On the Morphological Processing of Hue

Erchan Aptoula, Sébastien Lefèvre

On the morphological processing of hue

E. Aptoula, S. Lefèvre *

LSIIT UMR-7005 CNRS-ULP, Pôle API, Blvd Sébastien Brant, PO Box 10413, 67412 Illkirch Cedex, France

Abstract

Although polar colour spaces are being increasingly used in the context of colour mathematical morphology, mainly due to their intuitiveness, the processing of the circular hue band continues to be their main drawback. In this paper, we discuss the two principal problems concerning the morphological processing of hue, first its lack of a lattice structure, which we propose to introduce by means of a distance based formulation from multiple representative reference hue points. And second, the distinction of chromatic from achromatic pixels, since the hue component is of no significance for "low" saturation levels. For this purpose we present a new weighting scheme, based on the combination of saturation and luminance channels. Application results on texture classification, asserting the superior performance of this approach, are also included.

Key words: Colour mathematical morphology, hue ordering, hue weighting, polar colour spaces, chromaticity

1 Introduction

The extension of mathematical morphology (MM) to colour, and more generally to multivariate images is an open problem. Since the formalisation of the underlying theory [7,17,18], and the acceptance of complete lattices as the appropriate algebraic basis for morphological operators, several approaches have been proposed to this end, a comprehensive survey of which can be found in Ref. [4]. Despite the absence of a widely accepted solution, the field of colour morphology has matured since the 1990s, and certain trends have started to form. More precisely, the use of polar (or phenomenal) colour spaces has been steadily spreading within the community [2,3,10,14,20].

* Corresponding author. Tel: +33(0)3 90 24 45 70; fax: +33(0)3 90 24 44 55
Email address: lefevre@lsiit.u-strasbg.fr (S. Lefèvre).

However, their popularity, which is mainly due to the intuitiveness (from the point of view of human colour vision) of the luminance, hue and saturation based description of colour, is seriously shadowed by the problems related to the circular hue band. In particular, the hue, which represents the dominant wavelength of a colour, presents two major inconveniences as far as MM is concerned. First, as hues are angular values, they lack a natural lattice structure, in other words we cannot say for instance that "red is greater than blue". That is why, it has become frequent practice to order these angles with respect to their distances from a predefined reference value, assumed to represent the "dominant" wavelength of the image at hand. Nevertheless, an image can have obviously more than one dominant hue group. In this paper, we present a new way of introducing a lattice on the hue circle, and thus defining morphological operators, using multiple reference hue values that represent the given image in a more consistent way with respect to existing approaches.

Moreover, when processing the hue, it is imperative to be able to distinguish chromatic from achromatic pixels, since the hue value of the latter has no practical significance, and hence can disrupt eventually the end result. To this end, most authors have so far preferred to use the saturation levels as a measure of "colourfulness". Here we propose a sigmoid based model, making use of both saturation and luminance, as the "significance" of the first is directly related to the overall illumination level. The results of a texture classification test are also included, where the proposed sigmoid based model exhibits a superior performance compared against state of the art approaches.

The rest of the paper is organised as follows. Section 2 discusses initially the crucial choice of polar colour space. Then in section 3 the problems related to the ordering of hue values are elaborated and the proposed solution is detailed. Next, in section 4, the weighting of hues is studied with the purpose of distinguishing chromatic pixels, along with the results of comparative tests on texture classification. Finally section 5 is devoted to concluding remarks.

2 Choice of polar colour space

3D-polar coordinate colour spaces appeared as the result of attempts to describe the RGB cube in a more intuitive manner, from the point of view of human interpretation of colour, in terms of *luminance*, *saturation* and *hue*. While luminance $L \in [0, 1]$ accounts for the amount of light, saturation $S \in [0, 1]$ represents the purity of a colour. The values of the periodical hue interval $H \in [0, 2\pi]$ [1] (or $[0, 1]$ if they are normalised) on the other hand, denote the

[1] Note that for the sake of clarity normalised hue values have been used throughout this paper

dominant wavelength, with 0 corresponding to red.

Basically, polar colour spaces achieve this transformation by representing colours with respect to the achromatic axis of RGB. Nevertheless, several implementational variants are available for this single transformation, e. g. HSV, HSB, HLS, HSI, etc [6]. According to Hanbury and Serra [13], all of the aforementioned colour spaces have been developed primarily for easy numerical colour specification, while they are ill-suited for image analysis. Specifically, although they were initially designed as conic or bi-conic shaped spaces, later on their cylindrical versions were employed in practice, in order to avoid the computationally expensive (for that period) checking for valid colour coordinates. The passage from conic to cylindrical shape however resulted in many inconsistencies within these spaces, for instance by allowing fully saturated colours to be defined in zero luminance. Extensive details on this topic can be found in Ref. [13].

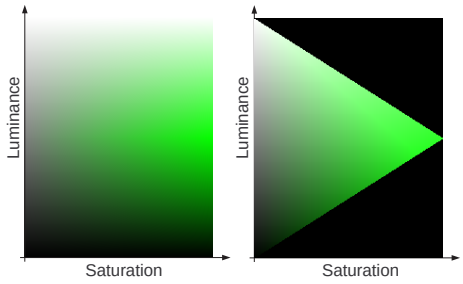

Fig. 1. Vertical semi-slice of the cylindrical HLS (left) and bi-conic IHLS (right) colour spaces.

Here we adopt the suggestion made in Ref. [13], and make our colour space choice in favour of the *improved HLS* space (IHLS), which employs the original bi-conic version of HLS, hence limiting the maximal allowed value for saturation in relation to luminance (Fig. 1). Further advantages of IHLS with respect to its counterparts include the independence of saturation from luminance, thus permitting the use of any luminance expression (e. g. RGB average, perceptual luminance, etc) and the comparability of saturation values. Although the IHLS colour space is chosen as the most appropriate representation of polar colour coordinates, and will be employed during the sequel of this paper, the hue related problems and solutions that are subsequently elaborated are by no means specific to IHLS, and can be used with other polar colour spaces with no or minimal modifications.

3 Ordering hues

In this section, we study the first problem concerning the application of morphological operators on the hue band of images, namely the absence of a lattice

structure.

