


Geometric Image Description Using
the Multiscale Orientation Field

TR91-036

August, 1991

Kah-Chan Low

The University of North Carolina at Chapel Hill
Department of Computer Science
CB#3175, Sitterson Hall
Chapel Hill, NC 27599-3175



UNC is an Equal Opportunity/Affirmative Action Institution.

Geometric Image Description Using the Multiscale Orientation Field

by

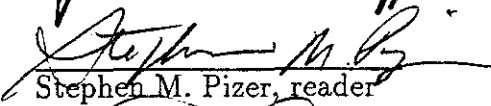
Kah-Chan Low

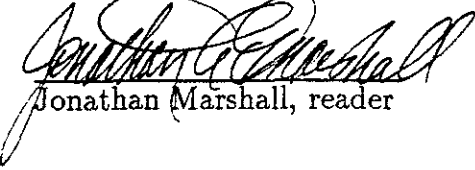
A thesis submitted to the faculty of the University of
North Carolina at Chapel Hill in partial fulfillment of the
requirements for the degree of Master of Science in the
Department of Computer Science.

Chapel Hill, 1991

Approved by:


James M. Coggins, advisor


Stephen M. Pizer, reader


Jonathan Marshall, reader

©1991
Kah-Chan Low
ALL RIGHTS RESERVED

KAH-CHAN LOW. Geometric Image Description Using the Multiscale Orientation Field
(Under the direction of JAMES M. COGGINS.)

Abstract

An investigation was conducted to determine the strengths and limitations of the Multiscale Orientation Field (MOF), which was proposed by Coggins as a multiscale representation of the local orientation at each pixel of an image. The MOF is obtained by filtering the original image with an orientation filter bank and subsequently combining the outputs using vector algebra. The design of the filter bank, together with the method of computing the MOF, provide a uniform orientation sensitivity along the orientation axis.

The MOF and its computation method was compared with similar algorithms developed by Grossberg, Hsieh, and Zucker using a number of criteria. The comparison provided a number of predictions regarding the utility of the MOF. The MOF's of several simple binary image structures were computed and the patterns of orientation vectors on these MOF's were observed and accounted for. These experiments provided direct verifications of the predictions made above.

ACKNOWLEDGMENTS

I wish to thank Professor James Coggins for his support, guidance and encouragement, and most of all for his enthusiasm, throughout the course of my research. He left me his door open so that I could discuss problems with him at any time. When things didn't work he offered me encouragement and advice, and this work could never have been completed without his guidance. I also wish to thank Professor Pizer and Marshall for their valuable ideas and suggestions in this work.

I also wish to thank Manish Pandey for allowing me to plagiarize much of the contents of this page! More seriously, the reason I plagiarize from Manish is because he describes my experience and feeling so eloquently that I couldn't have put it better myself.

CONTENTS

I. Introduction	1
II. Relevant Research	3
2.1 Grossberg's and Mingolla's Boundary Contour System	3
2.2 Hsieh's connectionist algorithm for image segmentation	4
2.3 Zucker's relaxation labeling approach towards oriented patterns	4
III. The filter bank	7
3.1 The Orientation Sensitivity Function	8
3.2 Computation of the MOF	8
3.3 Implementation of Coggins' method	10
3.4 Two possible models for analysis	11
IV. A Comparison Between Related Methods	12
4.1 Number of filters required	12
4.2 Computational requirement	13
4.3 Computational mechanism	13
4.4 Accuracy of orientation estimates	14
4.4.1 Straight lines and edges	14
4.4.2 Curved lines and edges	15
4.4.3 Corners and junctions	15
V. Experiments on Straight Edges	16
5.1 A Straight Step Edge	16
5.1.1 Magnitude profiles of orientation vectors	16
5.1.2 Phase angles of orientation vectors	20

5.1.3	Remarks	22
5.2	A Blurred Edge	23
5.2.1	Magnitude profiles	23
5.2.2	Phase profiles	25
5.2.3	Orientation estimation	25
5.3	The Effect of Noise on Orientation Vector Lengths	27
5.4	Summary	30
VI.	Curved Edges and Elongated Structures	32
6.1	The macro-geometry of an object	32
6.2	An Elliptical Disc	33
6.2.1	Vector profiles along the minor axis	33
6.2.2	Vector profiles along the major axis	37
6.2.3	Comments	43
6.3	Summary	45
VII.	Corners	50
7.1	The 30° corner	50
7.2	The 150° corner	56
7.3	Comments	57
VIII.	Summary and Evaluation	58
8.1	A final summary	58
8.1.1	Observations	58
8.1.2	Theoretical findings and evaluations	59
8.2	Future research directions	60
	Bibliography	62

LIST OF TABLES

6.1	Locations (columns) of magnitude troughs and peaks along row 127 of figure 6.1	35
6.2	Phase angles corresponding to magnitude peaks along row 127 of figure 6.1	37
6.3	Locations (rows) of magnitude troughs and peaks along columns 127 of figure 6.1	39
6.4	Phase angles corresponding to magnitude peaks along column 127 of figure 6.1 (see footnote on page 37 for explanations)	41
6.5	Locations (columns) of magnitude troughs and peaks along row 127 of figure 6.9	46
7.1	Orientation vectors in the immediate neighborhood of the corner in figure 7.1	54
7.2	Orientation vectors in the immediate neighborhood of the corner in figure 7.2	55

