A novel approach to hue ordering

David Connah[1], Graham D. Finlayson[1] and Marina Bloj[2]; 1) Department of Computer Science, University of East Anglia, UK; 2) Bradford Optometry Colour and Lighting Lab, School of Life Sciences, University of Bradford, UK.

Abstract

Applying morphological operators to colour images requires that the colour vectors can be ordered. In this paper we propose an alternative to traditional lexicographic ordering methods by replacing the hue component with a new component named anti-saturation, which is the distance from the given colour vector to the boundary of the colour cube. While anti-saturation does not permit a total ordering, as does the hue angle, it both varies smoothly with hue angle and does not contain any discontinuities, making it a plausible candidate for a third level of lexicographic ordering instead of hue-angle.

Background

To apply morphological operators to images requires that the pixel values can be ordered. The erosion and dilation operators, for example, are defined with relation to the greatest lower bound and lowest upper bound of a set of pixel values; clearly these quantities have no meaning without the notion of order. Greyscale images are, by definition, scalar-valued at each location; thus, the order of the pixels is well-defined. Colour images, however, are vector-valued at each pixel, so the order is not well-defined. As a result, in order to apply morphological operators to colour images, some ordering principle must be devised.

Conceptually the simplest method is simply to apply morphological operators in the three colour channels independently (a so-called *marginal* ordering [1]). A well-reported problem with this approach is that "false colours" can be generated; for example, a dilation operation performed at a blue-green boundary will expand the blue and green regions such that they overlap, creating a cyan colour that is not present in the original image (see Figure 1).

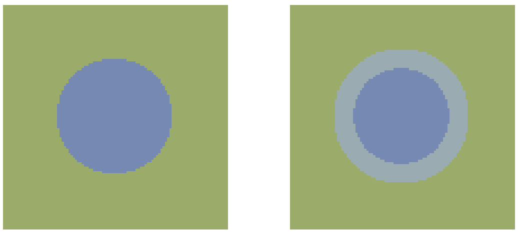

(a) Original image (b) Dilated image

Figure 1. *A colour image before (a) and after (b) dilation using a disk shaped structuring element and a marginal ordering scheme. By applying the operators in the three colour channels separately, false colours (i.e. colours that are not present in the original image) can be generated.*

Comer and Delp [2], suggest a simple *reduced* ordering strategy by projecting the 3D colour values onto a single 1D value; this can be the projection onto the first principal component of the image data, or a more perceptually relevant variable such as luminance. The problem with this approach is that there are still likely to be many boundaries in the image where adjacent colours have similar luminance values, and thus morphological operations cannot be performed [3].

A more involved approach, used extensively by Hanbury and Serra [4, 5, 6] is to use a lexicographic ordering (also called a *conditional* ordering). This approach is similar to that used in the dictionary to order words: ordering is first performed on a particular 1D variable, but if two colours are equal with respect to this variable, then they are compared based upon a second variable. A simple example could be to order based upon just the red-channel, but if the red value of two colours is equal, then order them based upon the green channel.

Rather than use RGB space, Hanbury and Serra make extensive use of more perceptually motivated colour-spaces such as HSV, HSL and CIELAB. Thus the primary ordering variable is luminance (or value), the second is saturation, and the third is hue, although they need not necessarily be applied in this order. The main problem with this approach is how to deal with the hue component. Hue is defined as an angle on the interval $[0^\circ, 360^\circ)$, or $[-180^\circ, 180^\circ)$; thus at some hue values there is an inevitable discontinuity, where two colours with almost identical hue values can be seen as polar opposites in terms of their order.

In this paper we propose an alternative to the traditional ordering of hues that does not contain a discontinuity. Instead of calculating hue angle we calculate the distance from a given colour vector to the boundary of the colour cube in the direction orthogonal to the luminance axis; we call this value *anti-saturation*. This value varies smoothly with hue, and has maximal values close to the major colour-space axes. Although it is not unique (two colours with different hue angles can have the same anti-saturation value), and thus does not facilitate a total-ordering of the colour-vectors, we argue that the resulting ordering of vectors is more natural. We present images that show the new method currently gives a plausible ordering of hues, in relation to hue angle, but that it could be improved by enhancing its overall smoothness.