In order to define morphological operators on an image $f : \mathcal{E} \to \mathcal{T}$, a complete lattice structure needs to be introduced to its intensity range $\mathcal{T}$. In other words, $\mathcal{T}$ must be a non empty set equipped with at least a partial ordering such that every non empty subset $\mathcal{P} \subseteq \mathcal{T}$ has a greatest lower bound $\bigwedge \mathcal{P}$ (*infimum*) and a least upper bound $\bigvee \mathcal{P}$ (*supremum*).

Consequently, as colour, and more generally multivalued images have vectorial intensities usually in $\overline{\mathbb{R}}^n$ or $\overline{\mathbb{Z}}^n$, with $n > 1$, given a vectorial ordering scheme "$<$", the vectorial versions of the two fundamental morphological operators, erosion ($\boldsymbol{\varepsilon}_b$) and dilation ($\boldsymbol{\delta}_b$) of a multivalued image $\mathbf{f}$, by a flat structuring element (SE) b, can be immediately derived by means of the vectorial extrema operators $\sup_v$ and $\inf_v$ based on the given ordering:

$$\boldsymbol{\varepsilon}_b(\mathbf{f})(\mathbf{x}) = \inf_{\substack{v \\ \mathbf{s} \in b}} \{\mathbf{f}(\mathbf{x} + \mathbf{s})\} \tag{1}$$

$$\boldsymbol{\delta}_b(\mathbf{f})(\mathbf{x}) = \sup_{\substack{v \\ \mathbf{s} \in b}} \{\mathbf{f}(\mathbf{x} - \mathbf{s})\} \tag{2}$$

Therefore, in order to benefit from the advantages of polar spaces, the ordering of IHLS colour vectors is necessary. Although luminance and saturation possess a ready to use lattice structure due to their natural scalar order, major problems arise as far as hue is concerned, since circular data lack such a structure.

Of course since hue values reside in $[0, 1]$, one can always order them using the scalar order, but there is no a priori reason for setting red ($h = 0$) as "less" than green ($h = 0.66$). The fixed origin and discontinuity at red, besides being an important flexibility constraint, also results in undesirable behaviours in this area of the hue circle. To illustrate this last point, Fig. 2 shows the dilation based on the scalar hue order, of a colour image containing all the possible hues. The abrupt discontinuity at red is what causes the edge in the middle of Fig. 2.

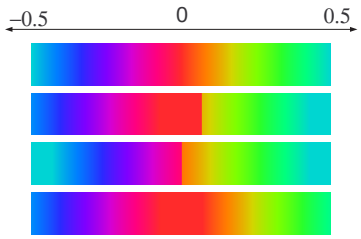

Fig. 2. From top to bottom, a colour image (500 × 70 pixels) containing fully saturated hues, its dilation with a 51 pixel wide horizontal line shaped SE based on the scalar order, the hue ordering (3) and the ordering (5), both with $h_0 = 0.0$.

Indeed, this situation has led some authors to even ignore this component as in Ref. [19], whereas the first attempt aiming to counter it was made by Peters [16], who employed a function $D(h_0, h)$ measuring the distance between the hue of the SE h_0 and the processed value h, with the aim of ordering them:

$$\forall\, h, h' \in [0,1],\ \ h \leq h' \Leftrightarrow D(h_0, h) \leq D(h_0, h') \tag{3}$$

where $D(\cdot,\cdot)$ is defined as:

$$\forall\, h, h' \in [0,1],\ \ D(h,h') = \begin{cases} h - h' + 1 \text{ if } & -1 \leq h - h' < -0.5, \\ h - h' \text{ if } & -0.5 \leq h - h' < 0.5, \\ h - h' - 1 \text{ if } & 0.5 \leq h - h' < 1. \end{cases} \tag{4}$$

Essentially, since hue values lack an inherent ordering relationship, any arbitrarily developed fixed order (e. g. *green* > *yellow* > *blue*, etc) is of limited practical use. The flexibility of a distance based approach, due to the use of an arbitrary hue origin, renders it adaptive to varying application needs, and that is why this type of hue orderings have met with acceptance. A variant has been later on proposed by Hanbury and Serra [10,11]. Specifically, according to them, Eq. (3) results in counter-intuitive operators, since for instance erosions would tend to enlarge objects of the reference h_0 (Fig. 2). Hence, they proposed using the ordering:

$$\forall\, h, h' \in [0,1],\ \ h \leq h' \Leftrightarrow h' \div h_0 \leq h \div h_0 \tag{5}$$

where, contrary to Eq. (3), hues closer to h_0 are considered greater and the angular distance $h \div h_0$ of h from h_0, is modified as:

$$\forall\, h, h_0 \in [0,1],\ \ h \div h_0 = \begin{cases} |h - h_0| & \text{if } \ |h - h_0| < 0.5 \\ 1 - |h - h_0| \text{ if } & |h - h_0| \geq 0.5 \end{cases} \tag{6}$$

Fig. 2 presents the result of a hue dilation based on the ordering of Eq. (5). As vectors are ordered with respect to their distance from $h_0 = 0.0$, the edge appearing in Fig. 2 is avoided within the red region. Furthermore, when combined with unsupervised reference hue computation methods [1,9], the ordering of Eq. (5) can be effectively used for general purpose morphological hue processing [11]. Let us note however that the term "ordering" is used loosely in the context of distance based reference hue computations, since they obviously do not verify the anti-symmetry constraint, as two distinct hue values can lie at equal distances to the reference hue; hence they constitute pre-orderings to be exact. One possible approach for satisfying the desired

anti-symmetry property, is to use an arbitrary ordering direction within the hue circle with the purpose of resolving such conflicts [12]. For instance if $h \div h_0 = h' \div h_0$ and $h \neq h'$, then the hue located first in the chosen ordering direction is considered smaller.