LIST OF FIGURES

1.1 The Artificial Vision System (taken from [1])	2
2.1 A block diagram describing the processing stages of Hsieh's algorithm (Taken from [6])	5
3.1 A photograph of $f_{\frac{\pi}{2}}(r, \theta; 16\sqrt{2}, \pi, \frac{\pi}{8})$	8
3.2 Mapping the filter responses to the Abstract Feature Space (Taken from [11])	9
5.1 A straight step edge	17
5.2 Vector magnitude curves of figure 5.1 along row 128 ($\sigma = \sqrt{2}, 2\sqrt{2}, 4\sqrt{2}, 8\sqrt{2}$) .	18
5.3 Thought experiment: the center of the X-orientation filter on top of the edge	19
5.4 Thought experiment: the center the X-orientation filter at an offset from the edge	19
5.5 Thought experiment: an edge with high curvature	20
5.6 Vector phase angles of figure 5.1 along row 128 ($\sigma = 1$)	21
5.7 A fuzzy edge	23
5.8 Vector magnitude curves of figure 5.7 along row 135 ($\sigma = 4\sqrt{2}, 8\sqrt{2}, 16, 16\sqrt{2}$) .	24
5.9 Vector phase angles of figure 5.7 along row 135 ($\sigma = \sqrt{2}, 4\sqrt{2}, 8\sqrt{2}, 16$)	26
5.10 A noisy and fuzzy edge	27
5.11 Vector magnitude curves of figure 5.10 along row 135 ($\sigma = \sqrt{2}, 2\sqrt{2}, 4\sqrt{2}$) . . .	28
5.12 Vector magnitude curves of figure 5.10 along row 135 ($\sigma = 8\sqrt{2}, 16\sqrt{2}$)	29
6.1 An elliptical disc centred at (127,127)	33
6.2 Vector magnitude curves of figure 6.1 along row 127 ($\sigma = 2, 2\sqrt{2}, 8, 8\sqrt{2}$) . . .	34
6.3 Vector magnitude curves of figure 6.1 along column 127 ($\sigma = 2, 2\sqrt{2}, 8, 8\sqrt{2}$) .	38
6.4 Filters of different sizes superimposed on an elongated structure	40

6.5	Filters of different sizes superimposed on a straight edge	40
6.6	Filter superimposed on a highly curved edge	42
6.7	Filter superimposed on a slightly curved edge	43
6.8	Vector field configuration ($\sigma=1$) in the vicinity of an edge curvature maximum in figure 6.1	44
6.9	A circular disc (centred at (127,127))	46
6.10	Vector magnitude curves figure 6.9 along row 127 ($\sigma = 2, 2\sqrt{2}, 8, 8\sqrt{2}$)	47
6.11	Orientation vector field configuration near the edge of the circular disc in figure 6.9 at $\sigma= 8\sqrt{2}$ (Dotted line explained on page 42)	48
7.1	A 30° corner	51
7.2	A 150° corner	51
7.3	Orientation vector field configuration for figure 7.1 at $\sigma=2$ (Dotted line ex- plained on page 42)	52
7.4	Orientation vector field configuration for figure 7.2 at $\sigma=2$ (Dotted line ex- plained on page 42)	53
7.5	Thought experiment on an acute corner	54
7.6	Thought experiment on an obtuse corner	55

Chapter 1

Introduction

The objective of this thesis is to investigate the configurations of the Multiscale Orientation Field (MOF) in the neighborhoods of a number of image structures. The MOF is formally defined in section 3.2. For now, the MOF can be understood as containing multi-scale representations of the local orientation at each pixel of an image. It is the output of a computational architecture proposed by Coggins called the Artificial Visual System (AVS) ([2]). Figure 1.1 is a schematic diagram showing the configuration of the AVS. The AVS consists of two components – a filter bank and a recombination algorithm which operates on the output of the filter bank to produce an inference about the scene.

The filter bank disperses information from the original image into separate bands of scale or orientation. In this thesis, the filter bank to be investigated consists of 10 sets of *orientation filters* (see Chapter 3). Each of the 10 sets are identical except for their sizes (or scales). The particular recombination algorithm used here combines the output of the filter bank using vector algebra to produce the MOF. This recombination algorithm is dependent on the specification of the orientation filter bank.