Method

The starting point is to define a three-component colour vector $\mathbf{c} = \{c_1, c_2, c_3\}$; we can denote this vector, for example, $\mathbf{c}_{RGB} = [c_R, c_G, c_B]$ or $\mathbf{c}_{HSV} = \{c_H, c_S, c_V\}$, depending upon which colour space is used to represent the vector. We now create a new three component colour representation, which we can term $\mathbf{c}_{LSA} = [c_L, c_S, c_A]$ for lightness[1], saturation, and anti-saturation. The three components of this representation can be computed as

[1]Here we do not refer to the perceptual attribute of lightness defined by the CIE, but use the term only to describe the achromatic nature of this value

follows

$$c_L = \frac{1}{3}\sum_i c_i, i \in \{R,G,B\}. \tag{1}$$

By defining an orthogonal projection operator $\mathbf{O}$ that projects a vector onto the space orthogonal to the vector $[1,1,1]^T$, the saturation component can be written as

$$c_S = \|\mathbf{O}\mathbf{c}_{RGB}\|_2 \tag{2}$$

where $\|\cdot\|_2$ computes the 2-norm. Given that the operator $\mathbf{O}$ projects $\mathbf{c}_{RGB}$ orthogonally to the lightness axis, the lightness and saturation components are fully independent of one another.

Now, in order to define the anti-saturation, we must first define the vector:

$$\mathbf{c}_{bound} = \alpha\mathbf{O}\mathbf{c}_{RGB} + c_L[1,1,1]^T \tag{3}$$

In words, this equation represents a point $\mathbf{c}_{bound}$ that is the sum of two vector components: the projection of $\mathbf{c}_{RGB}$ onto the luminance axis, and the projection of $\mathbf{c}_{RGB}$ onto the space orthogonal to the luminance axis. The positive scalar α is chosen such that the point $\mathbf{c}_{bound}$ lies on the boundary of the *RGB* colour cube. We clarify here that, given that *R*, *G* and *B* values fall on the range [0,1], the set of all *RGB* vectors fall within a cube, and it is the sides of this cube that form its boundary. The anti-saturation value is now given by:

$$c_A = \|\mathbf{c}_{bound} - (\mathbf{O}\mathbf{c}_{RGB} + c_L[1,1,1]^T)\| \tag{4}$$

i.e. it is the distance, in the plane orthogonal to the $[1,1,1]^T$ axis, from the RGB colour vector to the boundary of the colour cube.

A diagrammatic representation of the three colour variables is given in Figure 2, where just two colour axes are plotted for ease of interpretation. In this Figure the red vector represents a given colour vector, and the values of c_L, c_S, and c_A are shown with relation to the vector $[1,1,1]^T$.

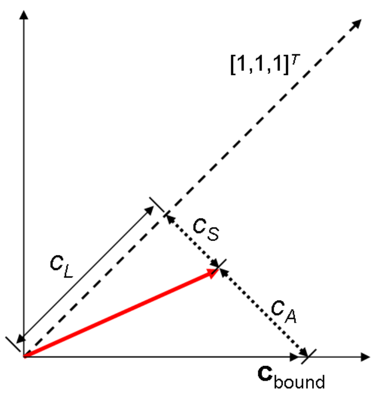

Figure 2. *Diagrammatic representation of the three colour variables for a given colour vector (shown in red); c_L is the length of the projection onto the $[1,1,1]$ axis, c_S is the saturation component, and c_A is the anti-saturation.*

Note that anti-saturation is dependent upon saturation, thus colours with low saturation also have high anti-saturation. As a result we cannot order images based upon anti-saturation alone, since this will result in vectors of higher saturation being considered of lower magnitude than vectors of higher saturation, which is counterintuitive. We therefore propose that the anti-saturation is used only after the saturation value in the lexicographic order.