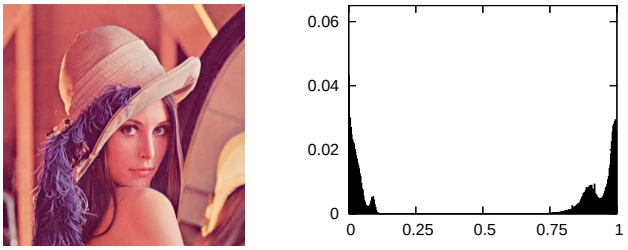

Fig. 3. The image Lenna (left, 512 × 512 pixels) and its normalised hue histogram (right).

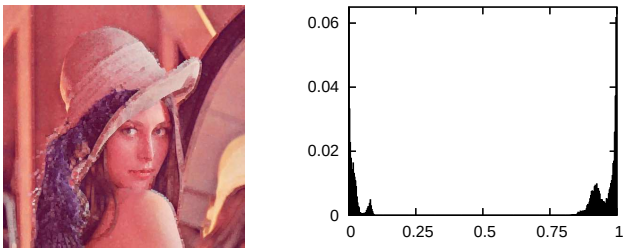

Fig. 4. The image Lenna (left), after the hue channel has been dilated with a square shaped SE of size 7 × 7 pixels using $h_0 = 0.0$ and ordering of Eq. (5), and the corresponding normalised hue histogram (right). Observe the concentration around h_0.

3.2 Multi-reference based hue ordering

Using a variable hue origin is certainly a practical way of introducing a complete lattice structure on the hue circle. However, as it will be subsequently detailed, this approach has certain limitations, the effects of which depend on the image at hand. Here we propose a more flexible new hue ordering based on distances from multiple origins. For the sake of simplicity, in the sequel all hue values are normalised to $[0, 1]$.

Let us consider the case of the Lenna image (Fig. 3 (left)). As it can be seen from its hue histogram (Fig. 3 (right)), there is a high concentration around the origin of the hue circle. And since reference hues are chosen usually as the average hue or the most frequent [1,9], it is considered pertinent to choose $h_0 = 0.0$. By accepting a certain wavelength within the red region of the spectrum as the reference hue, the nuances of red closer to this value increase their presence within the image by means of a dilation (Fig. 4). In other words "red > orange and red > magenta". But how about blue, and green or cyan?

Their relationship to the reference hue is *implicitly* assumed to be that of their position on the hue circle. Specifically, red is equidistant to blue and green, and furthest from cyan, based on the principle of colour complementarity.

This approach however assumes that the input image has at most a single dominant hue, which in the case of Lenna is red. Hence, in order to make the connection with binary morphology, the reference hue h_0 represents the "foreground" and its complement on the hue circle the "background". However, in the general case the presence of more than one not dominant hues is rather more frequent. Observe the image Cat (Fig. 5 (left)) for instance, where primarily three colours appear to be frequent: red/orange, green and blue, corresponding to the three large peaks in its hue histogram (Fig. 5 (right)). The reference hue that was obtained ($h_0 = 0.34$) through averaging, shows that green-yellow is the dominant wavelength, as it occupies a relatively important portion of the total image area. Consequently, during the dilation of the same image (Fig. 6), hues are ordered only based on their distance from green, even though orange and blue correspond to almost equally important regions/objects within the image.

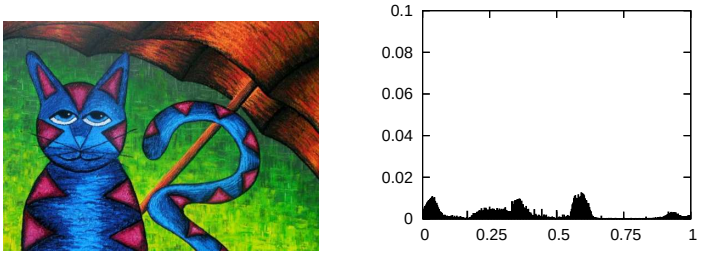

Fig. 5. The image Cat (by Gamze Aktan) (left, 576 × 420 pixels) and its normalised hue histogram.

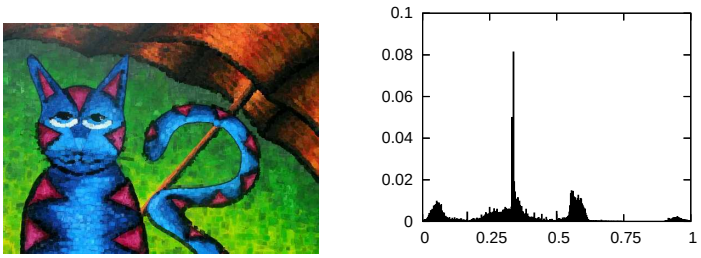

Fig. 6. The image Cat (left), after the hue channel has been dilated with a square shaped SE of size 7×7 pixels using $h_0 = 0.34$ and ordering of Eq. (5), and its corresponding hue histogram (right).

To counter this situation, one could eventually employ a region specific hue reference, for instance by computing h_0 independently for each hue region. And during ordering, the hues of a given region could be ordered according to their distances from the local reference hue. Nevertheless, this approach is hard to use in practice as it requires a pertinent segmentation step, that will provide the necessary region borders. Moreover, the image content may be of

such nature that regions may not even be distinguishable, and furthermore, how to order hues that belong to different regions? Alternatively, one can also adopt a pixel specific approach, where the reference hue is constantly updated according to the pixels under the SE. This use of a variable reference hue however, undermines several theoretical properties of the resulting operators, since "distant" pixels become no longer comparable.

The solution we propose here, is to use a hue ordering based on multiple references. Specifically, let us assume that we have obtained the k representative hue values $R = \{h_i\}_{1 \leq i \leq k}$. In the case of the image Cat (Fig. 5 (left)), the set R would contain roughly the hues red/orange, green/yellow, blue and magenta. Given two hue values, we can now order them based on their distances from each of the k references, and choose as greater the one closer to a representative hue value. More precisely:

$$R = \{h_i\}_{1 \leq i \leq k}\,, \;\; \forall\, h, h' \in [0,1], \;\; h \leq_H h' \Leftrightarrow \min_i \{h' \div h_i\} \leq \min_i \{h \div h_i\} \tag{7}$$

where the binary relation $\leq_H$ denotes the proposed ordering between hue values.