In the process of computing the MOF, some image features in the original image are inevitably destroyed while others are preserved, even enhanced. In order to use the MOF in an object-definition or segmentation algorithm, it is very important to have a thorough understanding of the different configurations of the MOF that might arise in response to different image structures. These configurations can be used to ascertain whether the MOF destroys, preserves, or enhances particular image features. If structure is destroyed, we must determine the extent to which the final goals are compromised by the loss. If the feature is

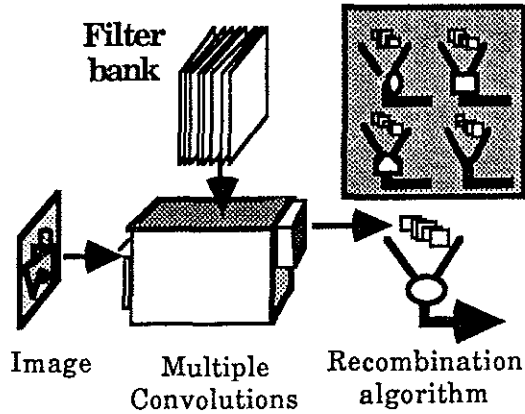


Figure 1.1: The Artificial Visual System (taken from [1])

enhanced, we must explore the manner in which the feature is enhanced. Closely tied to the question of feature destruction and enhancement is the utility of the MOF – its strengths and limitations need to be carefully mapped.

In this thesis, orientation vector field configurations that arise in response to a number of edges with different curvatures and corners will be observed and recorded across several scales. Configurations of the MOF corresponding to a straight step edge will be examined first, followed by a straight but blurred edge, a noisy edge, edges of an elliptical and a circular disc, and two types of corners. These image structures are increasingly difficult to capture using the MOF. The above certainly does not represent an exhaustive list of image structures whose representations need investigation and, if possible, characterization, if the properties of the MOF are to be thoroughly understood. However, it is necessary to have a good grasp of the representations of simple image structures such as these before those belonging to more complicated image structures are explored.

Chapter 2

Relevant Research

The visual field in primates is covered by overlapping receptive fields that come in a number of sizes and shapes. Research in the past two decades has revealed several important properties of receptive fields. In particular, knowledge of the filter-like behavior of receptive fields and the preponderance of receptive fields that are orientation sensitive ([7]) have led a number of researchers to formulate computer vision algorithms within the framework of receptive field-like filtering operators in which orientation sensitivity plays an important role.

2.1 Grossberg's and Mingolla's Boundary Contour System

In the Boundary Contour System (BCS) ([5]), Grossberg and Mingolla employs a neural net implementation to account for subjective contours in illusory images such as Kanizsa's triangle. In the particular implementation described in [5], an image was first filtered using four pairs of elongated step edge filters, with each pair sensitive to one of the four orientations of preference which are spaced 45° apart. Each pair of oriented edge filters consists of two filters with the same preferred orientation but different polarities. The output of these filters are fed, in successive stages, into five layers of "neural cells", with a feedback loop linking the second the fifth layers. Within these stages, the cells perform a series of positional competitions and cooperations. More relevant to this thesis, orientation competitions at a particular location occurs between cells on the second and third layers. During these competitions at a particular location, responses derived from filter pairs that are 90° apart cancel each other. In this fashion the final orientation preference of an edge can be determined or at least approximated.

More recently the BCS has been modified and further extended by Neumann and Steihl ([12],[13]) to include corner and junction detection.

2.2 Hsieh's connectionist algorithm for image segmentation

The goal of Hsieh's algorithm ([6]) is to segment both artificial and natural images. In particular, the algorithm is designed to have the capability of performing segmentation based not only on real edges but also on subjective contours. In devising the algorithm Hsieh adopts a connectionist approach. A block diagram depicting the processing involved can be found in figure 2.1. The algorithm allows for a multiscale approach but does not have any provision for cross-scale interactions, which could resolve ambiguities about boundary locations and grouping of image features into objects. Within each scale several sets of filters having different degrees of elongation are used. Each set of filters consists of four pairs of the first derivative of Gaussians, with each pair having a different preferred orientation. For each filter pair, the two constituent filters are identical except that they have opposite polarities. Explicit corner detection and a number of spatial coherence checks are also provided to counter the adverse effects of noise and false positive responses given by some of the edge filters in the presence of arbitrarily complicated image structures. Hsieh found it necessary to subtract responses of perpendicular filters. This operation is called *artifact cancellation*. Hsieh suggests that artifact cancellation counters the adverse effects of aliasing due to discrete sampling.

2.3 Zucker's relaxation labeling approach towards oriented patterns

An objective of Zucker's research described in ([16]) is to infer orientations at locations on images with oriented patterns on them. The algorithm uses an approximation to the Laplacian of elongated Gaussians with different preferred orientations as filters. A relaxation labeling algorithm ([8]) is applied on the output images of these filters to determine the preferred orientation at every location on the image. Zucker attempts to address the problem of estimating orientations at locations where the oriented patterns are curved. In addition, he distinguishes between two types of oriented patterns, which he calls Type-I and Type-II




Figure 2.1: A block diagram describing the processing stages of Hsieh's algorithm (Taken from [6])

patterns. Type-I patterns are characterized by well-defined and non-overlapping oriented contours; whereas in Type-II patterns, oriented contours are, at certain locations, often occluded by other contours which are oriented differently from the former. Contour lines in Type-II patterns are therefore more *implicit* than those of Type-I patterns, and interpolations across “gaps” are required to extract these contour lines.

