylevels correspond to the features we wish to extract.

The use of the Euclidean distance in this space does not imply that it is in any way perceptually uniform. Nevertheless, the magnitudes of the Euclidean distances give a good indication of the magnitudes of the colour differences. In many imaging applications, especially in multimedia, one does not have access to the calibration data necessary to convert to a perceptually uniform space, such as the CIE $L^*a^*b^*$ space. This colour difference approximation could therefore be useful in these cases.

The lexicographical order with saturation at the first level is useful in situations where one wishes to distinguish between colourful and non-colourful objects or regions. This could be, for example, in the isolation of a non-saturated phase amongst several others in geological or biological applications, or the extraction of colourful blobs from a grey background.

12 Conclusion

The 3D-polar (hue, saturation and brightness) coordinate colour representation systems currently in use are often unsuited to image processing and analysis. The principal reason is the artificial expansion of the natural (conic or bi-conic) shapes of these spaces into cylindrical form, by dividing each saturation value by the maximum saturation possible for the associated brightness. While this cylindrical shape is convenient for colour specification or choice applications, it is completely unsuitable for image processing and analysis for the following reasons:

- Colours which appear almost achromatic can receive high (or even maximal) saturation values.
- Because the saturation normalisation depends on the brightness function, these two coordinates are not independent.
- Comparison between saturation values is meaningless, as each saturation is normalised by a different factor.

We present three prerequisites for 3D-polar coordinate colour systems to be suitable for image processing and analysis, and then derive systems using the L_1 , L_2 and max – min norms. Transformation systems to and from the L_1 norm space and the improved HLS (IHLS) space are presented. The L_1 space transformations are linear and include a linear approximation of

the hue. This space should therefore be used if very efficient transformations are required. The I HLS space is more suited to image analysis tasks, as it permits a wide choice of brightness functions and does not approximate the hue. Any 3D-polar coordinate system is very closely tied to the RGB space, being simply a different representation of it. It therefore does not have any of the good properties of the $L^*a^*b^*$ or $L^*u^*v^*$ spaces, such as perceptual uniformity. Its advantages are that the conversion algorithm is very simple, and that no colour calibration data, such as the white point coordinates, are required. This calibration data is usually not available in multimedia applications, for example.

Three applications demonstrating the good properties of the suggested saturation measure are given. The calculation of statistics of the hue weighted by the saturation and the determination of the saturation-weighted hue histogram use the fact that the suggested saturation measure is always small for near-achromatic colours. The mathematical morphology application takes advantage of the ability to do comparisons between the saturation values.

In summary, much confusion and incompatibility between results in colour image processing and analysis could be avoided by the use of the suggested 3D-polar coordinate system.

A Proof of the independence of the proposed saturation and the brightness function

We give a proof of the independence of the brightness and saturation, stated in section 5. Consider the triangle shown in figure 5a, which contains all the points with the same hue as c . This triangle is reproduced in figure 9 to facilitate the following proof.

Proposition A.1. *The proposed saturation definition*

$$S = \frac{\|L(c) - c\|}{\|L[q(c)] - q(c)\|} \quad (72)$$

is independent of the choice of the brightness function.

Proof. Refer to figure 9. The iso-brightness lines between $L[q(c)]$ and $q(c)$ and between $L(c)$ and c are by definition parallel. By adding the line between f and c which is parallel to the line between w and $q(c)$, we create two similar triangles with vertices $f, c, L(c)$ and $w, q(c), L[q(c)]$ respectively. We hence have the relation

$$\frac{\|L(c) - c\|}{\|L[q(c)] - q(c)\|} = \frac{\|f - c\|}{\|w - q(c)\|} \quad (73)$$

This relation is true for any brightness function, as long as the iso-brightness surfaces it produces are parallel. Hence, the proposed saturation definition is independent of the brightness function. \square