Consequently, if the set R contains a single reference hue, such as in the case of the image Lenna, then Eq. (7) becomes obviously equivalent to Eq. (5) (Fig. 7 (left)). However, if on the contrary R contains multiple references, then during ordering we favour those hues that are closer to one of the representative references (Fig. 7 (middle)), hence dividing the hue circle into regions with each having its own extremum. In the case of the Cat image, given two nuances of green, the greater would be the one which is closer to the reference hue situated in the green region of the hue circle, whereas if we compare a nuance of blue against one of green, the one closer to its respective reference would be considered greater, hence favouring "representative" hues, in terms of their distance from their respective reference hues. This can be easily observed by examining the hue histogram of the image Cat, obtained after having applied a dilation based on this principle (Fig. 8 (right)). The concentration of hue values around the four major hues is obvious.

Since it is relatively hard to detect the visual differences among the previous images, consider the artificial example shown in Fig. 9 (left). It consists of two uniform blue and yellow regions, containing noise spots of green and magenta. The "most dominant" hue is evidently blue, hence, a dilation using this single reference leads to the removal of spots only from the upper area (Fig. 9 (middle)), since both green and magenta are closer to blue than yellow. On the other hand, both regions become noise free using the proposed approach in combination with the two dominant hues, i. e. blue and yellow (Fig. 9 (right)).

Similarly, in the case of an erosion, non-representative hues, in other words

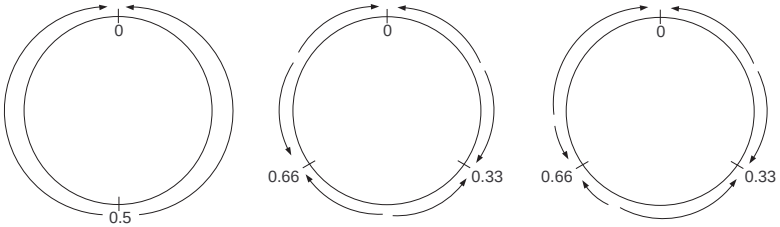

Fig. 7. From left to right, the hue circle with arrows showing the direction of increasing hues, in the case of a single reference hue ($h_0 = 0.0$), with multiple reference hues ($h_0 = 0.0$, $h_1 = 0.33$, $h_2 = 0.66$), and with multiple reference hues weighted according to arbitrary significances.

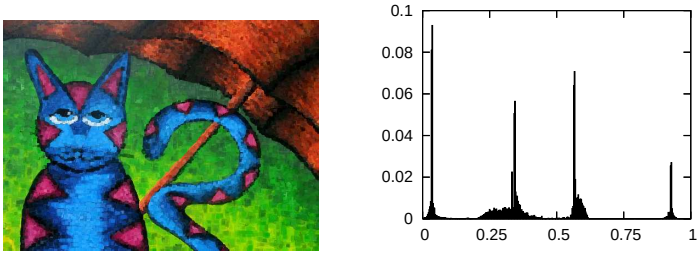

Fig. 8. The image Cat (left), after the hue channel has been dilated with a square shaped SE of size 7×7 pixels using ordering of Eq. (7), references $h_0 = 0.04, h_1 = 0.34, h_2 = 0.56, h_3 = 0.93$, and its corresponding normalised hue histogram (right).

hues that are distant from the reference points on the hue circle are favoured. Consequently, the approach based on a single reference erodes the spots in the yellow region, since both green and magenta are closer to yellow than blue (Fig. 10 (left)). A more consistent behaviour is achieved with the proposed ordering (Fig. 10 (right)).

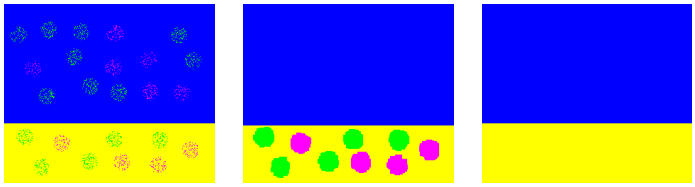

Fig. 9. From left to right, the original image (297×252 pixels), its hue dilation using ordering (5) and $h_0 = 0.66$, and the proposed ordering (7) with $h_0 = 0.66$, $h_1 = 0.33$, both with a square shaped SE of size 5×5 pixels.

Then again, likewise to Eq. (5), $\leq_H$ is too a pre-ordering, and unfortunately we can no longer use the arbitrary hue circle direction to resolve equivalences among distinct hues [12], as described in Sec. 3.1, since now we have multiple reference hues. That is why an alternative approach is required in order to satisfy the anti-symmetry property. One way of achieving this goal is to employ the order of importance among references, so that the hue closer to a "more important" reference hue is considered greater, and if that does not resolve

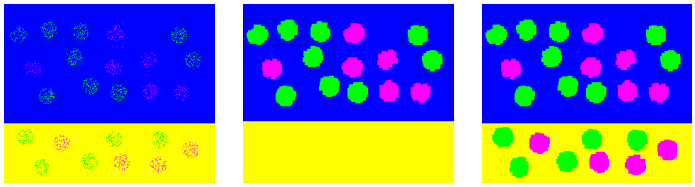

Fig. 10. From left to right, the original image, its hue erosion using ordering (5) and $h_0 = 0.66$, and the proposed ordering (7) with $h_0 = 0.66$, $h_1 = 0.33$, both with a square shaped SE of size 5×5 pixels.

the ambiguity, comparing the scalar hue values certainly does. Specifically:

$$h <_H h' \Leftrightarrow \begin{cases} \min_i \{h' \div h_i\} < \min_i \{h \div h_i\}, & \text{or} \\ \min_i \{h' \div h_i\} = \min_i \{h \div h_i\} \quad \text{and} \quad f(\text{ref}(h)) < f(\text{ref}(h')), & \text{or} \\ \min_i \{h' \div h_i\} = \min_i \{h \div h_i\} \quad \text{and} \quad f(\text{ref}(h)) = f(\text{ref}(h')) \quad \text{and} \quad h < h' \end{cases} \tag{8}$$

where $\text{ref}(h) = \arg\min_{h_i \in R} \{h \div h_i\}$ is the reference hue closest to a hue value h and $f(\cdot) : [0, 1] \rightarrow \mathbb{R}$ is the function associating to a reference hue its relative "importance" (detailed further later in this section).