As a step toward solving the problem of estimating orientations of curved contour lines, Zucker defines the orientation of a point on a curved line to be the orientation of the tangent touching that point. The orientation of a point on the curved line is then estimated by matching the configuration of the orientation field lines in the neighborhood of that point to that produced by seven types of *osculating circle* templates for each preferred orientation of the filters. Two different classes of osculating circle templates are introduced – one for Type-I patterns and the other for Type-II patterns. Each of the seven Type-I templates consists of simple osculating circles. In contrast, each of the Type-II templates consists of a family of osculating circles, which are the result of *laterally interpolating* an osculating circle. The matching process is formulated in terms of a *response matching problem*, which is solved via relaxation labeling. Details on the *compatibility function* used can be found in [14].

At this point, the relevance of the foregoing algorithms to this thesis is not immediately obvious. In Chapter 3, the filters used to produce the MOF will be mathematically specified. The exact method by which the MOF is computed will also be explained. After that, I shall explain the relevance of these algorithms and compare the computational method used in this thesis with the algorithms above.

Chapter 3

The filter bank

A filter set belonging to a particular scale (σ) in the filter bank employed in this thesis is an instance derived from the following general formula:

$$f_{\phi}(r, \theta; \sigma, \Theta, \omega) = kG_{\sigma}(r)S_{\phi}(\theta; \Theta, \omega) \quad \phi = \omega, 2\omega, \dots, \Theta \quad (3.1)$$

ϕ is a dependent parameter indicating the preferred orientation of an individual filter. (r, θ) are the polar coordinates of a point; k is a normalizing constant that makes the filter's total energy equal to unity; $G_{\sigma}(r)$ is an isotropic Gaussian with a standard deviation of σ ; and $S_{\phi}(\theta; \Theta, \omega)$ is the *orientation sensitivity function* for the filter, to be described in section 3.1. The parameters Θ and ω are explained in detail in [11]. Briefly, Θ is the angular periodicity of a filter class (or equivalently, the length of the orientation dimension on which the filter class is defined); that is, $f_{\phi}(r, \theta) = f_{\phi}(r, \theta + k\Theta)$ where k is an integer. Θ and ω totally specify the shape (but not the size, which depends on σ) of the filters in a filter set. The number of filters constituting a filter set is given by the expression $\frac{\Theta}{\omega}$. The filter bank used here is created with $\Theta = \pi$ and $\omega = \frac{\pi}{6}$. Ten different values of σ are used. The smallest σ is 1; the largest σ is $16\sqrt{2}$; and consecutive σ 's differ by a multiplicative factor of $\sqrt{2}$.

With these values of Θ and ω , the filter set corresponding to a particular σ has six constituent filters. The preferred orientations (ϕ) of these filters are: $\frac{\pi}{6}$, $\frac{\pi}{3}$, $\frac{\pi}{2}$, $\frac{2\pi}{3}$, $\frac{5\pi}{6}$, and π . Because $\Theta = \pi$, a filter with $\phi = \pi$ is identical to the one with $\phi = 0$ ([11]). For the sake of clarity, a filter belonging to a filter set produced with $\Theta = \pi$ and $\omega = \frac{\pi}{6}$ will henceforth be referred to as an *orientation filter*.




Figure 3.1: A photograph of $f_{\frac{\pi}{2}}(r, \theta; 16\sqrt{2}, \pi, \frac{\pi}{6})$

3.1 The Orientation Sensitivity Function

The orientation sensitivity function, whose complete derivation can be found in [2] and [11], is:

$$S_{\phi}(\theta; \Theta, \omega) = \begin{cases} \frac{\sin(\Omega - \frac{\Omega}{\omega}|\theta - \phi|)}{\sin \Omega} & \text{if } \phi - \omega < \theta < \phi + \omega \\ 0 & \text{otherwise} \end{cases} ; \quad \Omega \equiv \frac{2\pi\omega}{\Theta} \quad (3.2)$$

Substituting $\Theta = \pi$ and $\omega = \frac{\pi}{6}$, equation 3.2 becomes:

$$S_{\phi}(\theta; \pi, \frac{\pi}{6}) = \begin{cases} \frac{\sin(\frac{\pi}{3} - 2|\theta - \phi|)}{\sin \frac{\pi}{3}} & \text{if } \phi - \frac{\pi}{6} < \theta < \phi + \frac{\pi}{6} \\ 0 & \text{otherwise} \end{cases} \quad (3.3)$$

Figure 3.1 shows a photograph of the orientation filter $f_{\frac{\pi}{2}}(r, \theta; 16\sqrt{2}, \pi, \frac{\pi}{6})$. The filter derives its peculiar shape from its orientation sensitivity function, which is designed such that the orientation filter set to which the filter belongs provides equal support for all orientations around a circle, although the filter set covers only six discrete orientations ([2],[11]).

