

Constructing Cylindrical Coordinate Colour Spaces

Allan Hanbury

*Pattern Recognition and Image Processing Group (PRIP),
Institute of Computer-Aided Automation, Vienna University of Technology,
Favoritenstraße 9/1832, A-1040 Vienna, Austria*

Abstract

A cylindrical coordinate colour space (lightness, saturation/chroma, hue) is derived from an opponent colour space in the RGB space. It is shown how cylindrical coordinate colour models widely used in the literature are related to or can be reduced to the derived model, thereby contributing to creating a unified cylindrical coordinate colour model. In particular, the widely used saturation expression $\max(R, G, B) - \min(R, G, B)$ is derived from the proposed model. Properties of the derived chroma and saturation expressions are examined. Finally, some applications of cylindrical coordinate colour spaces are briefly reviewed.

Key words: hue, saturation, chroma, lightness, colour, cylindrical coordinates

1 Introduction

The representation of the colour coordinates of images in cylindrical coordinates (hue, saturation/chroma, lightness), is widely used in the image analysis and computer vision community. There are however two main difficulties associated with its use:

- (1) the large choice of available transformations from an RGB space, e.g. HSV (Smith, 1978), HSL, HMMD (Manjunath et al., 2002), HSB and HSI (Gonzalez and Woods, 1992).
- (2) the fact that some of these transformations have their saturation normalised by the lightness, making them unsuitable for some tasks.

The latter concern can be understood by looking at the saturation images shown for three colour models in Figure 1. The saturation for the HSV and HSL models represent percentages of the maximum saturation obtainable for

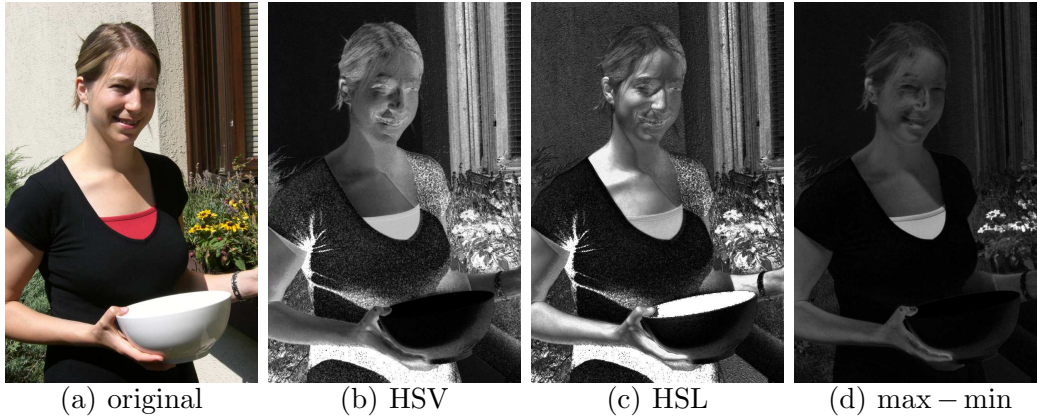


Fig. 1. (a) Original image. (b) HSV, (c) HSL and (d) max – min saturation.

a given lightness. This implies that in the lightness ranges where the saturation range is small, there is a large variation in the percentages. This can be seen by the “noisy” appearance of the dark image regions (e.g. shadows) for the HSV colour space and the dark and light image regions (e.g. reflections on the arm) for the HSL colour space. Furthermore, because the majority of RGB coordinates of the inside of the white bowl are (255, 255, 254), the saturation percentage for the bowl is high for the HSL saturation and low for the HSV saturation. Removing the saturation normalisation by lightness results in the saturation image shown in Figure 1d, which is more useful for many applications.

In this paper, we derive a cylindrical coordinate colour space from basic principles. We also show how existing colour spaces are related to or can be reduced to the derived colour space, thereby contributing to a unification of these colour spaces.