3.3 Computation of reference hues

Of course, a necessary step for the implementation of the hue ordering in Eq. (7), consists in obtaining the reference hue(s). Moreover, unless an expert is available, it needs to be carried out automatically. Finding the "important" histogram maxima, can be considered as a histogram clustering or segmentation problem. As such, several options can be used with this purpose, for example any of the known unsupervised clustering techniques (e. g. meanshift), or even a morphological watershed transform. Through empirical means, the following method was observed to provide robust results:

(1) Calculate the normalised hue histogram of size n, where bin h_j is denoted by $P(h_j)$ and $P_\mu = \sum_j P(h_j)/n$ is the histogram average
(2) Threshold the histogram using its average, i.e. set to 0 all $P(h_j)$ such as $P(h_j) < P_\mu$
(3) Find the m connected histogram sections with non-zero bins, i.e. all pairs $h_{\alpha,\beta} = (h_\alpha, h_\beta)$ such as both conditions $P(h_j) > 0$, $\forall h_j \in [h_\alpha, h_\beta]$ and $P(h_{\alpha-1}) = P(h_{\beta+1}) = 0$ hold
(4) For each section $h_{\alpha,\beta}$, calculate its maximum, as reference hue:

$$h_{\alpha,\beta}^{\max} = \arg \max_{h_j \in [h_\alpha, h_\beta]} P(h_j) \tag{9}$$

However, using this reference computation scheme, one can eventually obtain several histogram sections, each with its own maximum, hence resulting in an excessively subdivided hue circle. To counter this eventual undesirable situation, one can of course first filter the histogram, using for instance an opening with a horizontal SE, but a more efficient way consists in weighting the distance from each reference hue, with the "importance" of its originating histogram section. To illustrate this idea, consider the histogram of the Cat image (Fig. 5), which with the proposed reference hue computation method, results in four sections, representing roughly the hues red/orange, green, blue and magenta, with the last being of relatively lesser significance. In order to quantify the "importance" of a section $h_{\alpha,\beta}$, one can use either its length $l_{\alpha,\beta}$ or preferably its discrete sum or volume $w_{\alpha,\beta}$ respectively defined by:

$$l_{\alpha,\beta} = h_\alpha \div h_\beta \tag{10}$$

$$w_{\alpha,\beta} = \sum_{h_j \in [h_\alpha, h_\beta]} P(h_j) \tag{11}$$

For the sake of clarity, we will denote by h_i the reference hue related to the i^{th} hue section ($1 \leq i \leq m$) and use the notations l_i and w_i instead of $l_{\alpha,\beta}$ and $w_{\alpha,\beta}$. Thus we can modify the ordering of Eq. (7) relying on a set of k reference hues among n existing hue sections. The new definition accommodates the significance n_i of each reference hue h_i depending on the volume w_i of its related section:

$$\begin{aligned} R = \{(h_i, n_i)\}_{1 \leq i \leq k}, \forall\, h, h' \in [0,1], \\ h \leq h' \Leftrightarrow \min_i \{(h' \div h_i) \times 1/n_i\} \leq \min_i \{(h \div h_i) \times 1/n_i\} \end{aligned} \tag{12}$$

where n_i is defined as:

$$n_i = \frac{w_i}{\sum_{j=1}^{m} w_j} \tag{13}$$

By means of the factor n_i, we ensure that hues that are moderately close to a significant reference hue, are considered larger than those much closer to a reference hue of lesser importance, hence providing increased robustness against noisy histograms (Fig. 7 (right)). Note that n_i may also be used as the reference importance function in Eq. (8). If we return to our example, this approach leads to the decrease of magenta hues after dilation (Fig. 11), since they are much less frequent than the other three major hues.

A further example demonstrating the effect of the proposed orderings, as well as the difference between Eqs. (7) and (12) is shown in Fig. 12. The original image (Fig. 12 (left)), contains a uniform hue gradient between a central cyan square and a larger surrounding blue area. Using the aforementioned reference computation scheme based on the hue histogram, we obtain as references the two values representing cyan and blue. For the sake of clarity, Fig. 13 shows

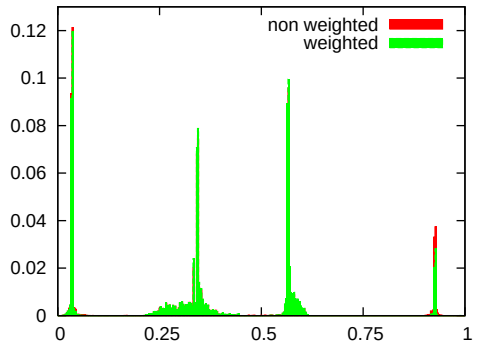

Fig. 11. The hue histograms of image Cat, after applying a dilation using orderings (7) and (12) along with references $h_0 = 0.04, h_1 = 0.34, h_2 = 0.56, h_3 = 0.93$, both with a square shaped SE of size 11×11 pixels.

additionally the uni-dimensional version of the hue gradient and the effect of the erosion and dilation operators, both with and without the use of weights. A relatively large SE has been used, in order to amplify the effect of each operator.

The dilation operator, applied based on Eq. (7) (Fig. 13 (left)), results in the discretisation of the gradient, by increasing each representative hue's influence range, hence creating the steep edge in the middle, showing the limit of each reference hue's attraction range (Fig. 7 middle). In other words, "representative" hues are favoured while the intermediate values are eliminated. Erosion on the other hand, favours the intermediate hue values between the two references, thus leading to the creation of a "middle" flat zone, and forcing the two gradient extremes edges to shrink.