3.2 Computation of the MOF

To recombine the output of a set of orientation filters, an abstract feature space, which is a polar-ruled plane, is defined. Each pixel is mapped into the abstract feature space using

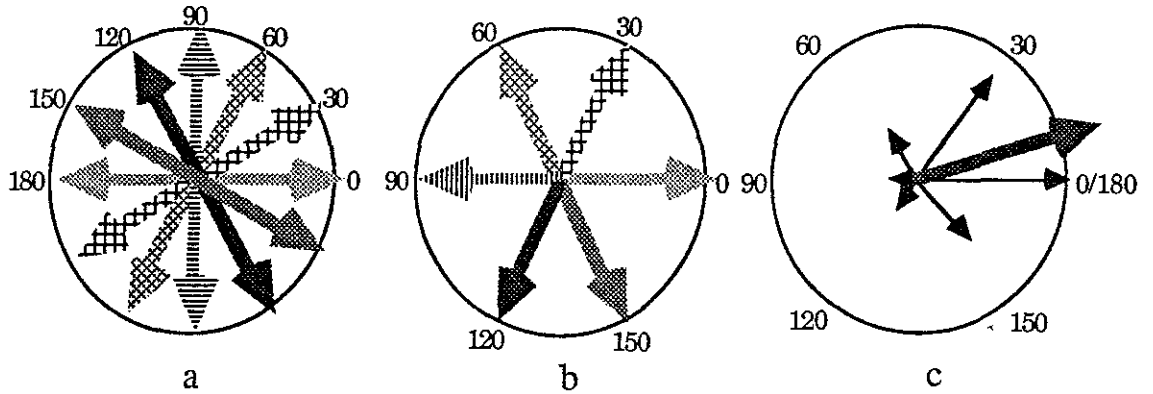


Figure 3.2: Mapping the filter responses to the Abstract Feature Space (Taken from [11])

the responses from each orientation filter set corresponding to a particular scale (σ). The mapping is performed as follows: at pixel (x, y) , the response ρ to filter f_ϕ (figure 3.2a) is mapped onto the abstract feature space as a vector $\vec{\rho}$ having magnitude ρ and phase angle 2ϕ (figure 3.2b). The phase angles of these vectors are evenly distributed from 0 to 2π in the abstract feature space just as the spatial orientations of the orientation filters are evenly distributed from 0 to π . These vectors in the abstract feature space, one corresponding to the output of each orientation filter, are summed to yield a resultant vector (figure 3.2c). The phase angle of the resultant vector is divided by 2 to map it back to the spatial domain. A vector field thus derived from an original image and an orientation filter set is called an *orientation vector field*. A collection of orientation vector fields derived from the same original image but using orientation filter sets of different scales are collectively known as the *Multiscale Orientation Field* (MOF). It should be noted that the specific way in which the filter responses are mapped onto the abstract feature space and back is dependent on the shape specification of the filters in the filter bank (i.e. Θ and ω). The general mapping formula can be found in [11]. For the sake of brevity, the specific mapping and vector computation described above will henceforth be referred as *Coggins' method*.

Due to the way in which the response of an orientation filter is mapped onto the abstract feature space, the responses of orientation filters that are perpendicular to each other will cancel each other. Pairs of orientation filters that are related in such a way will be referred to as *antagonistic filter pairs*. In addition, a member of an antagonistic filter pair is called the *antagonistic filter* of the other member in the same pair. A set of orientation filters consists of three antagonistic filter pairs.

Mutual cancellation of responses between antagonistic filter pairs (or pairs of filters with a similar configuration) seems to be a rather common technique employed in computations that involve orientations of image structures. For instance, Kass and Witkin ([9]) employ a method that closely resembles Coggins' method to compute orientations of oriented patterns. Kindred methods can also be found in Grossberg's BCS and Hsieh's algorithm described in Chapter 2. A comparison between Coggins' method and the two foregoing algorithms can be found in Chapter 4.

3.3 Implementation of Coggins' method

Coggins' method, if naively implemented, requires an original image to be filtered six times before an orientation vector field can be computed. Although the six filtering operations can be performed in parallel, there is certainly room for reducing the method's computational requirement. Since the computation of an orientation vector field is a linear process, the orientation filters can be designed such that the step of computing the orientation vector field can be partially incorporated into the filters themselves. Consider the step where the vector sum is performed in the abstract feature space. A convenient way to sum the vectors is to perform two independent summations in the x and y directions:

$$V_x = \sum_{i=0}^5 \cos\left(\frac{\pi}{3}i\right) \left[f_{\frac{\pi}{6}i} \otimes I \right] \quad (3.4)$$

$$V_y = \sum_{i=0}^5 \sin\left(\frac{\pi}{3}i\right) \left[f_{\frac{\pi}{6}i} \otimes I \right] \quad (3.5)$$

where V_x and V_y respectively are the resultant x and y vector components in the abstract feature space. $f_{\frac{\pi}{6}i}$ is the orientation filter with a preferred orientation of $\frac{\pi}{6}i$; and I denotes the original image. Finally, \otimes denotes the convolution operator. Because convolution is distributive over addition, equations 3.4 and 3.5 can be rewritten as:

$$V_x = \left[\sum_{i=0}^5 f_{\frac{\pi}{6}i} \cos\left(\frac{\pi}{3}i\right) \right] \otimes I \equiv F_x \otimes I \quad (3.6)$$

$$V_y = \left[\sum_{i=0}^5 f_{\frac{\pi}{6}i} \sin\left(\frac{\pi}{3}i\right) \right] \otimes I \equiv F_y \otimes I \quad (3.7)$$

It is obvious that F_x and F_y can be implemented as two filters. Thus, only two filters are needed to compute an orientation vector field. F_x and F_y will from now be respectively

referred to as the X- and the Y-orientation filters. Because of this simple manipulation, the computational and memory requirements are now reduced by $\frac{2}{3}$ over the naive implementation.