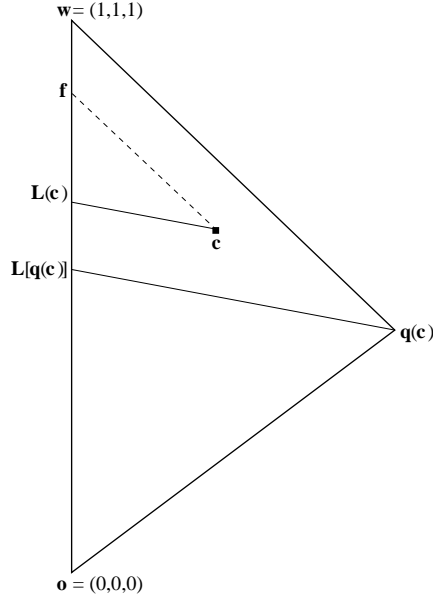


Figure 9: The triangle of figure 5a which contains all the points having the same hue as c . The achromatic axis is at the left, between the points o and w .

B Proof that $\max - \min$ is a semi-norm

Proposition B.1. *In a vectorial space E of finite dimension n and in which the vector \mathbf{x} has coordinates $\{x_i, 1 = 1, 2, \dots, n\}$, the quantity*

$$\delta(\mathbf{x}) = \max(\mathbf{x}) - \min(\mathbf{x}) \quad (74)$$

$$= \max\{x_i, i = 1, \dots, n\} - \min\{x_i, i = 1, \dots, n\} \quad (75)$$

is a semi-norm.

Proof. Firstly we observe that for any set of values $x_i \in]-\infty, +\infty[, 1 \leq i \leq n$, the quantity δ cannot be negative. For every $\mathbf{x} \in E$, we obviously have $\delta(\lambda\mathbf{x}) = \lambda\delta(\mathbf{x})$. It remains to show the validity of the triangular inequality associated with each pair of points $(\mathbf{x}, \mathbf{x}')$ in E by the relation

$$\delta(\mathbf{x} + \mathbf{x}') \leq \delta(\mathbf{x}) + \delta(\mathbf{x}') \quad (76)$$

We first show that

$$\max(\mathbf{x} + \mathbf{x}') \leq \max(\mathbf{x}) + \max(\mathbf{x}') \quad (77)$$

Take for the n th coordinate the one which maximises $\mathbf{x} + \mathbf{x}'$, i.e. $\max(\mathbf{x} + \mathbf{x}') = x_n + x'_n \geq x_i + x'_i, i = 1, \dots, n$. It is obviously not possible to have $x_n < x_i$ and $x'_n < x'_i$, with $i = 1, \dots, n$. Suppose therefore, for example that $\max(\mathbf{x}) = x_n$ and $\max(\mathbf{x}') = x'_j$ for an index $j \in 1, \dots, n$. It follows that

$$\max(\mathbf{x} + \mathbf{x}') = x_n + x'_n \leq x_n + x'_j = \max(\mathbf{x}) + \max(\mathbf{x}'). \quad (78)$$

One can prove in exactly the same way that $\min(\mathbf{x} + \mathbf{x}') \geq \min(\mathbf{x}) + \min(\mathbf{x}')$, from which it is possible to establish inequality 76 by subtraction, which completes the proof. \square

Acknowledgments

Thank you to Jesús Angulo and Walter Kropatsch for their extremely useful comments on the manuscript.

