Circular Statistics Applied to Colour Images

Allan Hanbury*

PRIP, TU Wien, Favoritenstraße 9/1832, A-1040 Vienna, Austria
hanbury@prip.tuwien.ac.at

Abstract *Three methods for summarising the characteristics of colour images are presented. They all rely on the ability to represent colours in terms of a 3D-polar coordinate system having independent brightness and saturation coordinates. The methods discussed are: hue statistics and saturation-weighted hue statistics, saturation-weighted hue histograms, and colour statistics histograms. The latter are a way of visualising hue, saturation and luminance information in a single uni-dimensional histogram, allowing a compact summary of a colour image.*

1 Introduction

A 3D-polar coordinate colour space is an alternative and often more intuitive representation of the RGB space. It is constructed by placing an axis, the *achromatic axis*, between the origin and the point $[R_{\max}, G_{\max}, B_{\max}]$ in the RGB space, and then defining 3D-polar coordinates with respect to this axis. A colour is then specified in terms of:

- *Brightness* or *luminance* $L \in [0, 1]$: gives the position on the achromatic axis (note that these terms do not represent the same concept, their precise definitions are given in [5]).
- *Hue* $H \in [0°, 360°]$: is an angular measure around the achromatic axis with respect to an origin at pure red. Because of the angular nature of this component, it requires specially adapted image processing operators.
- *Saturation* $S \in [0, 1]$: the distance from the achromatic axis.

Many of the plethora of available 3D-polar coordinate colour representations, such as HSV, HLS, HSI, HSB, etc., are ill-suited to image analysis and processing, having been developed primarily for easy numerical colour specification in computer graphics applications. The "natural" shapes of their colour gamuts have therefore been artificially expanded from cones and bi-cones into cylinders, which simplifies the checking of whether a specified colour is valid. We have suggested in [3, 5] the use of a 3D-polar coordinate colour space well suited to image analysis and processing tasks, which we called the *Improved Hue, Luminance and Saturation (IHLS) space*. Its main advantages over the commonly used 3D-polar coordinate spaces are:

- Achromatic or near-achromatic colours always have a low saturation value.
- The saturation and brightness coordinates are independent, as the normalisation of the saturation by the brightness function present in most of the systems having cylindrically-shaped colour gamuts has been removed.
- Comparisons between saturation values are meaningful, also due to the saturation normalisation having being removed.

An algorithm for calculating the 3D-polar coordinates of an RGB vector is given in the appendix. In addition, MATLAB routines implementing transformations between the RGB and IHLS spaces (summarised in [5]) are available at `http://www.prip.tuwien.ac.at/~hanbury`.

In this paper, we present three methods for summarising the characteristics of colour images, which demonstrate the usefulness of the IHLS space. They all rely on the first two of the three advantages of this space listed above[1]. The first method is the calculation of colour statistics (section 2), and the second is the calculation of saturation-weighted hue histograms (section 3). Finally, in section 4, we apply colour statistics to the generation of histograms which include information on the hue, luminance and saturation.

2 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 [1] should be calculated for it. After reviewing how the standard circular statistics are applied to the hue channel, we introduce the concept of saturation-weighted hue statistics.

2.1 Hue statistics

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

$$\overline{H} = \arctan\left(\frac{B}{A}\right) \tag{1}$$

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

[1]The third advantage is particularly useful for mathematical morphology applied to colour images.