3.4 Two possible models for analysis

As a result of the above mathematical manipulation, an orientation filter set can be regarded as consisting of six orientation filters or two composite filters. This means that visualizations (as used in “thought experiments”) can be performed using either the six-filter model or the two-filter model as circumstances warrant. In addition, as mentioned in section 3.1, individual orientation filters are designed such that the orientation filter set to which a filter belongs provides equal support for all orientations around a circle. This means that it is possible to mentally “rotate” the set of the filters arbitrarily to simplify visualization, as long as the relative positions of the component filters to each other remain unchanged.

Chapter 4

A Comparison Between Related Methods

The four computational frameworks defined by Grossberg and Mingolla, Hsieh, Zucker, and Coggins share an underlying theme of orientation determination of image structures. To this end all of them employ receptive field-like filters to perform preliminary feature extraction. On the other hand, the actual architectures of the four frameworks differ, and the results of orientation determination are used for different purposes.

In this chapter, Coggins' method will be compared with the three algorithms discussed in Chapter 2 from the following perspectives: 1) number of filters needed; 2) computational requirement; and 3) computational mechanism; 4) the accuracy with which orientations of several general image structures can be estimated.

4.1 Number of filters required

Eight filters are used in the algorithms defined in [5] and [16]. Hsieh's algorithm needs at least eight filters for each scale. Coggins' method only requires two filters at each scale. Because an orientation vector field is a linear combination of the outputs of the orientation filters, a reduction of the number of filters needed can be achieved by partially building the vector computing stage directly into the filters (section 3.3). It is difficult to imagine how such an incorporation can be implemented in the cases of BCS and Zucker's algorithm, since the operations involved are highly non-linear. In Hsieh's algorithm, the filter responses are partially retained for performing corner/junction detections and spatial coherence checks. In fact,

Hsieh's requirement reveals not so much a shortcoming than a trade-off – incorporating the vector sum computation in the design of two composite filters means that responses of individual orientation filters are unavailable for further processing. These individual orientation filter responses might be useful for distinguishing image structures that are indistinguishable in the MOF.

4.2 Computational requirement

Of the three algorithms described in Chapter 2, only Zucker's algorithm can be readily used for comparison with Coggins' method under this criterion. Both BCS and Hsieh's algorithm require enormous computational power, but they do much more than orientation determination, and in both cases, it is not easy to isolate orientation measurement from the rest of the computational frameworks. In contrast, Zucker's algorithm is totally devoted to orientation determination. To understand the difference in computational requirement between Zucker's and Coggins' algorithms, a précis of the steps involved in the former is in order.

In Zucker's algorithm, eight elongated Laplacian-of-Gaussian filters having different preferred orientations are convolved with the original image. Each pixel is then matched against seven osculating circle templates (section 2.3) by relaxation labeling. The time-complexity of Zucker's algorithm is dependent on both the size of the original image and the size of the neighborhoods over which relaxation labeling is performed.

In contrast, the computation of an orientation vector field is a non-iterative, one-step process. No neighbor pixel needs to be consulted in the process. Thus, the computational cost increases only linearly with the number of pixels in the original image. Recall also that Coggins' algorithm requires two filters (per scale), while eight are used in Zucker's algorithm.

4.3 Computational mechanism

In computing the MOF, the responses of the orientation filters are mapped onto an abstract feature space such that responses derived from an antagonistic filter pair cancel each other (section 3.2). The counterparts of this mechanism exist in the BCS and Hsieh's (but not Zucker's) algorithm.

In the BCS, this mechanism, known as orientation competition (section 2.1), constitutes

one of the two competitive stages, which in turn make up an integral part of the neural net feedback loop. Consequently, unlike the computation of the MOF, orientation competition can potentially be performed many times before the neural net reaches an equilibrium state. In addition, orientation competition does not destroy or discard orientation variables that initially contain responses of individual filters with different preferred orientations. This is borne out by the need to store (and possibly update) the values of individual responses for subsequent rounds of orientational competition. In contrast, in Coggins' method, responses of individual filters can be totally discarded after computation of the MOF is completed.

In Hsieh's algorithm, artifact cancellation (section 2.2) is an one-step operation. In contrast to the BCS, only one variable will be retained for a pair of antagonistic filters after artifact cancellation has been performed. The final value of the variable is the difference between the responses of its corresponding antagonistic filter pair, and the orientation of that final value is the preferred orientation of the edge filter possessing the larger response in the antagonistic filter pair. Thus, Hsieh uses an approach that is half-way between Coggins' and BCS' approaches.

Hsieh comments on the loss of image structure information due to artifact cancellation. In particular, artifact cancellation destroys corner and junction information. Therefore Hsieh implements an explicit processing stage for corner and junction detection. Since Coggins' method does not retain resultant responses of antagonistic filter pairs as does Hsieh's algorithm, I expected Coggins' method to lose of image structure information, especially around corners and junctions.

