Notes in English on

Mathematical Morphology on the Unit Circle

with applications to hues and to oriented textures

THESIS

presented to the
École Nationale Supérieure des Mines de Paris
by

Allan Hanbury

to obtain the title of

DOCTOR of MATHEMATICAL MORPHOLOGY

Preface

This document presents a skeletal outline of the corresponding Ph.D. thesis [7]. English documents are referenced where they exist, and slightly more detail is given where no corresponding document exists. At the moment, only chapter 1, the introduction, is fully translated. This document may be expanded in the fullness of time if demand is high enough.

Allan Hanbury
Fontainebleau
February 2002

Contents

1. Introduction

Data in the form of angles or two-dimensional orientations, called *circular data*, are often found in the analysis of the natural world, for example wind directions, directions of departure of birds or animals from a point of release, or orientations of fracture planes in rocks. The statistical analysis of circular data is a well researched subject [5, 14], but even though circular data are also found in the context of image analysis, the development of methods for processing them correctly has received less attention. For colour images, the hue component of cylindrical coordinate representations of colour spaces is an angular value. For this reason, this component has properties which are completely different to those of the two other components, saturation and luminance. Nevertheless, we often ignore these differences and apply the same algorithms to the three components. Two-dimensional direction fields are also found in image analysis, for example in the analysis of oriented textures and of vector fields which describe movement between adjacent frames of a sequence of images. A Fourier transform of any image also produces a vector field. The final example given here is that of a spectrogram, the result of a series of windowed Fourier transforms on a one-dimensional signal, which can also be visualised as a vector field. With the results of a Fourier transform, we have a tendency to apply filters only to the magnitudes of the vectors, leaving their directions (phases) aside. Why this reticence to touch the angular values? Why do we ignore the angular nature of certain data by treating them as linear data[1]?

A simple way of visualising circular data are as points on the unit circle, the circle of unit radius.With this representation, we can immediately see the characteristics which make these data so difficult to process. They are cyclic — adding 2π or one of its multiples to a circular coordinate always results in the same direction. Furthermore, an obvious origin does not exist — any position on the unit circle is equal to any other. This property was very useful to King Arthur when he chose a round table for his knights, but it is rather limiting in the context of mathematical morphology, as one cannot impose an order of magnitude on values on the unit circle. In this thesis, possible solutions to problems posed by these properties in the context of mathematical morphology are discussed.

The notion of rotational invariance is important when one is working with circular data. In a set of directions, the numerical value of the coordinate of each direction depends on the position of the origin. If an operator acting on this set always returns the same direction independently of the position of the origin (note that this is not necessarily the same numerical value), the operator

[1] The practitioners of image analysis are not the only people to err in this direction. At one time, the statistics of wind direction data were calculated using standard linear statistics, requiring the later development of methods to correct these errors.

is said to be *rotationally invariant*. For data on the unit circle, for which an obvious origin does not exist, this property is desirable. We therefore indicate which of the morphological operators which we develop satisfy this property.

We continue this introduction with a summary of the structure of the document (section 1.1), followed by a brief introduction to mathematical morphology (section 1.2).

1.1. Structure of the thesis

The thesis is divided into two parts. In the first part, we consider the treatment of circular data in the context of image processing. We begin by developing, in chapter 2, morphological operators suited for use on circular data. The first applications of these operators are presented in chapter 3, in the context of the analysis of oriented textures, a domain in which it is often possible to treat the circular data independently. We then consider the case where an angular coordinate is a component of a vector, the case of colour images represented in a cylindrical coordinate system. To try and simplify the choice of a colour vector representation from among the plethora of available colour spaces of this type, we present, in chapter 4, the development of a cylindrical coordinate colour space well adapted for use in image processing and analysis. The application of vectorial morphological operators in this space, as well as in the L*a*b* space, is presented in chapter 5.