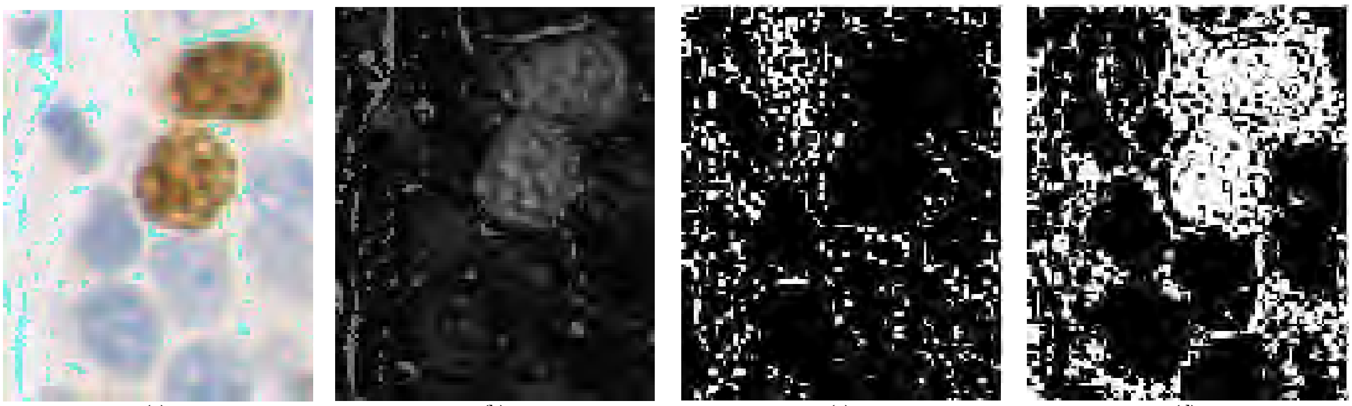

(a) (b) (c) (d)

Figure 1: (a) Colour image. (b) Saturation of image (a). (c) Pixels of image (a) with hue values in the interval $20°$ on each side of the non-weighted hue mean $\left[\overline{H} - 20°, \overline{H} + 20°\right]$. (d) Pixels of image (a) with hue values in the interval $\left[\overline{H}_S - 20°, \overline{H}_S + 20°\right]$ around the saturation-weighted hue mean.

where

$$A = \sum_i \cos H_i, \quad B = \sum_i \sin H_i \tag{2}$$

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

$$R = \sqrt{A^2 + B^2} \tag{3}$$

and its *mean length* is

$$\overline{R} = \frac{R}{n} \tag{4}$$

The mean length can be interpreted as the length of the actual resultant vector R compared to the length of the resultant vector which would be obtained if all the H_i were identical. 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. The *circular variance* is defined as

$$V = 1 - \overline{R}$$

This variance differs from the standard linear statistical variance in being limited to the range $[0, 1]$, however it is similar in that lower values of V represent less dispersed data. Further measures of circular data distribution are given in [1].

2.2 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 colours with zero saturation, 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 2 by

$$A_S = \sum_{i=1}^{n} S_i \cos H_i, \quad B_S = \sum_{i=1}^{n} S_i \sin H_i \tag{5}$$

and replacing A and B in equation 1 by A_S and B_S. We denote by $\overline{H}_S$ the resultant saturation-weighted hue mean. For the mean length (equation 4), there are two possibilities. One can either divide the length of the resultant vector by the length of the vector obtained if all the vectors were aligned, i.e. the sum of the S_i values:

$$\overline{R}_S = \frac{\sqrt{A_S^2 + B_S^2}}{\sum_{i=1}^{n} S_i} \tag{6}$$

Alternatively, one can use

$$\overline{R}_n = \frac{\sqrt{A_S^2 + B_S^2}}{n} \tag{7}$$

in which the length of the resultant vector is compared to the length obtained if all the vectors had had the same direction and maximum saturation. This has the effect of weighting the mean length by the saturation values. For example, for a set of n identical vectors of hue H and saturation S, the mean length $\overline{R}_S$ would be 1 irrespective of the value of S, whereas $\overline{R}_n$ would be equal to S. The value of $\overline{R}_n$ therefore gives an indication of the saturations of the vectors which gave rise to the hue mean as well as an indication of the angular dispersion of the vectors.

In practice, for images which contain only strongly saturated colours, there is not a significant difference between the values of the weighted and unweighted hue means. Figure 1a shows an image in which this difference is important. As is visible in figure 1b, the saturation of the two brown cells is higher than the saturation of the surroundings. For this image, the unweighted hue mean is $\overline{H} = 326.9°$, and the saturation-weighted hue mean is $\overline{H}_S = 19.7°$.

To show the difference, thresholds on the hue component were calculated for the intervals $\left[\overline{H} - 20^\circ, \overline{H} + 20^\circ\right]$ and $\left[\overline{H}_S - 20^\circ, \overline{H}_S + 20^\circ\right]$, and these are shown in figures 1c and 1d 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.