4.4 Accuracy of orientation estimates

4.4.1 Straight lines and edges

Since the orientation filters are mathematically formulated to provide equal sensitivity for all orientations ([2],[11]), it can be safely predicted that, within the limits of pixelization errors, the MOF will yield the exact orientations of straight lines and edges. In contrast, the filters used by Grossberg and Mingolla, Hsieh, and Zucker are not similarly corrected. As a result, at least for cases involving straight lines and edges, the foregoing three algorithms, however sophisticated, are bound by sampling theory to contain errors due to the discrete sampling of orientation. In Chapter 5, I shall demonstrate the efficacy of the MOF in providing accurate

orientation estimates for straight edges.

4.4.2 Curved lines and edges

Grossberg and Mingolla and Hsieh do not explicitly cover this subject. Zucker's algorithm, however, devotes a large amount of computational power to the accurate determination of the orientations of curved lines. In fact, the enormous computational requirement of Zucker's algorithm is a direct consequence of Zucker's belief that accurate orientation determination of curved, linear structures can not be taken for granted. His sophisticated treatment does enhance the robustness of his algorithms in the face of curved linear structures. However, an upper bound is inherently imposed on the accuracy of the estimates because of the discretization of orientation values that is necessitated by the relaxation labeling technique.

In order to maintain uniformity across all pixels, Coggins' method does not have a special provision for dealing with curved lines and edges. If we accept Zucker's proposition that special provisions have to be made in order to perform good orientation estimates of curved lines and edges, then it is quite likely that the absence of such a provision in Coggins' method is going to limit the MOF's ability to capture curved image structures. I shall attempt to verify this perceived shortcoming in Chapter 6.

4.4.3 Corners and junctions

Because of the discontinuous nature of corners and junctions, naive differentiation-based methods always encounter difficulties in their presence. To overcome this problem, Hsieh has included in his algorithm an explicit stage for detecting corners and junctions. Grossberg and Mingolla are less clear regarding their handling of corners and junctions, but it appears that at corners and junctions with balanced edge strengths, the final responses for all preferred orientations at these locations approach zero, implying that there is no preferred orientation. Zucker excludes corners and junctions altogether from his problem domain. Similarly, Coggins' method gives no special consideration to these structures. This decision limits the MOF's ability to represent corners and junctions. In Chapter 7, the effects two types of corners have in their neighborhoods on the configurations of the MOF will be noted. I shall also discuss the degree to which the MOF is compromised by the lack of corner and junction consideration.

Chapter 5

Experiments on Straight Edges

In this chapter the MOF configurations in the vicinities of several straight edges are observed. The straight edges used are a noiseless step edge, a noiseless blurred edge, and a noisy blurred edge. Where warranted, observations on these MOF are explained. In addition, the accuracy of orientation estimation of the noiseless blurred edge using its MOF will also be demonstrated.

5.1 A Straight Step Edge

In this section, the MOF of a straight step edge will be investigated. Figure 5.1¹ is first filtered with the filter bank as described in Chapter 3. The MOF is computed as explained in section 3.2.

5.1.1 Magnitude profiles of orientation vectors

Figure 5.2 shows a graph of vector magnitude profiles along row 128 across several scales (σ 's). Although the curves look different, vectors at column 127, which is a part of the edge in the image, vanish in all cases. The existence of these vanishing points can be accounted for by performing a thought experiment in which the X-orientation filter at a particular scale is superimposed on figure 5.1 such that the center of the filter is located at the edge as in figure 5.3. The Y-orientation filter is not taken into consideration because of the symmetry

¹Locations on an image are specified by the coordinates (*row,column*). (0,0) is the top left corner of the image. This convention applies to *all* subsequent images.




Figure 5.1: A straight step edge

of the image in figure 5.3. Responses derived from different parts of the Y-filter exactly cancel each other out.

As the X-orientation filter moves away from the edge in a perpendicular direction to the orientation of the edge into the bright half of the image (figure 5.4), the response cancellation effect weakens, as the filter now covers the bright half considerably more than the dark half. Consequently, the computed vector magnitude becomes larger. The computed magnitude peaks when 75% of the area under the filter is within the bright half. After that, the response cancellation effect makes a comeback and vector magnitudes vanish when the filter is totally inside the bright half.

Thus, in the vicinity of an edge a vanishing point can be found sandwiched between two magnitude peaks. As a matter of fact, as we shall see in section 5.2.1, the term “vanishing point”, as applied to describe the abrupt diminution of vector magnitudes at an edge, is a misnomer. Due to a number of reasons, vector magnitudes in general do not vanish at these “vanishing points” at the edges. Consequently, “local minima” is a more appropriate term for describing these points. For the sake of clarity and brevity, a local magnitude minimum sandwiched between two magnitude peaks both of which can be ascribed to the same edge shall from now on be referred to as a *magnitude trough*.