In the second part, we present the results of a study of a current industrial problem: the automatic matching of planks of wood to make panels which are visually æsthetic. This study was done in the framework of the European project SCANMATCH. After an introduction (chapter 7), we present image processing algorithms which extract features from an image of a plank of wood which are useful for finding a good match. The colour features are treated in chapter 8, and the texture features in chapter 9. The colour of a plank is represented in a system of cylindrical coordinates of the type discussed in chapter 4. Wood texture is a classic example of an oriented texture, which implies that the algorithms developed in the first part of the thesis are directly applicable. The main difference between the approaches to the processing of the texture in the two parts of the thesis is that in the first part, we are concerned with defect detection in textures, whereas in the second part, we define features which allow the automatic classification of textures. The algorithms of the second part are also less experimental, having already been tested in practice in a prototype machine designed to classify wood according to colour and texture characteristics. The results of the application of these algorithms to databases of images captured under industrial conditions are also presented.

1.2. Mathematical morphology

Mathematical morphology [18, 19] comprises a set of non-linear operators which act on images by the use of structuring elements. A brief overview of the mathematical morphology operators used in this thesis is given below. We restrict ourselves to the specific case of an image $f : Z \rightarrow \overline{\mathbb{R}}$ defined on a sub-section Z of the two-dimensional Euclidean space $\mathbb{R}^2$ (the notation $\overline{\mathbb{R}}$

indicates the set $\mathbb{R} \cup \{-\infty, \infty\}$). In other words, the function assigns a grey-level to each point $x \in Z$ of the image. In practice, the Euclidean space is replaced by a discrete space $\mathbb{Z}^2$. For the case of digital images, the values taken by $f(x)$ are often limited to integers. For 8 bit greyscale images

$$f : Z \to \{0, 1, 2, \ldots, 255\}$$

and for binary images

$$f : Z \to \{0, 1\}$$

For colour images, the function f associates a vector with each point x. The morphological treatment of this type of image is discussed in chapter 5.

1.2.1. Erosion and dilation

These operators make use of a *structuring element* B, which is a set of points with a specified origin. The notation B_x indicates a structuring element which has been translated so that its origin is at position x in Z. An erosion of f by structuring element B is defined as

$$\varepsilon_B f(x) = \{\inf [f(y)], y \in B_x\}$$

applied to every point $x \in Z$, where $\inf$ is the infimum operator. The dilation of f by B is defined as

$$\delta_B f(x) = \{\sup [f(y)], y \in B_y\}$$

applied to every point $x \in Z$, where $\sup$ is the supremum operator. An alternative notation for the erosion of f by B is $f \ominus B$, and the dilation by B is $f \oplus B$.

As the square structuring element is used extensively, the notation S_i will be used to indicate a square structuring element of size i, which has sides of length $2i + 1$. So S_1 is a 3×3 square, S_2 is a 5×5 square, etc.

1.2.2. Opening and closing

An *opening* of f by B is defined as

$$\gamma_B f = (f \ominus B) \oplus B$$

and a *closing* of f by B is defined as

$$\varphi_B f = (f \oplus B) \ominus B$$

These two operators are increasing and idempotent. Furthermore, the opening is anti-extensive and the closing is extensive.

1.2.3. Top-hats

The top-hats (usually used only on greyscale images) are used to locate structures in the image which are smaller than the structuring element. Two top-hats are defined, the *white top-hat* is the difference between f and its opening

$$WTH_B(f) = f - \gamma_B f$$

and the *black top-hat* is the difference between the closing of f and f

$$BTH_B(f) = \varphi_B f - f$$

1.2.4. Reconstruction

Morphological reconstruction is based on *geodesic dilation*. A geodesic dilation involves two images: a marker image f and a mask image g, both defined on the same domain Z. The geodesic dilation of f with respect to the mask image g is defined as

$$f \oplus_g B = (f \oplus B) \wedge g$$

or the point-wise minimum between the dilation of f and g.

If one iterates a geodesic dilation on a bounded image, a point will be reached after which further changes in the marker image are impeded by the mask image. The geodesic reconstruction R_g is based on this principle. It is defined as the iteration until convergence of geodesic dilations of f with respect to g by the elementary structuring element (S_1). If we use the notation $\delta_g^{(n)}(f)$ to indicate n iterations of a geodesic dilation of f with respect to g using an elementary structuring element, then the geodesic reconstruction of f with respect to g is

$$R_g(f) = \delta_g^{(m)}(f)$$

such that $\delta_g^{(m)}(f) = \delta_g^{(m+1)}(f)$. For binary images, this has the effect of completely reconstructing the connected components in the mask image g which have at least one non-zero pixel in common with the marker image f. This operation is therefore an opening on the mask image g. When the marker image contains only one pixel at point x, the reconstruction is the connected component of g which contains point x. This operation is called the *connected point opening* and is written as $\gamma_x(g)$.