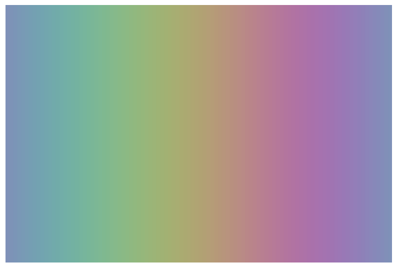

(a) Original spectrum

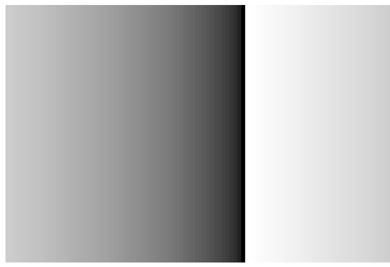

(b) Hue-angle value

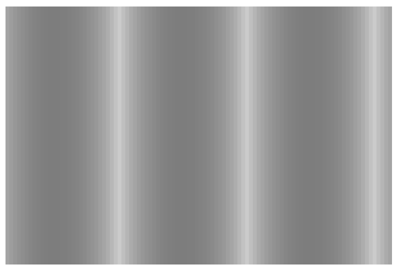

(c) Anti-saturation value

Figure 3. *An equiluminant and equisaturation spectrum for different hue values (a) and two different orderings represented by greyscale images. The ordering in (b) shows the traditional hue angle (scaled from 0 to 255), while the image in (c) shows the anti-saturation image (also scaled from 0 to 255).*

In Figure 3 we show a greyscale visualisation of the method for an equi-luminant, equi-saturation, colour spectrum (i.e., an image where the lightness and saturation components are constant). The first of the two greyscale images shows the hue angle (scaled from 0 to 255); here there is a clear discontinuity which occurs at the origin from which the hue angle is calculated. The second image shows the anti-saturation; here the discontinuity has disappeared, and we argue that the greyscale rendering is more similar in look and feel to the original colour spectrum.

(a) Original image

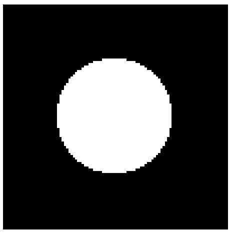

(b) Anti-saturation component of (a)

(c) Dilated image

Figure 4. *The same colour image as shown in Figure 1 both before (a) and after (c) dilation using a disk shaped structuring element and a lexicographic ordering scheme based on anti-saturation. The image in (b) shows the anti-saturation component of (a), and since both colours in (a) have equal lightness and saturation, this image is used to order the colours. When compared to Figure 1 it can be seen that the false colours are no longer present .*

In Figure 4 we show the same circular disk arrangement as shown in Figure 1. In Figure 4, however, the dilation operation is performed using a lexicographic ordering based upon the lightness, saturation and anti-saturation, rather than the marginal ordering as shown in Figure 1. We note that since the blue and green colours both have equal hue and saturation components, the ordering is done solely based upon anti-saturation value. In this example, the use of a lexicographic ordering has removed the problem of false-colours.

Experimental

Given that most images contain a large amount of luminance and saturation variation, these components will naturally have more weight in a lexicographic ordering procedure that uses luminance (or a luminance correlate) and saturation as the primary components. Thus, to test our method directly, we apply it to images that are both equi-luminant and equi-saturation. To generate such images we take a standard *RGB* image and firstly force it to be equi-luminant by dividing each *RGB* vector by its sum, i.e. $R \rightarrow \frac{R}{R+G+B}, G \rightarrow \frac{G}{R+G+B}$ and $B \rightarrow \frac{B}{R+G+B}$. We then force the image to be equi-saturation by mapping each *RGB* vector, along a line of constant hue, to a circle in the plane $R+G+B=1$ whose centre is the point $\left[\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right]$; i.e. a circle in the equi-luminant plane, all points on which have equal saturation. Since lines of equal hue are not defined for achromatic colours, such colours were mapped arbitrarily to a blue colour. The result of these operations is shown in Figure 5 where the original image in (a) is shown with the its equi-luminant and equi-saturation counterpart in (b).