An alternative saturation weighting for the hue, which shifts the hue values around the circle, is described in [4], where it is used in the context of colour ordering for mathematical morphology.

3 Saturation-weighted 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 which the hue has little significance. However, as the saturation term of commonly used 3D-polar colour representations is essentially useless in discriminating between achromatic and chromatic colours, a number of heuristics, summarised by Stokman and Gevers [8], have been used. Tico et al. [10], 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 of the IHLS space can be directly used as such a weight. In building the saturation-weighted hue histogram for the specific case of the hues (measured in degrees) having been rounded to the nearest integer, the total in bin θ ($\theta \in [0^\circ, 1^\circ, \ldots, 360^\circ]$) of the histogram is simply calculated as

$$W_\theta = \sum_x S_x \delta_{\theta H_x} \quad (8)$$

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 comparison of a histogram calculated directly on the hue channel of a colour image to the saturation-weighted hue histogram described by equation 8 is given in figure 2. Figures 2b, 2c and 2d are respectively the hue, saturation and luminance channels of figure 2a. Upon examining the hue image, one can see that for the low saturation (black or white) regions, the hue can take on a wide range of values, which do not correspond to any significant colour differences in the colour image. The simple hue histogram, calculated by counting the number of times each hue value occurs in the hue channel (i.e. equation 8 with S_x set permanently to unity), is shown in figure 2e. The most noticeable property of this histogram is the presence of equally-spaced spikes (the spike at the zero hue bin has a height of about 30000). These spikes are due to the discretisation errors introduced by using a polar coordinate system on a discrete grid, first pointed out by Kender [6], and their heights are amplified for images containing large numbers of achromatic pixels. The saturation-weighted hue histogram is shown in figure 2f. It is clear that the amplitudes of the spikes have been greatly reduced, leading to a more easily interpretable histogram. As an example, we have extracted the pixels of the hue image which correspond to two of the most pronounced peaks in this histogram. Figure 3a is the result of a threshold which extracts pixels having hue values between 210° and 240°, and figure 3b shows those pixels having hue values between 30° and 50°. One can see that these correspond to the colours of the most highly saturated regions in the image, but these peaks are not easily discernible on the simple hue histogram. The additional background pixels extracted by these thresholds could easily be removed by combining the hue threshold with a threshold on the saturation.

A saturation-weighted luminance histogram can be calculated analogously by using the inverse weighting $(1 - S_x)$, thereby privileging the low saturation (achromatic) pixels.

4 Colour statistics histograms

Histograms are a very common and useful tool in the analysis of greyscale images. In this section, we present a similar tool for use with colour images, a histogram which encodes the dominant hue as a function of the luminance, for which the saturation-weighted hue statistics introduced above are used.

Building a greyscale histogram is very simple, one counts the number of pixels at each greylevel. For colour images, we propose to compute the saturation-weighted hue mean and mean length for the pixels at each luminance level. Given a colour image in the IHLS space, the luminance values are first quantised into $N+1$ levels labeled by $\ell = \{0, 1, 2, \ldots, N\}$. Then, for each value of ℓ, the following circular statistics descriptors are calculated

$$A_{S\ell} = \sum_x S_x \cos H_x \delta_{L_x \ell}, \; B_{S\ell} = \sum_x S_x \sin H_x \delta_{L_x \ell} \quad (9)$$

$$\overline{H}_{S\ell} = \arctan\left(\frac{B_{S\ell}}{A_{S\ell}}\right) \quad (10)$$

$$\overline{R}_{n\ell} = \frac{\sqrt{A_{S\ell}^2 + B_{S\ell}^2}}{\sum_x \delta_{L_x \ell}} \quad (11)$$

where H_x, L_x and S_x are the hue, luminance and saturation at position x in the image, and the sums are over all the pixels in the image. These equations are essentially the saturation-weighted hue statistical functions with the addition of Kronecker deltas to limit them to only one luminance level. Equation 11 corresponds to equation 7.