References

- [1] V. Barnett. The ordering of multivariate data. *Journal of the Statistical Society of America A*, 139(3):318–354, 1976.
- [2] Thierry Carron. *Segmentations d’images couleur dans la base Teinte-Luminance-Saturation: approche numérique et symbolique*. PhD thesis, Université de Savoie, 1995.
- [3] Vicent Caselles, Bartomeu Coll, and Jean-Michel Morel. Geometry and color in natural images. *Journal of Mathematical Imaging and Vision*, 16:89–105, 2002.
- [4] Mary L. Comer and Edward J. Delp. Morphological operations for colour image processing. *Journal of Electronic Imaging*, 8(3):279–289, 1999.
- [5] Commission Internationale de l’Éclairage. *International Lighting Vocabulary*. Number 17.4. CIE, 4th edition, 1987.
- [6] Claire-Hélène 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.
- [7] Claire-Hélène Demarty and Serge Beucher. Color segmentation algorithm using an HLS transformation. In *Proceedings of the International Symposium on Mathematical Morphology (ISMM ’98)*, pages 231–238, 1998.
- [8] N. I. Fisher. *Statistical Analysis of Circular Data*. Cambridge University Press, 1993.
- [9] T. Gevers and A. W. M. Smeulders. A comparative study of several color models for color image invariant retrieval. In *Proceedings of the First International Workshop, IDB-MMS*, pages 17–26, August 1996.
- [10] Allan Hanbury. The taming of the hue, saturation and brightness colour space. In *Proceedings of the Seventh Computer Vision Winter Workshop, Bad Aussee, Austria*, 2002.
- [11] Allan Hanbury and Jean Serra. Mathematical morphology in the HLS colour space. In Tim Cootes and Chris Taylor, editors, *Proceedings of the British Machine Vision Conference 2001*, pages 451–460. BMVA, 2001.
- [12] Allan Hanbury and Jean Serra. Mathematical morphology in the L*a*b* colour space. Technical Report N-36/01/MM, CMM, Ecole des Mines de Paris, 2001.
- [13] Allan Hanbury and Jean Serra. Morphological operators on the unit circle. *IEEE Transactions on Image Processing*, 10(12):1842–1850, December 2001.

- [14] ITU-R Recommendation BT.709. Basic parameter values for the HDTV standard for the studio and for international programme exchange. Geneva: ITU, 1990.
- [15] James M. Kasson and Wil Plouffe. An analysis of selected computer interchange color spaces. *ACM Transactions on Graphics*, 11(4):373–405, October 1992.
- [16] John R. Kender. Saturation, hue and normalised color: Calculation, digitization effects, and use. Technical report, Department of Computer Science, Carnegie-Mellon University, 1976.
- [17] Haim Levkowitz and Gabor T. Herman. GLHS: A generalised lightness, hue and saturation color model. *CVGIP: Graphical Models and Image Processing*, 55(4):271–285, July 1993.
- [18] G. Louverdis, M. I. Vardavoulia, I. Andreadis, and Ph. Tsalides. A new approach to morphological color image processing. *Pattern Recognition*, 35:1733–1741, 2002.
- [19] Francisco Ortiz, Fernando Torres, Jesús Angulo, and Santiago Puente. Comparative study of vectorial morphological operations in different color spaces. In *Proc. SPIE on Intelligent Robots and Computer Vision (IRCV) XX : Algorithms, Techniques and Active Vision*, volume 4572, pages 259–268, 2001.
- [20] R. A. Peters II. Mathematical morphology for angle-valued images. In *Non-Linear Image Processing VIII*, volume 3026. SPIE, February 1997.
- [21] Charles Poynton. Frequently asked questions about color. URL: <http://www.inforamp.net/~poynton/PDFs/ColorFAQ.pdf>, 1999.
- [22] Charles Poynton. Frequently asked questions about gamma. URL: <http://www.inforamp.net/~poynton/PDFs/GammaFAQ.pdf>, 1999.
- [23] William K. Pratt. *Digital Image Processing*. Wiley, 1991.
- [24] Gaurav Sharma and H. Joel Trussel. Digital color imaging. *IEEE Transactions on Image Processing*, 6(7):901–932, July 1997.
- [25] Tian-Yuan Shih. The reversibility of six geometric color spaces. *Photogrammetric Engineering and Remote Sensing*, 61(10):1223–1232, October 1995.
- [26] Alvy Ray Smith. Color gamut transform pairs. *Computer Graphics*, 12(3):12–19, 1978.
- [27] J. R. Smith. *Integrated Spatial and Feature Image Systems: Retrieval, Compression and Analysis*. PhD thesis, Columbia University, 1997.
- [28] Pierre Soille. *Morphological Image Analysis: Principles and Applications*. Springer-Verlag, 1999.
- [29] Harro Stokman and Theo Gevers. Error propagation due to color space transforms. In *Proceedings of SPIE, San Jose*, January 2000.