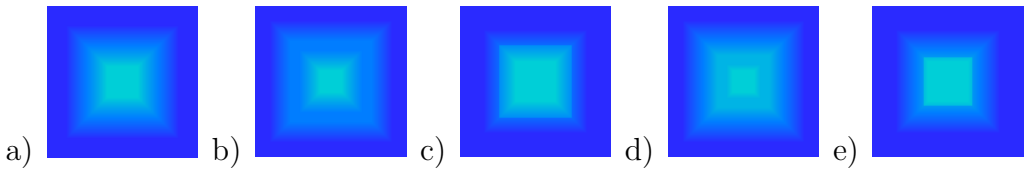

Fig. 12. From left to right, the original image (418×418 pixels), its hue erosion (b,d) and dilation (c,e) using orderings of Eqs. (7) and (12), both with references $h_0 = 0.66$, $h_1 = 0.5$ and their respective weights $n_1 = 0.8$ and $n_2 = 0.2$ along with a square shaped SE of size 25×25 pixels.

By taking into account the much smaller size of the central cyan area with respect to the outer blue, we obtain the results depicted in Fig. 13 (right). While dilation and erosion behave identically as in the unweighted case, the influence range of each reference hue is radically modified in favour of the obviously more dominant blue. Consequently, given a noisy hue channel which could lead to the eventual unsupervised detection of unimportant hue references, the weights incorporated within the ordering of Eq. (12) can decrease substantially their influence on the end result.

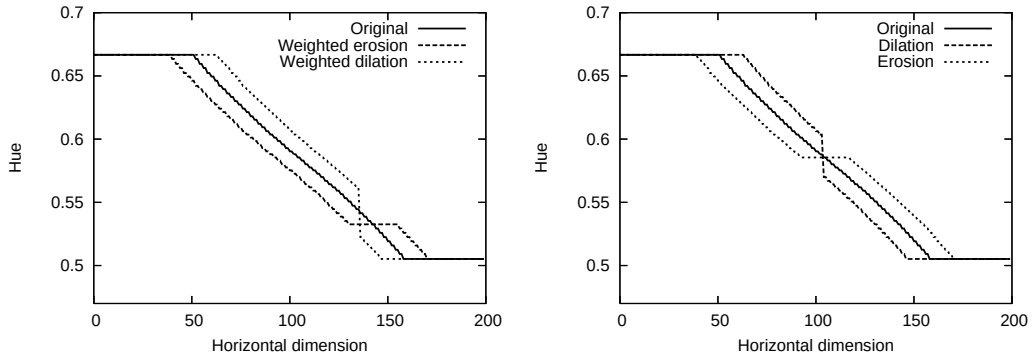

Fig. 13. The uni-dimensional gradient of Fig. 12, along with its dilated and eroded versions with (right) and without (left) using weights.

Prior to any practical use however, one last point that remains is the disturbance caused within the histogram by hue values corresponding to achromatic pixels. This problem may be resolved by using a function making it possible to distinguish pixels of chromatic value from those that are achromatic. Of course this is no deterministic procedure, and a continuous model is necessary. We choose to employ a trivial linear model, where the value of saturation provides a direct indication on the importance of hue. Specifically, given a coefficient $c(x, y) \in [0, 1]$ (i. e. the saturation at x, y) indicating the colourfulness level of a colour pixel at position x, y, the histogram can be constructed by incrementing the corresponding bin by c instead of 1. In other words, the total value in bin h_i of the hue histogram, is calculated as:

$$P(h_i) = \sum_{x,y} c(x, y)\delta_{h_i h_{x,y}} \tag{14}$$

where the sum is over all the pixels at x, y, $h_{x,y}$ is the hue value at this position and δ_{ij} denotes the Kronecker delta function.

4 Weighting of hue

In this section, we review the distinction of chromatic from achromatic pixels and propose a solution subsequently tested with a texture classification application against state of the art approaches.

Although no predefined order among wavelengths exists, by shifting the origin of the hue circle with arbitrary reference hues and using distance measurements, one can effectively introduce a complete lattice structure, thus resolving the first of the two problems linked with morphological hue processing. The second is not limited with the domain of MM, and concerns the distinction of chromatic from achromatic pixels. In particular, since the significance of the notion of colour depends directly on its "colourfulness" or saturation level,

it becomes necessary to be able to detect achromatic pixels, as they do not carry any information about colour, and hence need to be ignored during the processing of the hue band. Of course, since a simple discrete model based on a threshold is obviously impractical, a continuous relationship model is required, in order to determine the "importance" of each pixel's hue.

Multiple models have been reported in the literature for describing the relationship between the importance of hue and its saturation, using some sort of transition function. After briefly presenting some of these models, here we introduce a sigmoid based transition function based on both saturation as well as luminance. Although the IHLS space is assumed for all expressions, they can be trivially modified to accommodate other conic polar colour spaces as well.

4.1 Saturation based weighting models

One way of favouring saturated pixels during hue extrema computation has been proposed by Hanbury and Serra [8], and consists in replacing the original hue values with their weighted versions as follows:

$$h_i^* = \begin{cases} \sup[h_i, 0.25(1 - s_i)] \text{ if } & 0 \leq h_i \leq 0.25, \\ \inf[h_i, 0.25(1 + s_i)] \text{ if } & 0.25 < h_i \leq 0.5, \\ \sup[h_i, 0.25(3 - s_i)] \text{ if } & 0.5 < h_i \leq 0.75, \\ \inf[h_i, 0.25(3 + s_i)] \text{ if } & 0.75 < h_i < 1. \end{cases} \tag{15}$$

where hues $h_i \in [0, 1[$, saturations $s_i \in [0, 1]$; and for an arbitrary reference hue h_0 in the previous expression, h_i must be replaced with:

$$h_i \rightarrow \begin{cases} h_i - h_0 & \text{if } \quad h_i - h_0 \geq 0 \\ 1 + h_i - h_0 & \text{if } \quad h_i - h_0 < 0 \end{cases} \tag{16}$$

This weighting scheme works so that highly saturated pixels keep their initial hue, while pixels with lower saturation values have their hues moved to positions closer to $h_0 + 0.25$ and $h_0 - 0.25$, in order to reduce their probability of being chosen as extrema.

Another, more general-purpose model has been proposed by Carron in [5]:

$$s \in [0, 1], \;\; \gamma(s) = \frac{1}{\pi} \cdot \left(\frac{\pi}{2} + \arctan(\beta(s - s_0)) \right) \tag{17}$$

where s denotes saturation, and β and s_0 are fine tuning parameters of the resulting sigmoid. More precisely, their suggested values are $\beta = 0.7$ and $s_0 =$

0.196. Consequently, this transition function provides a "smoother" alternative to carrying out a discrete discrimination by means of a threshold.