We therefore have two histograms of colour information, the mean hue and its associated mean length as a function of luminance. These could conceivably be used directly in image matching and database retrieval applications. For visualisation purposes, these two histograms can very simply be combined into a single histogram, in which the height of the bar at luminance ℓ corresponds to the mean length $\overline{R}_{n\ell}$, and its colour is given by the fully saturated colour corresponding to the mean hue $\overline{H}_{n\ell}$. As the mean hue associated

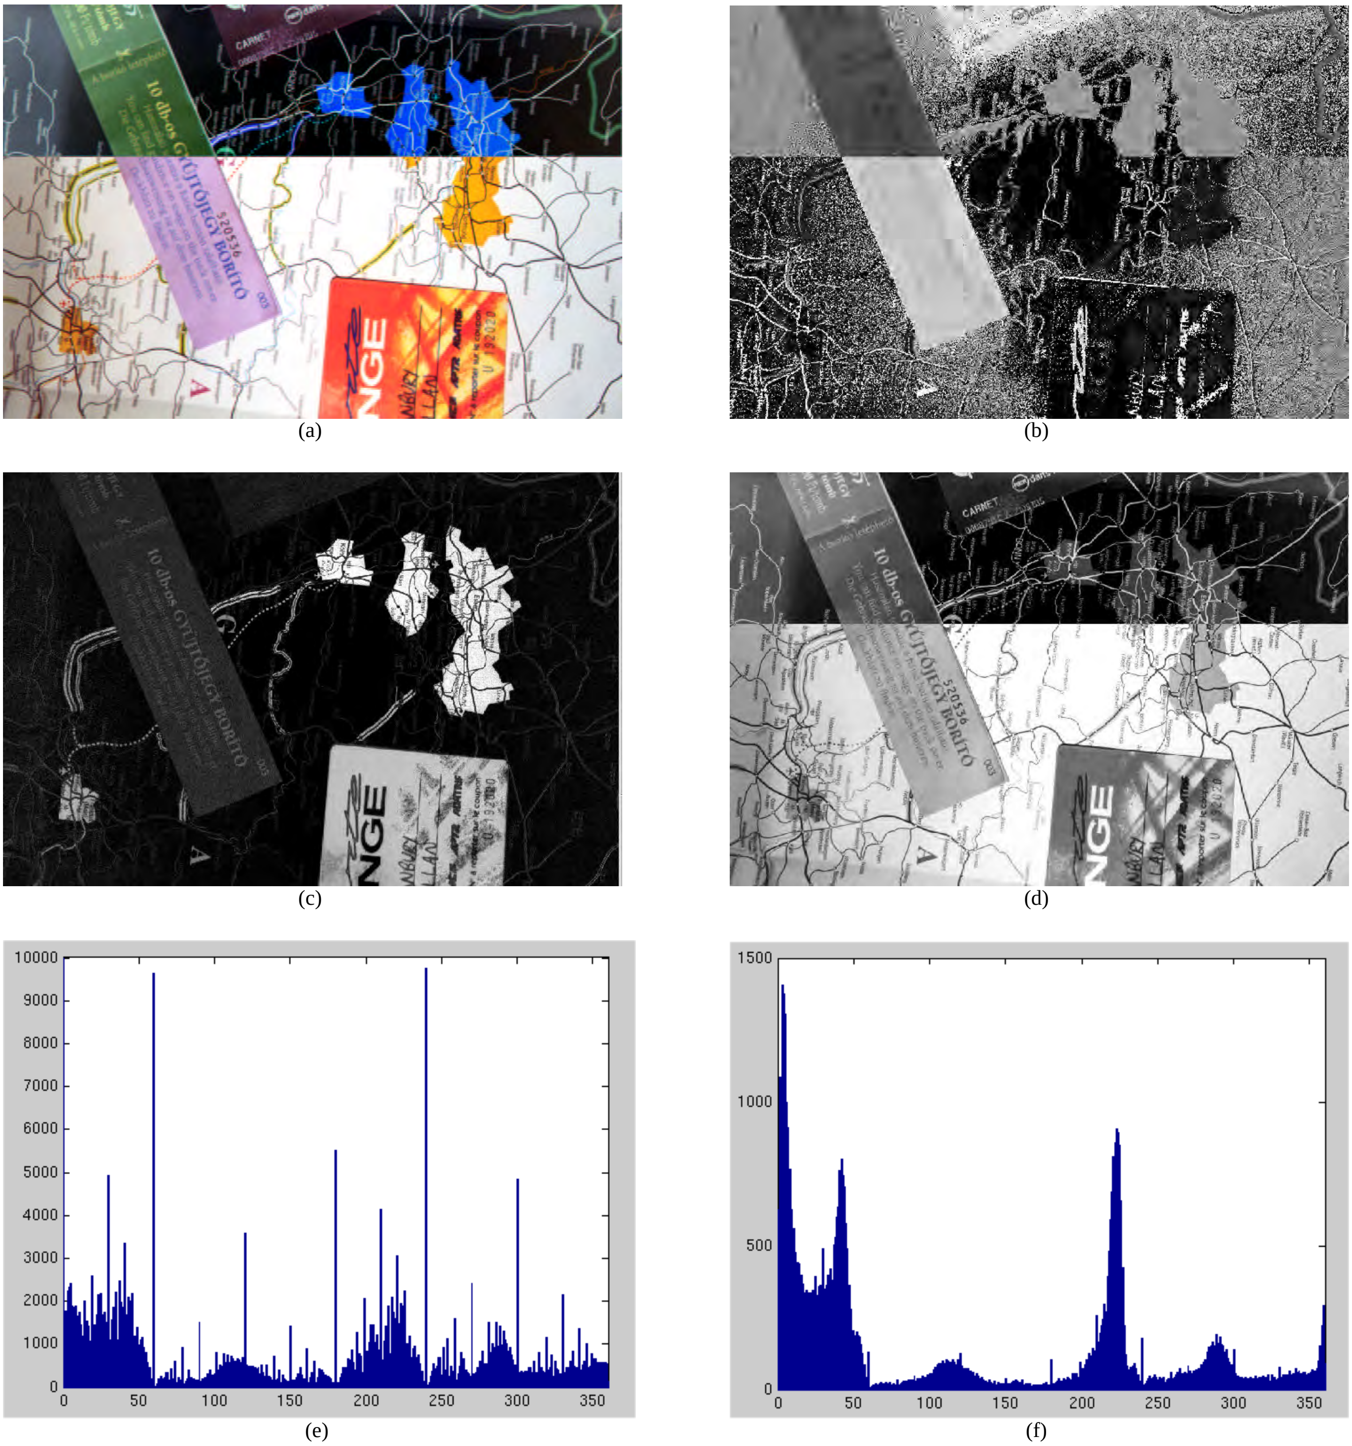

Figure 2: (a) Colour image. (b) Hue of image (a). (c) Saturation of image (a). (d) Luminance of image (a). (e) Hue histogram. (f) Saturation-weighted hue histogram.

with a very low mean length value does not give much information, we set the colours of the bars with a mean length below a threshold to the greylevel corresponding to the associated luminance value (for the examples shown here, this threshold was set to 0.05).

An example of these colour statistics histograms is given in figure 4. Figures 4b, 4c and 4d are respectively the hue, saturation and luminance of figure 4a. The smoothing and artifacts visible in the hue and saturation images are present because the initial colour image was JPEG compressed. The colour statistics histogram of this image is shown in figure 4e, for which the luminance quantisation parameter $N = 100$. One sees clearly from the histogram that the dominant colour for the low luminance parts (the vegetation) is green, the very uniform (i.e. with high mean length) blues associated with the water, and the reds of the highly luminous roofs. For comparison, the colour statistics histogram of the image in figure 2a is shown in figure 4f. In