- [30] Marius Tico, Taneli Haverinen, and Pauli Kuosmanen. A method of color histogram creation for image retrieval. In *Proceedings of NORSIG*, 2000.

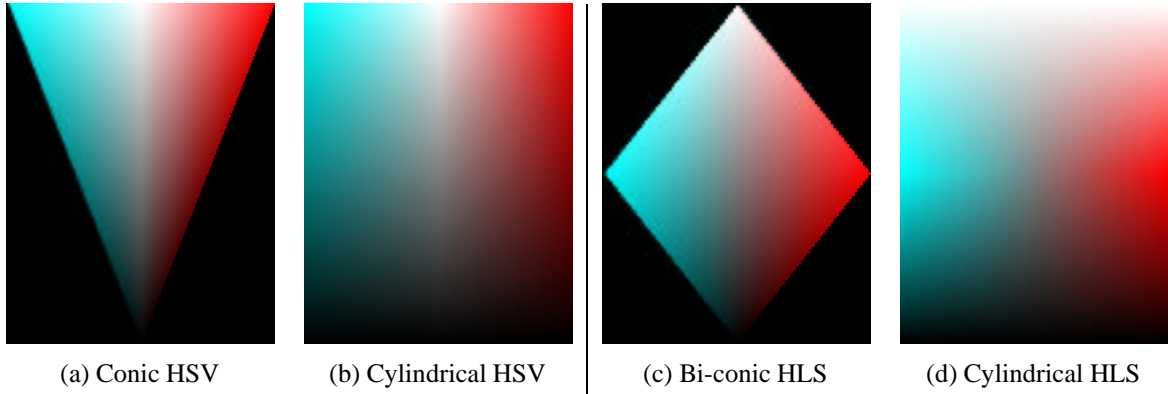


Figure 10: Slices through the conic and cylindrical versions of the HSV and HLS colour spaces. The brightness increases from bottom to top, and the saturation increases from the centre (achromatic axis) outwards. Colours to the right of the central achromatic axis have hues of 0° , and colours to the left have hues of 180° .

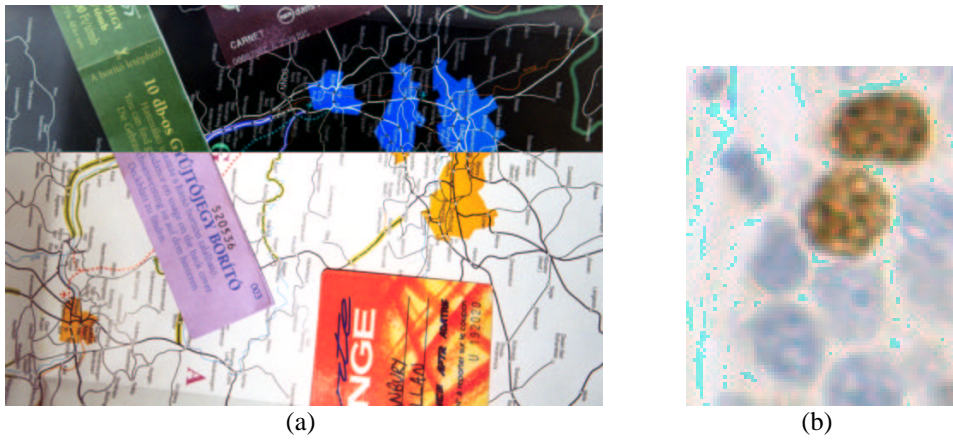


Figure 11: Example images.



Figure 12: (a) Colour image (size 544×360 pixels). (b) A morphological closing of image (a) using a lexicographical order with saturation at the first level.