The derivation follows that of the cylindrical coordinate version of the CIELAB colour space (Fairchild, 1997). For the CIELAB colour space, an opponent colour space (OCS) is created in the XYZ colour space. The axes represent lightness (L^*), and the red-green (a^*) and blue-yellow (b^*) opponent colours modelled on the human visual system. A non-linear transform and weighting on these coordinates is used to impose the approximate perceptual uniformity property on the CIELAB space. A chromaticity coordinate (a^*, b^*) is represented in polar coordinates by the chroma C_{ab}^* (the Euclidean distance from the lightness axis) and the hue angle h_{ab} expressed in degrees starting from the positive a^* axis (red) and turning in an anti-clockwise direction.

In order to be able to convert from the RGB space to the CIELAB space, one needs the colour coordinates of the primary colours of the image sensor (for the RGB to XYZ conversion), as well as the coordinates of the scene illuminant (for XYZ to CIELAB). This information is often not available for images that are analysed, for example arbitrary images obtained from the internet. While it

is possible to estimate these values, one adds extra uncertainty to the analysis process by doing this. This shows that building a cylindrical coordinate colour system directly in the RGB space is often desirable, as one has access to polar chromaticity coordinates through a simpler transform requiring no estimated values.

We follow a similar process to the CIELAB transformation: (1) Creation of an OCS in the RGB space, (2) Conversion to cylindrical coordinates. We do not attempt to impose any type of perceptual uniformity on the space. While this leads to a completely usable colour space, it does not account for the saturation expression $\max(R, G, B) - \min(R, G, B)$ that is often used. We therefore additionally demonstrate how this saturation expression can be derived from this space.

The main contributions of the paper are:

- (1) a sound derivation of a cylindrical coordinate colour space within the RGB space.
- (2) the unification of several existing cylindrical coordinate colour spaces by demonstrating how they reduce to or are related to the derived space.

The paper is organised as follows. Section 2 presents the derivation of the cylindrical coordinate colour space from an opponent colour space. Section 3 concentrates on the chroma and saturation terms. A summary of the transform and inverse transform between the RGB and cylindrical coordinate spaces is given in Section 4. A short review of some applications is given in Section 5. Section 6 concludes.

2 Derivation based on an Opponent Colour Space

One performs this transformation by first building an OCS in the RGB space and then representing it in cylindrical coordinates. We do not specify if the RGB coordinates used in this paper are gamma corrected or not. More on the effect of gamma correction can be found in Angulo and Serra (2007) and Serra (2005).

2.1 Opponent Colour Space

We take an RGB space with coordinates $R \in [0, 1]$, $G \in [0, 1]$, $B \in [0, 1]$. To convert to an OCS, a new axis is placed in the RGB space. This axis is usually between the black $(0, 0, 0)$ and white $(1, 1, 1)$ corners of the RGB cube, and hence contains all colours for which $R = G = B$. The coordinate on this axis

gives a measure of the *lightness* I , and we therefore refer to it as the *lightness axis*. A plane perpendicular to this axis and with origin at the intersection with the lightness axis is chosen and all RGB coordinates are projected onto this plane. The position on the plane gives information about the chromaticity of a pixel, and it is therefore referred to as the *chromatic plane*.

A simple OCS is given by Plataniotis and Venetsanopoulos (2000):

$$I = R + G + B \tag{1}$$

$$c_1 = 2B - R - G \tag{2}$$

$$c_2 = R - G \tag{3}$$

This OCS has the disadvantage that the red, green and blue corners of the RGB cube have different Euclidean distances from the origin of the chromatic plane than the cyan, magenta and yellow corners. This leads to a non hexagonally-shaped area of valid RGB coordinates in the chromatic plane.

The opponent colour space used by Lambert and Carron (1999) is normalised so that the six abovementioned corners of the RGB cube projected onto the chromatic plane all have unit Euclidean distance from its origin. The projection of the RGB cube onto the chromatic plane therefore has the shape of a hexagon. This transformation is defined as¹:

$$I = \frac{1}{3}(R + G + B) \tag{4}$$

$$c_1 = R - \frac{1}{2}G - \frac{1}{2}B \tag{5}$$

$$c_2 = \frac{\sqrt{3}}{2}(G - B) \tag{6}$$

The area of the chromatic plane containing c_1 and c_2 coordinates corresponding to colours in the RGB unit cube is shown by the hexagon in Figure 2a. The circle circumscribing the hexagon has a radius of 1. We adopt the notation in which RGB vectors projected onto the chromatic plane take a subscript p . Hence, for example, $\mathbf{r}_p = (1, 0)$ and $\mathbf{g}_p = (-1/2, \sqrt{3}/2)$ represent respectively the coordinates of the projections onto the chromatic plane of pure red $\mathbf{r} = (1, 0, 0)$ and pure green $\mathbf{g} = (0, 1, 0)$.