Figure 5 shows representations of (b) in terms of their anti-saturation, (c), and hue-angle, (d), values. If anti-saturation is a more natural ordering scheme than hue-angle, we expect the antisaturation value in (c) to show a better representation than the hue-angle image in (d). In (d) it is clear that the second hat from the right contains hues that fall either side of the hue-angle discontinuity, thus different parts of the same hat are assigned opposite hue-angle values (shown as black and white in this representation). This particular problem is reduced in the anti-saturation representation, where the same hat does not contain polar opposite values. Also, the second hat from the left has a very similar hue to the second hat from the right in (b), but is made black in the hue-angle image reproduction; this again shows the problem of the discontinuity in the hue angle. We point out, however that the anti-saturation image is relatively noisy when compared to the hue-angle image. Furthermore, while it may better represent the content of the original image in (a) in terms of image detail (such as the writing on the hats and the grain on the wood), it seems somewhat too detailed to be an accurate representation of (b), the image from which it is derived, where these features are not apparent. In this respect, the hue-angle image in (d) gives a better representation.

For the same image we can evaluate the effect of morphological operators using a lexicographic ordering. Firstly we define the four principal operations of morphology: these are dilation, erosion, opening and closure; which are commonly represented using the symbols $\oplus$, $\ominus$, $\circ$, and $\bullet$ respectively. For a greyscale image, given by the function $f(x,y)$, and a structuring element given by a set h, the four principal operations are given by:

$$(f \oplus h)(x,y) = \sup_{(r,s)\in H} \{f(x-r, y-s) + h(r,s)\} \tag{5}$$

$$(f \ominus h)(x,y) = \inf_{(r,s)\in H} \{f(x-r, y-s) - h(r,s)\} \tag{6}$$

$$f \circ h = (f \ominus h) \oplus h \tag{7}$$

$$f \bullet h = (f \oplus h) \ominus h \tag{8}$$

where inf denotes the *infimum*, or greatest lower bound, of the set, and sup denotes the supremem, or least upper bound.

In Figure 6 we show the effect of a closing operation on the hats image from Figure 5 (b), using a disk-shaped structuring element. This closing operation is performed on the hue component of the image, and is shown for both anti-saturation in 6 (a) and hue-angle in 6 (b). A closing operation is defined as a dilation followed by an erosion, and is intended to close small gaps, leaving large structures in-tact and removing noise. A direct comparison of Figures 6 (a) and (b) shows that, while the anti-saturation ordering achieves the desired goal to some degree, the final image is relatively noisy, and the shape of the hats is better maintained by the hue-angle based ordering.

We believe that the principal problem with the proposed method can be seen in Figure 7. Here, the anti-saturation and hue-angle values (both scaled from 0 to 1) are plotted for an equi-luminant, equi-saturation hue circle (like that shown in Figure 3). Here it is clear that, while the hue-angle (dashed red line) has a single discontinuity, it is is generally smooth (i.e. has a small gradient) away from this discontinuity. The anti-saturation, on the other hand, is a continuous function, but is relatively less smooth.

(a) Original image

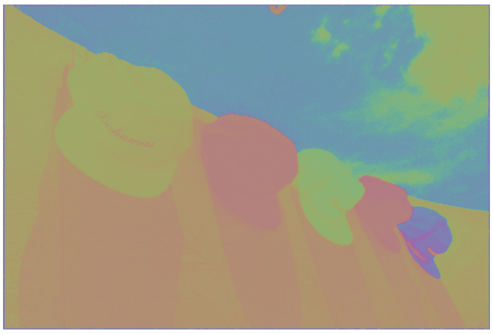

(b) Equi-luminant and equi-saturation image

(c) Anti-saturation value

(d) Hue angle value

Figure 5. *An image of some hats (a), and an equi-luminant equi-saturation version of the same image (b) (see text for details of the transform). The anti-saturation value at each pixel is then shown in (c), while the hue-angle value for each pixel is shown in (d).*