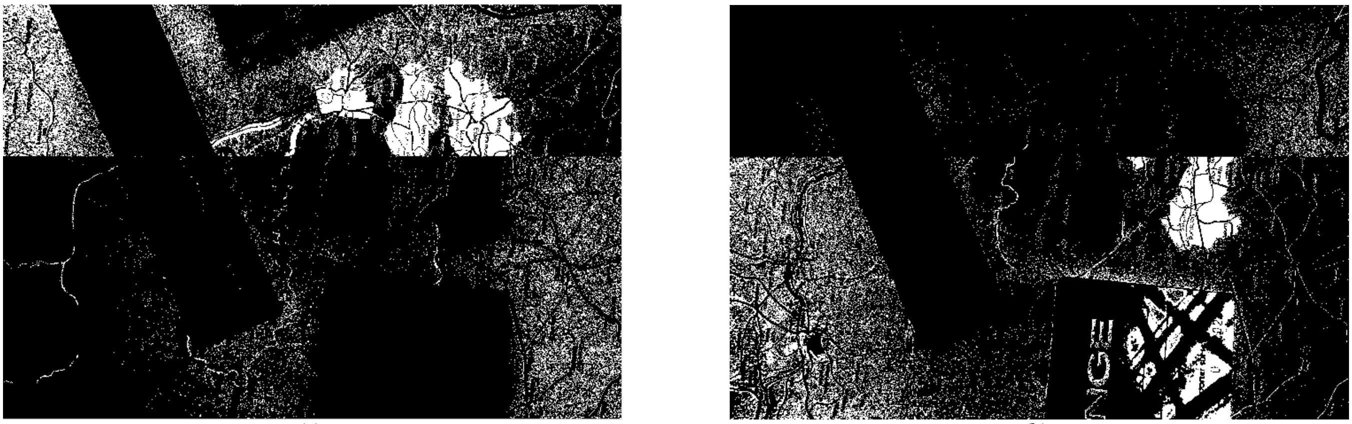

(a) (b)

Figure 3: (a) Hue pixels of figure 2b with values between 210° and 240°. (b) Hue pixels of figure 2b with values between 30° and 50°.

this histogram, the low luminance blues and purples in the top part of the colour image are clearly separated from the high luminance yellows and oranges in the bottom part of the image.

5 Conclusion

When dealing with hue data, it is of course imperative that circular statistics be used. For example, if one calculates the standard linear mean for the hue of an image containing hue values only in the red region of the hue circle (i.e. on either side of the hue origin), this hue mean would be positioned on the opposite side of the circle, in the same place as for an image containing only cyan hues. In this paper, we suggest three methods of statistically describing colour image data. The first is saturation-weighted hue circular statistics, an extension of standard circular statistics allowing the saturation to be taken into account with the hue, the second is saturation-weighted hue histograms, and the third involves combining colour statistics with luminance histograms. It should be noted that the use of these three methods relies on the "well-behaved" saturation coordinate of the IHLS space, and that they will not give usable results if used in conjunction with the standard cylindrically-shaped 3D-polar coordinate colour spaces, such as HLS, HSV, etc.

Applications of these statistical colour descriptors can be found, for example, in image retrieval from databases, for which colour matching is an important component [7]. The use of the first three statistical moments (mean, variance and skewness) of each channel in the HSV colour space as a nine-component image colour feature vector has been suggested [9]. However, standard linear instead of circular statistical formulae were used for the hue channel. Given the possible incorrect hue means that could appear through this, as illustrated above, experiments are currently underway to investigate the magnitude of the improvement in retrieval results when the correct statistical formulae are used. The suggested saturation-weighted hue histograms have the advantage over those described in [10] of having no parameters to be set by the user. The difference in retrieval performance remains to be investigated. Finally, the use of these statistical descriptors in conjunction with colour invariant features [2] should still be explored. The implementation of colour statistics histograms in image analysis software as a simple description of colour images could also be useful.

A RGB to IHLS transform

The simplest algorithm for calculating the 3D-polar coordinates of a vector $\mathbf{c} = (R, G, B)$ in the RGB space is given here. The coordinates calculated are the luminance Y, the saturation S and the hue H (H' is an interim result in the calculation of H).

$$\begin{aligned}
Y &= 0.2126R + 0.7152G + 0.0722B \\
S &= \max(R, G, B) - \min(R, G, B) \\
H' &= \arccos\left[\frac{R - \frac{1}{2}G - \frac{1}{2}B}{(R^2 + G^2 + B^2 - RG - RB - BG)^{\frac{1}{2}}}\right] \\
H &= \begin{cases} 360^\circ - H' & \text{if } B > G \\ H' & \text{otherwise} \end{cases}
\end{aligned}$$

References

[1] N. I. Fisher. *Statistical Analysis of Circular Data*. Cambridge University Press, 1993.