However, the drawback of these aforementioned approaches is that they assume saturation to be an absolute measure of "colourfulness", whereas from the point of view of human colour vision, the perception of colour and hence the importance of hue, is strongly related to the luminance levels. This can be trivially observed in the cylindrical HLS colour space (Fig. 1 (left)), where extreme saturation values lose totally their importance at low and high luminance levels. And that is one of the reasons why the widely employed HLS colour space is inadequate for quantitative colour image processing. An extensive study on this topic can be found in [13], where the bi-conic IHLS space is proposed in order to counter this problem. Specifically, within the IHLS colour space, as it can be observed from Fig. 1 (right), the maximum allowed saturation value (S_{max}) depends linearly on the luminance (L) level: $S_{max} = 2\,(0.5 - |L - 0.5|)$.

Consequently, in order to avoid the pitfall of ignoring luminance, one needs to employ either an adequate colour space such as IHLS or take luminance into account during the weighting of hue. Furthermore, even if the IHLS space is used, one can still wonder if the linear dependence between S_{max} and L is optimal for the task under consideration.

4.2 Proposed model

In the light of the previous remarks, we propose a combined weighting model for hue, which takes into account both saturation and luminance by means of a double sigmoid function. As far as saturation is concerned, we use the following sigmoid:

$$s \in [0,1], \;\; g(s) = \frac{1}{1 + exp(-k(s - s_0))} \tag{18}$$

where the parameter s_0 is the offset controlling the point for which $g(s) = 0.5$, and k controls the slope. As shown by the plot of $g(s)$ in Fig. 14 (left), for large values of k, the sigmoid becomes equivalent to a threshold function at the offset s_0, which "brutally" designates the hues having saturations $s < s_0$ as achromatic, and they are totally ignored, whereas for $k = 1$, all hues are treated almost linearly with respect to their saturation, thus attributing a relatively high significance to hues of lower saturation levels. Similarly to Carron [5], we find it more pertinent to employ an intermediate form with $k = 10$, which results in smoother transitions, between the chromatic and achromatic classes. As far as the offset is concerned, it was empirically set as $s_0 = 0.5$.

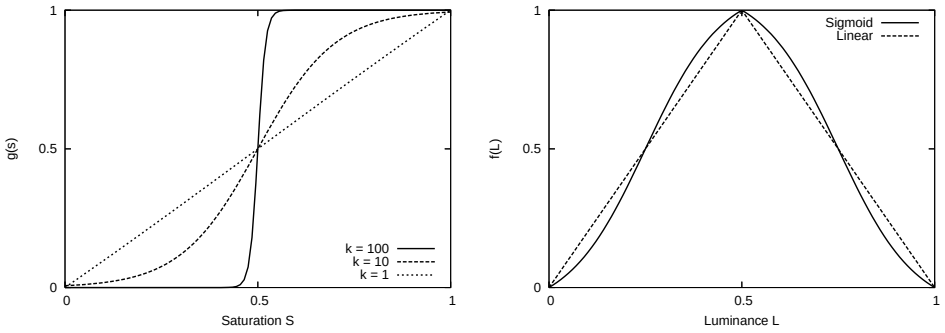

Fig. 14. Plots of weighting functions for the importance of colour in relation to saturation (left) and in relation to luminance (right).

As to luminance, in an attempt to further benefit from the smooth transition of the sigmoid function, we use it with the aim of detecting the luminance levels for which saturation is important. Specifically, we use the following form with this purpose:

$$l \in [0,1], \;\; f(l) = \begin{cases} \frac{1}{1+exp(-k_L(l-l_l))} & \text{if} \quad l \leq 0.5 \\ \frac{1}{1+exp(k_L(l-l_u))} & \text{if} \quad l \geq 0.5 \end{cases} \tag{19}$$

where the slope $k_L = 10$, and the lower and upper offsets are respectively $l_l = 0.25$ and $l_u = 0.75$. Its plot is given in Fig. 14 (right). The arguments of $f(l)$ were set empirically, and divide the luminance range roughly in three regions, with the middle corresponding to important saturation levels.

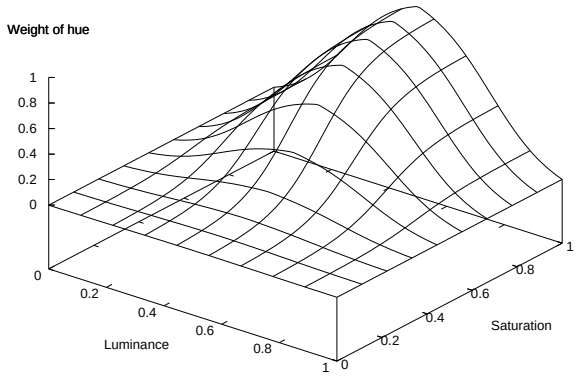

Fig. 15. The proposed combined luminance and saturation weight for hue.

The final weight α of hue is derived based on both saturation and luminance, as:

$$l, s \in [0,1], \;\; \alpha = f(l) \times g(s) \tag{20}$$

Fig. 15 shows its general allure. Consequently, hue values are considered as "important" or as corresponding to chromatic pixels, only when both satu-

ration is relatively high and luminance is at medium levels, hence making it possible for saturation to reach its maximal range. In order to show the practical interest of the proposed weighting scheme, application results follow.

4.3 Application results

In this section we present the results that have been obtained from a texture classification test, employing the aforementioned hue weighting methodologies.

The Outex13 colour texture database has been used with this purpose. It consists of 68 types of textures (Fig. 16) [15], with 20 examples from each, making a total of 1360 images acquired under identical illumination conditions. The test and training sets contain both 680 images while the classification is realised by means of a kNN classifier employing the closest neighbour ($k = 1$), computed with the Euclidean distance. As to the hue weighting models, we compare the proposed double sigmoid of equation (20) (COMBINED), against the absence of any weighting mechanism (NOWEIGHT), the scheme presented in equations (15) and (16) (HANBR), the direct use of saturation (LINEAR), the single sigmoid of saturation (18) (SIGMOID) as well as the model of equation (17) (CARR), with their parameters tuned to their aforementioned values. The textures of Outex13 are described by a single 360-dimensional feature vector consisting of their hue histogram. The histogram is constructed using equation (14), where $c(x, y)$ is replaced with the corresponding value provided by the weighting model in use.