The RGB unit cube is shown in Figure 3. It is divided into six sectors which

¹ In Lambert and Carron (1999), the hue angle turns in an unconventional way, with green at $4\pi/3$ and blue at $2\pi/3$. We have modified the transformation by exchanging the signs in Equation 6 so that the hue angle turns anti-clockwise, as for the polar CIELAB transform.




Fig. 2. (a) Chromatic plane. (b) 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° .

are projected onto the six sectors of the hexagon on the chromatic plane. The following equation gives the sector of a colour with coordinates $\mathbf{c} = (R, G, B)$ based on the order of magnitude of its coordinates

$$\lambda(\mathbf{c}) = \begin{cases} 0 & \text{if } R > G \geq B \\ 1 & \text{if } G \geq R > B \\ 2 & \text{if } G > B \geq R \\ 3 & \text{if } B \geq G > R \\ 4 & \text{if } B > R \geq G \\ 5 & \text{if } R \geq B > G \end{cases} \quad (7)$$

The cube edge corresponding to each sector and its corresponding hexagon edge on the chromatic plane are indicated by the italic numbers in Figures 3 and 2a.

In sector 1, the line between \mathbf{r} and \mathbf{y} in the RGB cube is projected onto the hexagon edge between \mathbf{r}_p and \mathbf{y}_p on the chromatic plane. It is therefore clear that this edge corresponds to a line of RGB coordinates having R and B constant and only G varying. The same holds for all lines in the chromatic plane parallel to the $\mathbf{r}_p\mathbf{y}_p$ edge and lying within the first hexagon sector. This can be generalised to all sectors of the hexagon as follows: for a given sector of the hexagon, lines parallel to the hexagon edge of the sector and included in the sector correspond to lines in the RGB space for which the coordinates




Fig. 3. The RGB unit cube. The italic numbers indicate the edges corresponding to the six sectors into which the RGB cube is divided.

with greatest and least magnitude for the sector (see Equation 7) are constant and only the coordinate of intermediate magnitude varies.

An opponent colour space is also used by van de Weijer et al. (2005), with the refinement that the lightness axis is between $(0, 0, 0)$ and the coordinate of the colour of the illuminant in RGB space $\mathbf{c}^i = (\alpha, \beta, \gamma)^T$. Using $\mathbf{c}^i = (1, 1, 1)$ produces an OCS similar to Equations 4 to 6, but with rotated chromaticity axes and different normalisation.

2.2 Polar representation

A 3D-polar representation of an OCS is obtained by keeping the lightness I and expressing c_1 and c_2 in polar coordinates:

$$h = \arctan\left(\frac{c_2}{c_1}\right) \quad (8)$$

$$c = \sqrt{c_1^2 + c_2^2} \quad (9)$$

When substituting the OCS in Equations 5 to 6 into the above equations, pure red has hue $h = 0^\circ$ (similar to the CIELAB space). Furthermore, the chroma $c = 1$ for the six corners of the hexagon on the chromatic plane. Substituting Equations 4–6 into Equation 9, one obtains

$$c = \sqrt{R^2 + G^2 + B^2 - RG - RB - GB} \quad (10)$$

and into Equation 8, one obtains

$$h = \arctan \left(\frac{\sqrt{3}(G - B)}{2R - G - B} \right) \quad (11)$$

Using the alternative definition from Lambert and Carron (1999)

$$h = \begin{cases} \arccos \left(\frac{c_1}{c} \right) & \text{if } c_2 \geq 0 \\ 360^\circ - \arccos \left(\frac{c_1}{c} \right) & \text{if } c_2 < 0 \end{cases} \quad (12)$$