[2] 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.

[3] A. Hanbury. The taming of the hue, saturation and brightness colour space. In *Proceedings of the 7th CVWW, Bad Aussee, Austria*, 2002.

[4] A. Hanbury and J. Serra. Mathematical morphology in the HLS colour space. In T. Cootes and C. Taylor, editors, *Proceedings of the British Machine Vision Conference 2001*, pages 451–460. BMVA, 2001.

[5] A. Hanbury and J. Serra. A 3D-polar coordinate colour representation suitable for image analysis. Technical Report PRIP-TR-077, TU Wien, 2002.

[6] J. R. Kender. Saturation, hue and normalised color: Calculation, digitization effects, and use. Technical report, Department of Computer Science, Carnegie-Mellon University, 1976.

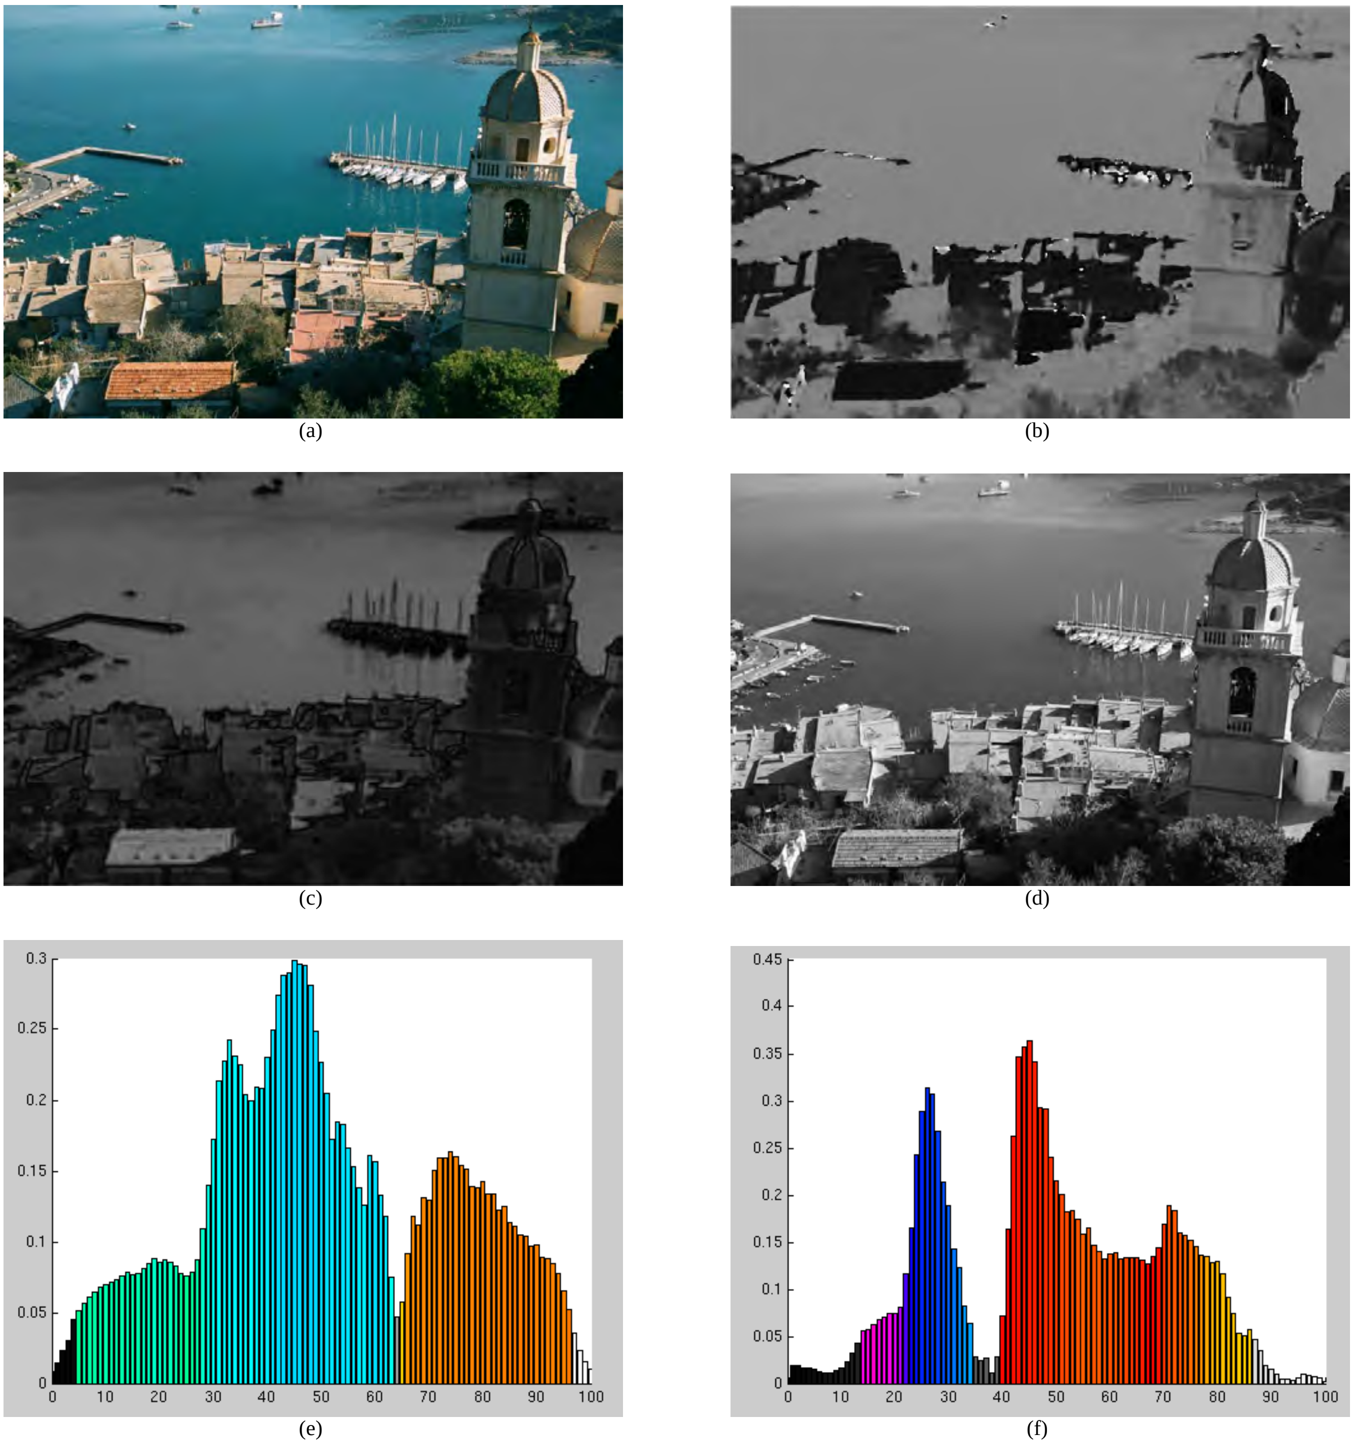

Figure 4: (a) Colour image (from the University of Washington content-based image retrieval database). (b) Hue of image (a). (c) Saturation of image (a). (d) Luminance of image (a). (e) Colour statistics histogram of image (a). (f) Colour statistics histogram of figure 2a.

[7] R. Schettini, G. Ciocca, and S. Zuffi. A survey of methods for colour image indexing and retrieval in image databases. In R. Luo and L. MacDonald, editors, *Color Imaging Science: Exploiting Digital Media*. John Wiley, 2001.

[8] H. Stokman and T. Gevers. Error propagation due to color space transforms. In *Proceedings of SPIE, San Jose*, January 2000.

[9] M. Stricker and M. Orengo. Similarity of color images. In *Storage and Retrieval for Image and Video Databases (SPIE)*, pages 381–392, 1995.

[10] M. Tico, T. Haverinen, and P. Kuosmanen. A method of color histogram creation for image retrieval. In *Proceedings of NORSIG*, 2000.