A common use of the reconstruction operator is in the opening by reconstruction, defined as

$$\gamma_B^R f = R_f(f \ominus B)$$

which, for a binary image has the effect of keeping only those connected components in f which do not disappear when eroded by B.

1.2.5. Area opening

In a binary image, an *area opening* γ_λ removes all the connected components with a surface area (pixel count) less than a certain threshold λ. It can be shown that this is equivalent to the union of all openings with connected structuring elements whose surface areas are equal to λ

$$\gamma_\lambda = \bigvee_i \{\gamma_{B_i} | \, B_i \; is \; connected \; and \; Area\,(B_i) = \lambda\}$$

but formulations more convenient for use in digital algorithms exist.

Part I.

The unit circle

2. The unit circle

2.1. Circular data and the unit circle

Some definitions of operations such as addition, subtraction and difference between values on the unit circle are given.

2.2. Circular statistics

The circular statistical measures of circular mean μ, mean length $\overline{R}$ and circular variance V, taken from [5], are presented.

2.3. Mathematical morphology on the unit circle

A brief introduction. Most of the material presented in the following four sections is from the article [13].

2.4. Mathematical morphology with the choice of an origin

This section is based on [6, chapter 4].

2.5. Pseudo-dilation and pseudo-erosion

This section is based on [13, section III].

2.6. Circular centred morphology

The circular centred gradient and top-hat operators are discussed in [13, section II].

2.7. Indexed partitions

The labelled openings (sections 2.7.5 and 2.7.6) are discussed in [13, section IV].

2.8. Conclusion

3. Defect detection in an oriented texture

3.1. Principal properties of textures

The four texture classes defined by Rao [16] are presented with examples. We then discuss the perceptual properties of textures: anisotropy, orientation, etc. [3].

3.2. Defects in textures

Various definitions of defects in textures are discussed.

3.3. Oriented texture

We present the algorithm of Rao and Schunck [15] for computing orientation and coherence images for summarising oriented textures.

3.4. Defect detection with the top-hat operator

A first example of defect detection in oriented textures is presented: the circular centred top-hat operator applied to some Brodatz textures (taken from [4]).

3.5. Defect detection in textiles

A second example of the use of orientation information for detecting defects in oriented textures, also taken from [4].

3.6. Defect detection in wood

Some examples of the use of the circular centred gradient and the labelled opening operators to detect defects in wood characterised by anomalies in the vein orientations. These defects are essentially the knots and the small light patches (called “maille” in French).

3.7. Segmentation of an oriented texture

A brief consideration of the use of the circular centred gradient operator as the first step in the segmentation of an oriented texture (into regions of homogeneous orientation) by the watershed operator.

3.8. Conclusion

4. Colour spaces

4.1. Colour spaces in cylindrical coordinates

This section presents a generalisation of the colour spaces represented in cylindrical coordinates (HSV, HSI, etc.). It is presented (in less detail) in [8, sections 3 and 4].

4.2. The L*a*b* space

The section is essentially a direct translation of [11, chapter 2].

4.3. Use of these spaces

A summary of the first half of this section is presented in [8, section 2].

4.4. Summary

5. Applications to colour spaces

5.1. Colour statistics

This section is a translation of [6, chapter 3].

5.2. Vectorial mathematical morphology

A presentation of the standard vectorial orders discussed in [2]: *marginal order*, *reduced order*, and *conditional* or *lexicographical order*.

5.3. Lexicographical order in the TYS space

This is essentially a translation of [6, chapter 5], which is summarised in [10]. The TYS space is a generalised cylindrical coordinate space represented by hue (T for “teinte”), luminance (Y) and saturation (S).

5.4. Weighting function in the L*a*b* space

This is essentially a translation of [11, chapter 3], summarised in [12]. The main difference is the improvement in the positioning of the negative “charges”. In [11], they are placed at equal *angular* distance, requiring that their magnitudes be proportional to the distance from the axis to get a good potential function. In the thesis, they are all of equal magnitude and placed at equal *Euclidean* distance from each other.