$$\text{where } \frac{c_1}{c} = \left(\frac{R - \frac{1}{2}G - \frac{1}{2}B}{\sqrt{R^2 + G^2 + B^2 - RG - RB - GB}} \right) \quad (13)$$

one obtains the hue expression derived in a more geometrical way for the HSI model presented in Gonzalez and Woods (1992). Pratt (1991) presents a similar transformation with the hue origin at blue and a different normalisation that he calls the *IHS colour coordinate system*. With the OCS defined in van de Weijer et al. (2005), $h = 60^\circ$ for pure red.

As it is clear that the lightness I is independent of the chromaticity coordinates c_1 and c_2 , one may replace Equation 4 by any measure of lightness.

3 Derivation of a saturation term

The *chroma* (Equation 9) is the Euclidean distance (L_2 norm) of a colour projected onto the chromatic plane from the origin. It assumes its maximum value at the six corners of the hexagon projected onto the chromatic plane. Colours with hues which are not multiples of 60° have smaller maximum chroma values limited by the hexagon edges in Figure 2a. We use the term *saturation* to describe chroma normalised so that the maximum value is 1 for all hue values². In other words, the hexagon projected onto the chromatic plane is slightly deformed into a circle of unit radius by a normalisation factor, so that the saturation assumes its maximum value for all points with projections on the edges of the hexagon.

² This is an abuse of the term saturation as defined in Fairchild (1997), but it is used to be compatible with the image analysis literature.

3.1 Derivation of saturation expressions

Two equivalent saturation expressions can be derived in the chromatic plane. The saturation definition above implies that one obtains the saturation by dividing the the chroma c by the length of the line with the same hue extended to the edge of the hexagon, or $s = c/c_{\max}$. The red-yellow sector of the hexagon in the chromatic plane is reproduced in Figure 2b, in which the upper vertex is at the origin $(0, 0)$, 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 2b that

$$c_{\max} = \frac{\sqrt{3}}{2 \sin(120^\circ - H)} \quad (14)$$

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

$$H^* = H - k \times 60^\circ \text{ where } k \in \{0, 1, 2, 3, 4, 5\} \text{ so that } 0^\circ \leq H^* \leq 60^\circ \quad (15)$$

The value of the saturation $s \in [0, 1]$ is then

$$s = \frac{2c \sin(120^\circ - H^*)}{\sqrt{3}} \quad (16)$$

This definition is particularly useful for obtaining an efficient inverse transformation.

The alternative definition is obtained through the use of similar triangles. Using Figure 2b, we project point \mathbf{c}_p onto the \mathbf{or}_p axis along the line parallel to $\mathbf{r}_p\mathbf{y}_p$, giving the point \mathbf{c}_r . By similar triangles: $s = \frac{\|\mathbf{c}_p\|}{c_{\max}} = \frac{\|\mathbf{c}_r\|}{1}$, where $\|\cdot\|$ denotes the L_2 norm on the chromatic plane. On the line \mathbf{or}_p , one has $G = B$. Substituting this into Equation 10, one obtains either $s = R - G$ or $s = R - B$. Along the lines parallel to $\mathbf{r}_p\mathbf{y}_p$, only the G coordinate varies (see Section 2.1). Hence only the value of $R - B$ is the same at both points \mathbf{c}_p and \mathbf{c}_r . This is the same as $\max(R, G, B) - \min(R, G, B)$ for sector 1. The derivations for the other five sectors are done similarly. This leads to the general expression for the saturation

$$s = \max(R, G, B) - \min(R, G, B) \quad (17)$$

It can easily be shown that this expression is a norm in the chromatic plane.