Fig. 16. Examples of the 68 textures of Outex [15].

The classification accuracies that have been obtained are shown in table 1. In general, the performances are relatively close, with the exception of HANBR which deviates largely from the average. The most probable explication would be the fact that this weighting scheme has been designed specifically for extrema computation, hence is not suitable for general purpose distinction of

Table 1
Classification accuracies (%) for outex13, using hue histograms calculated with different weight models.

Method	NOWEIGHT	HANBR	LINEAR	CARR	SIGMOID	COMBINED
Accuracy	82.35	70.59	83.53	83.62	83.82	**84.85**

chromatic pixels. Moreover, as expected the lack of any weighting mechanism is the second worst among them, while even the linear use of saturation is sufficient for improving this result. The sigmoid based models however have outperformed their linear counterpart, even though with an almost negligible difference. The additional participation of luminance, along with saturation, in determining the importance of hues has had a minimal but positive effect on the results.

5 Conclusion

The two main problems concerning the processing of colour images' hue band have been studied and solutions have been presented. The lack of a lattice structure within the hue circle as well as the absence of an inherent order among colours has made dynamic reference based hue orderings adaptive and efficient solutions in this regard. However, they make use of a single reference value, which is implicitly assumed to represent their input.

We have presented a more general solution, employing multiple references, capable of taking into account images with arbitrary hue distributions. The resulting morphological operators, were shown to be equivalent to their mono-reference based versions, in the case of images dominated by a single hue group, whereas with multiple hue groups, a more frequent case in practice, the proposed approach has exhibited a more consistent performance, as asserted by application examples.

Furthermore, another hue related problem, the distinction of chromatic pixels from achromatic ones has also been examined and a new sigmoid based weighting model has been presented. Contrarily to previous approaches that make exclusive use of the saturation component, it was chosen to exploit additionally the luminance component, since the perception of colourfulness is directly related to the luminance levels. This double sigmoid was compared against multiple alternative approaches in the context of texture classification, were it has achieved the overall best accuracy rate.

References

[1] J. Angulo, Morphologie Mathématique et Indexation d'Images Couleur. Application à la Microscopie en Biomédicine, Ph.D. thesis, Ecole des Mines, Paris, France (2003).

[2] J. Angulo, Unified morphological color processing framework in a lum/sat/hue representation, in: C. Ronse, L. Najman, E. Decencière (eds.), Mathematical Morphology: 40 Years On, vol. 30 of Computational Imaging and Vision, Springer-Verlag, Dordrecht, 2005, pp. 387–396.

[3] J. Angulo, J. Serra, Morphological coding of color images by vector connected filters, in: IEEE Proceedings of the 7th International Symposium on Signal Processing and Its Applications (ISSPA'2003), vol. 1, Paris, France, 2003.

[4] E. Aptoula, S. Lefèvre, A comparative study on multivariate mathematical morphology, Pattern Recognition 40 (11) (2007) 2914–2929.

[5] T. Carron, Segmentations d'Images Couleur dans la Base Teinte-Luminance-Saturation: Approche Numérique et Symbolique, Ph.D. thesis, Université de Savoie, France (1995).

[6] R. C. Gonzalez, R. E. Woods, Digital Image Processing, 2nd ed., Addison-Wesley, New York, 1992.

[7] J. Goutsias, H. J. A. M. Heijmans, K. Sivakumar, Morphological operators for image sequences, Computer Vision and Image Understanding 62 (3) (1995) 326–346.

[8] A. Hanbury, Morphologie Mathématique sur le Cercle Unité: avec Applications aux Teintes et aux Textures Orientées, Ph.D. thesis, Ecole des Mines, Paris, France (2002).

[9] A. Hanbury, Circular statistics for colour images, in: 8th Computer Vision Winter Workshop, Valtice, Czech Republic, 2003.

[10] A. Hanbury, J. Serra, Mathematical morphology in the HLS colour space, in: T. Cootes, C. Taylor (eds.), 12th British Machine Vision Conference, Manchester, UK, 2001.

[11] A. Hanbury, J. Serra, Morphological operators on the unit circle, IEEE Transactions on Image Processing 10 (12) (2001) 1842–1850.

[12] A. Hanbury, J. Serra, Mathematical morphology in the CIELAB space, Image Analysis and Stereology 21 (3) (2002) 201–206.

[13] A. Hanbury, J. Serra, Colour image analysis in 3d-polar coordinates, in: Proceedings of DAGM, Magdeburg, Germany, 2003.

[14] G. Louverdis, M. Vardavoulia, I. Andreadis, P. Tsalides, A new approach to morphological color image processing, Pattern Recognition 35 (8) (2002) 1733–1741.

[15] T. Maenpaa, M. Pietikainen, Classification with color and texture: jointly or separately?, Pattern Recognition 37 (8) (2004) 1629–1640.

[16] A. Peters, Mathematical morphology for angle-valued images, in: Proceedings of SPIE, Nonlinear Image Processing VIII, vol. 3026, 1997.

[17] C. Ronse, Why mathematical morphology needs complete lattices, Signal Processing 21 (2) (1990) 129–154.

[18] J. Serra, Anamorphoses and function lattices, in: E. R. Dougherty (ed.), Mathematical Morphology in Image Processing, chap. 13, Marcel Dekker, New York, 1993, pp. 483–523.

[19] M. I. Vardavoulia, I. Andreadis, P. Tsalides, Vector ordering and morphological operations for colour image processing: fundamentals and applications, Pattern Analysis and Applications 5 (3) (2002) 271–287.

[20] M. Wheeler, M. A. Zmuda, Processing color and complex data using mathematical morphology, in: IEEE Proceedings of the National Aerospace and Electronics Conference, vol. 1, Dayton, USA, 2000.