5.5. Conclusion

6. Conclusion

In this work, our main contribution to the state of the art of image analysis is the development of morphological operators for circular data and the demonstration of the affinities between the processing of oriented textures and colour images. There are certainly extensions which remain to be studied. We envisage theoretical extensions, including the development of methods for applying the rotationally invariant operators to vectorial images, and the development of morphological reconstruction operators for colour images. Other domains of application for the operators developed are in the processing of sound represented in the form of a spectrogram, and of electron microscope images.

Part II.

SCANMATCH

7. Introduction

This second part of the thesis presents results obtained in the framework of the European project SCANMATCH. This project was concerned with the development of an automated process for the assembly of planks of wood according to various æsthetic criteria with the aim of constructing wood panels for use in the furniture industry.

7.1. Project description

The various project requirements are introduced. The most important information is the definition of the three traditional wood texture classes: *quarter* (quartier), *false-quarter* (faux-quartier) and *slab* (dosse), with examples shown in figure 7.2.

7.2. Databases

The image databases are described. The two largest are collections of 152 colour images of wild cherry (merisier) planks, and 146 colour images of oak (chêne) planks. Various smaller databases of pine and oak were also used.

7.3. Pre-processing of wood images

The luminance band was threshold to separate the wood from the background. Colour characteristics were calculated within the thresholded region, and texture characteristics within the smallest rectangle surrounding the thresholded region.

7.4. Automatic classification

A basic introduction to statistical classification algorithms, in particular the Mahalanobis distance classifier and performance evaluation by leave-one-out cross-validation (see [17] for details).

8. Classification and matching by colour

8.1. Wood colour properties

We use a method of colour reconstruction from luminance proposed by Albiol [1] to demonstrate that each species of wood belongs to a single colour category. In other words, with the chosen model of red, green and blue as a function of luminance (equation 8.1), only six parameters are needed to rebuild the colours of all the images of a species from their corresponding luminance images (as long as all the images are captured under the same lighting conditions).

8.2. Colour features

We demonstrate that the TLC space has less correlation between the colour channels than the RGB space. We suggest describing the global wood colour by a vector of six values, the means of each of the TLC channels and their corresponding standard deviations. The standard deviations are used to characterise the vein contrast (an example is given in figure 8.6).

8.3. Colour classification experiments

We begin by testing the level of agreement of human classifiers. Five colour classes are defined for wild cherry, and the level of agreement between three human classifiers is given in table 8.2.

The vector of colour characteristics is used in a Mahalanobis distance classifier to classify all the images in the wild cherry and oak databases. The actual class (ground truth) for each piece is chosen based on the results of the manual classification experiments (three human classifiers for wild cherry, one for oak). The confusion matrices are given in table 8.3 (five classes for wild cherry, four classes for oak).

Finally, the results of a matching experiment using the Mahalanobis distance (with a diagonal covariance matrix) between colour feature vectors to find the three planks in the database with the colour closest to a selected plank are presented in figure 8.8. For these results, the plank with the colour closest to the user-selected plank is chosen, then the plank closest to this second plank is found, etc.

8.4. Conclusion

9. Classification and matching by texture

The veins are the main wood texture feature. They are characterised by their local orientation and inter-vein spacing.

9.1. Vein orientation

The Rao and Schunck algorithm (chapter 3) is used to calculate the local vein orientation. In this section, the effect on the results of the neighbourhood size and of the variance of the Gaussian filter is investigated.

9.2. Inter-vein spacing

An algorithm for calculating the distance between the veins using the co-occurrence matrix is presented. It is also described in [9].

9.3. Texture features

We have two summary images of the wood texture: the orientation image and the inter-vein spacing image. In order to reduce the number of inputs to an automated classifier, twelve texture features are calculated from these images. The first ten are calculated from the orientation summary image, and the last two from the inter-vein spacing summary image.

9.4. Texture classification experiments

The procedure is similar to the one followed for the colour classification experiments (section 8.3). We begin by comparing the classification of wild cherry images by four human classifiers into the six texture classes sketched in figure 9.1. The ground truth is then defined based on these experiments.

We examine, by the use of plots, the capability of two global features (mean circular direction μ and circular mean length $\overline{R}$) to separate five of the texture classes (class 1 was omitted due to a lack of examples in the databases). Some of the outliers are discussed. We then apply a Mahalanobis distance classifier to a vector of five texture features to perform a classification into five texture classes, with the resulting confusion matrices shown in table 9.4.