3.2 Relation to existing cylindrical coordinate colour spaces

The Generalised LHS (GLHS) model by Levkowitz and Herman (1993) calculates the lightness-normalised saturation for a lightness of the form $L = w_{\min} \min(R, G, B) + w_{\text{mid}} \text{mid}(R, G, B) + w_{\max} \max(R, G, B)$, where $w_{\min} + w_{\text{mid}} + w_{\max} = 1$. The HSL and HSV models are obtained from the GLHS model for specific values of these weights. By changing the definition of the saturation to remove the lightness normalisation in the derivation of this model, one obtains the above unique saturation expression (Equation 17) for all possible lightness expressions (Hanbury, 2003). This saturation is called *Diff* in the HMMD colour space (Manjunath et al., 2002).

An alternative chroma based on the L_1 norm is derived and used in Hanbury and Serra (2003), Serra (2005) and Angulo and Serra (2007). For completeness, we present it here. The lightness I is defined by Equation 4. The chroma is then:

$$c_{L1} = \begin{cases} \frac{3}{2} (\max - I) & \text{if } I \geq \text{med} \\ \frac{3}{2} (I - \min) & \text{if } I \leq \text{med} \end{cases} \quad (18)$$

where $\max = \max(R, G, B)$, $\min = \min(R, G, B)$ and $\text{med} = \text{median}(R, G, B)$.

3.3 Comparison of the saturation and chroma formulations

We compare the distributions of the saturation and chroma formulations discussed: the $\max - \min$ saturation expression (equation 17), the L_2 norm chroma (equation 10), and the L_1 norm chroma (equation 18). 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. The saturation (chroma) values of each point are calculated (as floating point values), and then rounded to the nearest integer (only necessary for the L_2 norm chroma). Histograms showing the distribution of these integer values over 256 levels are shown in Figure 4.

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 rounding of its floating point values. 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 and the histogram is more regular.




Fig. 4. The saturation and chroma histograms.

4 Transformations to and from the cylindrical coordinate space

As a summary, efficient algorithms to calculate the lightness, hue and saturation from RGB coordinates are given here. The inverse transformation is also derived. MATLAB routines implementing the transformations are available on <http://www.prip.tuwien.ac.at/~hanbury>.

4.1 RGB to cylindrical coordinates

One calculates a lightness measure, a saturation or chroma measure and a hue. Due to the independence of the lightness and the saturation, one is free to choose any lightness measure, e.g. different coefficients in the sum in Equation 4, such as those for the luminance measure in Poynton (1997). Saturation using Equation 17 or chroma using Equation 10 or 18 may be used. The hue may be equivalently calculated using Equation 11 or Equations 12 and 13. Alternatively, one may use the hue measure based on the L1 norm and avoiding trigonometric functions described in Angulo and Serra (2007).

For example, using Equations 4, 17 and 11 gives

$$i = \frac{1}{3} (R + G + B) \tag{19}$$

$$s = \max(R, G, B) - \min(R, G, B) \tag{20}$$

$$h = \arctan \left(\frac{\sqrt{3}(G - B)}{2R - G - B} \right) \quad (21)$$

4.2 Cylindrical coordinates to RGB

One possible inverse transformation is given here. One first calculates the chroma values from the saturation values (using Equation 16)

$$c = \frac{\sqrt{3}s}{2 \sin(120^\circ - H^*)} \quad (22)$$

where H^* is given by Equation 15. From the chroma, one calculates

$$c_1 = c \cos(h) \quad (23)$$

$$c_2 = c \sin(h) \quad (24)$$

For the case where the hue is undefined: $c_1 = c_2 = 0$. To get the R , G and B values, the following are used (derived from Equations 4 to 6):

$$R = I + \frac{2}{3}c_1 \quad (25)$$

$$G = I - \frac{1}{3}c_1 + \frac{1}{\sqrt{3}}c_2 \quad (26)$$

$$B = I - \frac{1}{3}c_1 - \frac{1}{\sqrt{3}}c_2 \quad (27)$$

5 Applications

Because the saturation or chroma terms are not percentages, it is possible to directly compare them. This allows them to be used directly in colour morphology. Colour morphology using lexicographical cascades with the max – min saturation is discussed in Hanbury and Serra (2001a), while lexicographical cascades in the L1 norm colour space are discussed in Angulo (2007). More information on colour morphology can be found in a recent comparative study by Aptoula and Lefèvre (2007).