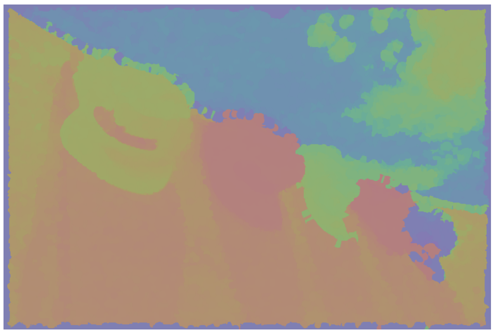

(a) Ordering using anti-saturation

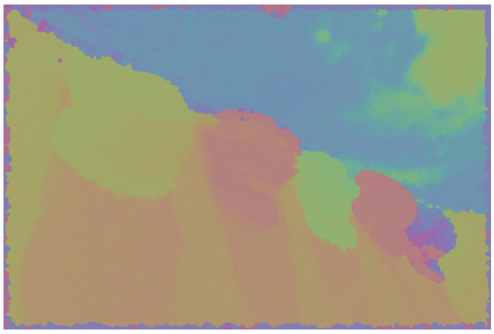

(b) Ordering using the hue angle

Figure 6. *An equi-saturation, equi-luminant image of some hats (see Figure 5 b) for the original), subjected to a morphological closing operation, (a) using anti-saturation to perform the ordering, and (b) using the hue-angle.*

This means that a small change in hue (i.e. a shift along the x-axis of Figure 7) results in a relatively large shift in anti-saturation value compared to the hue-angle; this is true for most hues except, of course, for those close to the discontinuity in hue-angle. We believe that this lack of smoothness accounts for the increased noise seen in Figure 5 (c) and 6 (a).

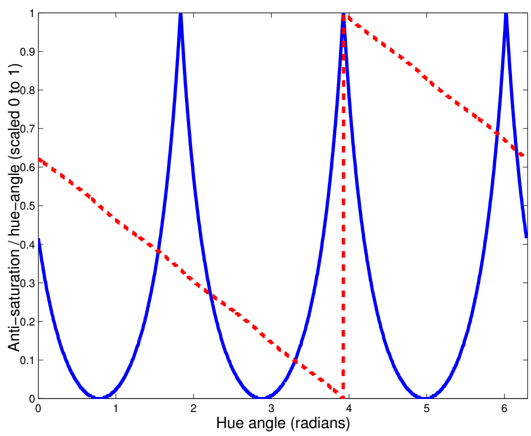

Figure 7. *A 1D representation of hue-angle (red dashed line) and anti-saturation value (blue solid line) for different points on an equi-saturation, equi-luminant hue circle; this Figure is effectively a cross-section of Figures 3 (b) and (c) respectively.*

Conclusions

We have proposed a novel approach to hue ordering, which we term anti-saturation, and investigated its application to image representation and to ordering for morphological colour operators. The proposed method is both continuous, and gives a natural representation of a hue-gradient. Practically, however, it can suffer from an overall lack of smoothness when compared to the hue-angle. Furthermore, preliminary results in a psychophysical investigation into hue ordering suggest that an ordering different from anti-saturation may provide a better model of preferred hue-ordering in human observers.

Acknowledgements

This work is supported by joint EPSRC grants number EP/E012248/1 and EP/E12159/1.

References

[1] V Barnett. The ordering of multivariate data. *Journal of the Royal Statistical Association - A*, **139**(3), 318–354 (1976).

[2] M. Comer and E. Delp. Morphological operations for color image processing (1999).

[3] C. L. Novak and S.A. Shafer. Color edge detection. In *Proceedings DARPA Image Understanding Workshop, Vol. 1.*, pp. 35–37, Los Angeles, CA, USA (1987).

[4] Allan G. Hanbury and Jean Serra. Morphological operators on the unit circle. *IEEE Transactions on Image Processing*, **10**(12), 1842–1850 (2001).

[5] Allan G. Hanbury and Jean Serra. Mathematical morphology in the hls colour space. In *Proceedings of the 12 British Machine Vision Conference* (2001).

[6] Allan G. Hanbury and Jean Serra. Mathematical morphology in the cielab space. *Image Anal Stereol*, **21**, 201–206 (2002).

Biography

David Connah received a BSc in Biology (1997), and an MSc in Machine Perception and Neurocomputing (1998), both from the University of Keele. He studied for his PhD at the Colour and Imaging Institute at the University of Derby before working as a Post-Doc researcher in the Norwegian Colour Research Lab in Gjovik. He is currently a senior research associate at the University of East Anglia, where he works in Prof. Graham Finlayson's Colour Group.