Finally, the results of two automated texture matching algorithms are presented. The first, called the “Distance” algorithm, is analogous to the colour matching algorithm, being the calculation of the Mahalanobis distance (with a diagonal covariance matrix) between vectors containing the twelve texture features. The second, the “Règles” (rules) algorithm, is based on the definition of texture compatibility rules by the user. We present two rules which define the allowed relation between vein orientations and spacings on the adjoining edges of two boards. The results of using these two algorithms to find the three planks closest to one picked from the database are shown in figure 9.17. For these results, the plank with the texture closest to the user-selected plank is chosen, then the plank closest to this second plank is found, etc.

9.5. Conclusion

10. Matching by colour and texture

In this section, the colour matching algorithm is combined with the texture matching by rules algorithm to perform matching by colour and texture. After the user chooses a plank, the algorithm searches for all the planks in the database which are compatible according to the texture rules, and among these planks, finds the one closest in colour. The results are presented in figure 10.1.

11. Conclusion

A prototype machine which uses the classification algorithms for colour and texture has been constructed. Nevertheless, an automated *matching* machine is more difficult to design as the matching process is more difficult to manage. For example, one never knows if the next plank which is scanned by the machine will be compatible with one of the panels under construction in a finite number of output boxes. This type of dynamic panel construction process and its relation to the colour and texture compatibility parameters still needs to be studied in detail.

Bibliography

[1] Antonio Albiol. Chrominance restoration from luminance by statistics matching. In *Proceedings of the 1997 IEE Conference on Image Processing and its Applications*, July 1997.

[2] V. Barnett. The ordering of multivariate data. *Journal of the Statistical Society of America A*, 139(3):318–354, 1976.

[3] Dmitry Chetverikov. Fundamental structural features in the visual world. In *Proceedings of the International Workshop on Fundamental Structural Properties in Image and Pattern Analysis*, pages 47–58, 1999.

[4] Dmitry Chetverikov and Allan Hanbury. Finding defects in texture using regularity and local orientation. *Pattern Recognition*, 35(10), September 2002.

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

[6] Allan Hanbury. Lexicographical order in the HLS colour space. Technical Report N-04/01/MM, Centre de Morphologie Mathématique, Ecole des Mines de Paris, 2001.

[7] Allan Hanbury. *Morphologie Mathématique sur le Cercle Unité: avec applications aux teintes et aux textures orientées*. PhD thesis, CMM, Ecole des Mines de Paris, 2002.

[8] Allan Hanbury. The taming of the hue, saturation and brightness colour space. In *Proceeding of the Seventh Computer Vision Winter Workshop, Bad Aussee, Austria*, 2002.

[9] Allan Hanbury and Frédéric Gardeux. A quantitative description of wood texture. In *Proceedings of the International Conference on Quality Control by Artificial Vision (QCAV'2001)*, pages 200–205. Cépaduès-Editions, May 2001.

[10] Allan Hanbury and Jean Serra. Mathematical morphology in the HLS colour space. In Tim Cootes and Chris Taylor, editors, *Proceedings of the British Machine Vision Conference 2001*, pages 451–460. BMVA, 2001.

[11] Allan Hanbury and Jean Serra. Mathematical morphology in the L*a*b* colour space. Technical Report N-36/01/MM, Centre de Morphologie Mathématique, Ecole des Mines de Paris, 2001.

[12] Allan Hanbury and Jean Serra. Mathematical morphology in the L*a*b* space. In *Proceedings of the 8th European Congress for Stereology and Image Analysis*, 2001.

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

[14] K. V. Mardia. *Statistics of Directional Data*. Academic Press, 1972.

[15] A. R. Rao and B. G. Schunck. Computing oriented texture fields. *CVGIP: Graphical Models and Image Processing*, 53(2):157–185, March 1991.

[16] A. Ravishankar Rao. *A taxonomy for texture description and identification*. Springer-Verlag, 1990.

[17] B. D. Ripley. *Pattern Recognition and Neural Networks*. Cambridge University Press, 1996.

[18] J. Serra. *Image Analysis and Mathematical Morphology*. Academic Press, London, 1982.

[19] Pierre Soille. *Morphological Image Analysis: Principles and Applications*. Springer-Verlag, 1999.