The saturation or chroma measurements in this colour space have been used as a weight differentiating between achromatic and chromatic pixels. The

saturation-weighting based colour gradient (Angulo and Serra, 2003), defined as

$$\nabla \mathbf{f}(x) = s(x) \times \nabla_c h(x) + [1 - s(x)] \times \nabla i(x) \quad (28)$$

gives a higher weight to the hue gradient $\nabla_c h(x)$ if the saturation is high and a higher weight to the lightness gradient $\nabla i(x)$ if the saturation is low. The notation ∇_c represents a circular centred gradient which takes the angular nature of the hue component into account (Hanbury and Serra, 2001b). This type of saturation-weighting has also been useful in the calculation of hue statistics. As the hue is an angular value this should be done using circular statistics (Fisher, 1993). The mean direction of a set of angular values is calculated as follows: given a set of n angular values $\{a_1, a_2, \dots, a_n\}$, one takes for each value a_i a unit vector $\hat{\mathbf{a}}_i$ in the 2D plane with direction a_i . The mean direction is the direction of the vector resulting from taking the vector sum $\sum_i \hat{\mathbf{a}}_i$. While this can be directly applied to calculating a mean hue, hues associated with low saturations should contribute less to the mean hue. This is done by using a weighted vector sum as follows: given n hue/saturation pairs $\{(h_1, s_1), (h_2, s_2), \dots, (h_n, s_n)\}$, the mean hue direction is the direction of the vector resulting from taking the weighted vector sum $\sum_i s_i \hat{\mathbf{h}}_i$, where $\hat{\mathbf{h}}_i$ is the unit vector with direction h_i . These weighted hue statistics have been successfully used in maintaining a background model for video surveillance (Blauensteiner et al., 2006).

Angulo and Serra (2007) have made extensive use of the L_1 norm chroma in the creation of 2D lightness/chroma histograms. Alignments in these histograms have been shown to correspond to interesting regions of an image, such as highlights.

The saturation measures normalised by lightness, as found in the standard HSV and HSL models, also have their uses. One example is in a visualisation application where one wishes to replace the lightness or saturation of an image by another function in order to enhance a certain aspect of the image. Changing the lightness or chroma/saturation in the space derived in this paper leads to the possibility that a colour may be moved out of the valid colour gamut. An inverse transform to RGB could then lead to invalid RGB coordinates (i.e. outside the RGB cube). An application example is setting all pixels in an image with saturation above a threshold to have maximum saturation (Demarty and Beucher, 1998). While the saturation thresholding should be done using one of the saturation/chroma expressions presented in this paper, it is easier to transform the pixel colours in the HSL or HSV space. Because the saturation is expressed as a percentage in these spaces, one simply sets the saturation of the relevant pixels to the maximum possible value without the possibility of moving the colour outside the gamut. Another application is replacing the lightness band of a colour image of a painting by a greyscale

image of the same painting captured in the infrared band, thereby making the relation of the underdrawing³ to the colours more visible (Kammerer et al., 2004).

6 Conclusion

We have derived a cylindrical coordinate colour representation based on an opponent colour space in the RGB colour space. While a number of opponent colour spaces are available in the literature, there is no standardised axis orientation, leading to a large variability in the colour represented by the hue origin. Furthermore, the proposed spaces tend to have different normalisations leading to differently shaped regions obtained by projecting the RGB cube onto the chromatic plane. We have proposed using an OCS which has its axes oriented so that a hue of zero corresponds to red, as for the CIELAB space and the spaces based on the GLHS model (Levkowitz and Herman, 1993). Furthermore, the proposed OCS is normalised so that the area projected onto the chromatic plane is in the form of a hexagon with each corner at a unit Euclidean distance from the centre.

On the chromatic plane, we have shown how, in addition to using the distance from the centre as a measure of chroma, one can derive the expression for saturation $s = \max(R, G, B) - \min(R, G, B)$. While the chroma can only assume a maximum value of 1 for hue values which are multiples of 60° , saturation can assume a maximum value of 1 for all hue values. We have also shown how the saturation of the GLHS model reduces to this saturation if the normalisation by lightness is removed.

From the discussion of applications, it is clear that both models without saturation normalisation by lightness (such as the proposed model), as well as models with lightness-normalised saturation (such as the GLHS model) have appropriate application areas.

References

- Angulo, J., 2007. Morphological colour operators in totally ordered lattices based on distances: Application to image filtering, enhancement and analysis. *Computer Vision and Image Understanding* 107, 56–73.
- Angulo, J., Serra, J., 2003. Color segmentation by ordered mergings. In: *Proc. of the Int. Conf. on Image Processing*. Vol. II. pp. 125–128.

³ The underdrawing is a preliminary sketch made by an artist when planning a painting. It can often be made visible by infrared reflectography.

- Angulo, J., Serra, J., 2007. Modelling and segmentation of colour images in polar representations. *Image and Vision Computing* 25 (4), 475–495.
- Aptoula, E., Lefèvre, S., 2007. A comparative study on multivariate mathematical morphology. *Pattern Recognition* 40, 2914–2929.
- Blauensteiner, P., Wildenauer, H., Hanbury, A., Kampel, M., 2006. Motion and shadow detection with an improved colour model. In: *Proceedings of the IEEE Int. Conf. on Signal and Image Processing (ICSIP)*. pp. 627–632.
- Demarty, C.-H., Beucher, S., 1998. Color segmentation algorithm using an HLS transformation. In: *Proceedings of the International Symposium on Mathematical Morphology (ISMM '98)*. pp. 231–238.
- Fairchild, M. D., 1997. *Color Appearance Models*. Addison-Wesley.
- Fisher, N. I., 1993. *Statistical Analysis of Circular Data*. Cambridge University Press.
- Gonzalez, R. C., Woods, R. E., 1992. *Digital Image Processing*. Prentice Hall.
- Hanbury, A., 2003. A 3D-polar coordinate colour representation well adapted to image analysis. In: *Proceedings of the Scandinavian Conference on Image Analysis (SCIA)*. pp. 804–811.
- Hanbury, A., Serra, J., 2001a. Mathematical morphology in the HLS colour space. In: *Proc. of the British Machine Vision Conference 2001*. pp. 451–460.
- Hanbury, A., Serra, J., December 2001b. Morphological operators on the unit circle. *IEEE Transactions on Image Processing* 10 (12), 1842–1850.
- Hanbury, A., Serra, J., 2003. Colour image analysis in 3D-polar coordinates. In: *Proc. of the DAGM'03 conference*. pp. 124–131.
- Kammerer, P., Hanbury, A., Zolda, E., 2004. A visualisation tool for comparing paintings and their underdrawings. In: *Proc. of the EVA*. pp. 148–153.
- Lambert, P., Carron, T., 1999. Symbolic fusion of luminance-hue-chroma features for region segmentation. *Pattern Recognition* 32, 1857–1872.
- Levkowitz, H., Herman, G. T., 1993. GLHS: A generalised lightness, hue and saturation color model. *CVGIP: Graphical Models and Image Processing* 55 (4), 271–285.
- Manjunath, B., Salembier, P., Sikora, T. (Eds.), 2002. *Introduction to MPEG-7: Multimedia Content Description Interface*. Wiley.
- Plataniotis, K. N., Venetsanopoulos, A. N., 2000. *Color Image Processing and Applications*. Springer.
- Poynton, C., 1997. Frequently asked questions about color. URL: <http://www.poynton.com/PDFs/ColorFAQ.pdf>.
- Pratt, W. K., 1991. *Digital Image Processing, 2nd Edition*. Wiley.
- Serra, J., 2005. Morphological segmentations of colour images. In: *Mathematical Morphology: 40 years on — Proceedings of the 7th International Symposium on Mathematical Morphology*. pp. 151–176.
- Smith, A. R., 1978. Color gamut transform pairs. *Computer Graphics* 12 (3), 12–19.
- van de Weijer, J., Gevers, T., Geusebroek, J.-M., 2005. Edge and corner detection by photometric quasi-invariants. *IEEE Trans. Pattern Analysis and Machine Intelligence* 27 (4), 625–630.