Figure 5.2: Vector magnitude curves of figure 5.1 along row 128 ($\sigma = \sqrt{2}, 2\sqrt{2}, 4\sqrt{2}, 8\sqrt{2}$)




Figure 5.3: Thought experiment: the center of the X-orientation filter on top of the edge

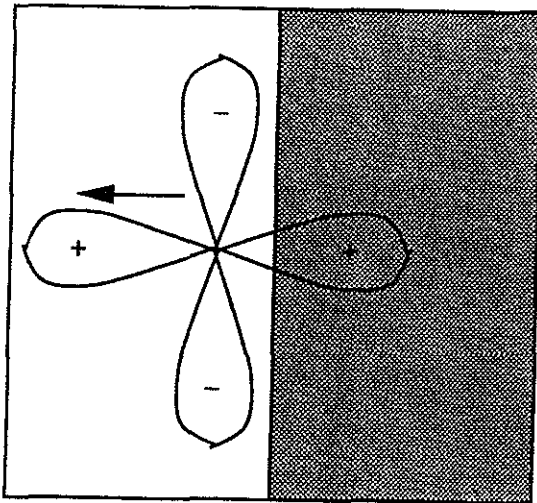


Figure 5.4: Thought experiment: the center the X-orientation filter at an offset from the edge




Figure 5.5: Thought experiment: an edge with high curvature

Magnitude troughs are most pronounced in the case of straight edges (zero curvature), or at least edges that are approximately so with respect to the sizes of the filters used. They become progressively less well defined as the curvature (with respect to the sizes of the filters) of an edge increases, because there simply doesn't exist a point in the vicinity of the edge where exactly $\frac{1}{2}$ of the area under a filter covers each side of the edge. In the case portrayed in figure 5.5 magnitude troughs are all but absent.

5.1.2 Phase angles of orientation vectors

Figure 5.6 depicts the vector phase angles of row 128 of the orientation vector field at $\sigma = 1$. It can be observed that the majority of the vectors belonging to pixels on the bright (left) side of figure 5.1 have phase angle of 90° . The dark (right) side of figure 5.1 generally produces vectors having phase angle of 0° . There are, however, a few columns of maverick vectors on the bright side of figure 5.1 that have phase angles of 0° . Similarly, some vectors corresponding to the dark side of figure 5.1 have phase angles equal to 90° .

Whereas it makes intuitive sense that vectors on the bright side of figure 5.1 that are near the edge have phase angles of 90° , there are two observations that are counter-intuitive. First, the existence of "maverick vectors" is indeed bothersome. Second, phase angles of vectors near the edge on the dark side of figure 5.1 read 0° , instead of 90° .

We see in figure 5.2 that these maverick vectors have very small magnitudes. Hence, the reliability with which we can interpret the phase angles of these vectors as the orientations




Figure 5.6: Vector phase angles of figure 5.1 along row 128 ($\sigma=1$)

of underlying image structures is low. As we shall see, the phase angle of a long orientation vector is not necessarily reliable either. With this notion in mind, the maverick vectors are much less bothersome now, since their phase angle values are virtually meaningless. However, low reliability of phase angles does not imply that they are useless, since they represent a certain amount of information regarding the image structures at these locations.

Since the magnitude of a vector is related to the reliability of its computed phase angle, it is desirable to make vector magnitude an absolute quantity. In this way, similar image structures belonging to different images can be expected to yield similar vectors in terms of both magnitude and phase angle. With this in mind, the grey scale of all images used in the experiments of this thesis are normalized to between 0 and 255. Otherwise, vector magnitudes of corresponding pixels derived from two images that are identical in all aspects except for their grey level ranges are going to be different.

5.1.3 Remarks

Recall the counter-intuitive result that the resultant vectors corresponding to the dark side of figure 5.1 generally have a phase angle of 0° instead of 90° . This result is not at all surprising on closer observation – the orientation filters are constructed to respond to bright oriented structures rather than dark ones. As a result, dark oriented structures respond to the filters in a way that bright structures which are oriented with an angular offset of 90° would have. Because of this ambiguity, if an orientation vector is examined alone, it is impossible to determine with absolute certainty the orientation of the underlying image structure that gives rise to this vector.

From the preceding observations, we can see that two striking features arise near and at a straight step edge. We see that if we march along the direction perpendicular to the orientation of the edge, we encounter a local minimum where the vector at that point vanishes. This point is preceded and followed by a local vector magnitude peak. In addition, there is a sudden switch of vector orientation by 90° at the edge. These two observations apply equally well for vector fields derived from orientation filters belonging to different σ 's.

It must be noted that these two observations should not be taken as a characterization of the behavior of orientation vectors in the presence of edges in general. What has been used here is an ideal step which seldom occurs in real images. The next step is to observe whether these observations still hold if the constraints hitherto imposed are loosened.




Figure 5.7: A fuzzy edge

5.2 A Blurred Edge


Perfect step edges almost never occur in real images. As the first step towards a more realistic modeling of real edges, the preceding perfect step edge is Gaussian-blurred to make it look “fuzzy”. At the same time, it is rotated 35° in the anti-clockwise direction (figure 5.7). The intention here is to see how accurately the orientation vector phase angles reflect the underlying orientation of image structures. The same filtering and vector field computation processes are then performed on this image.

5.2.1 Magnitude profiles

Vector magnitude profiles produced by filters of different sizes along row 135² are plotted on figure 5.8. These graphs are quite similar in shape to the corresponding ones in figure 5.2. As in figure 5.2, a magnitude trough – which occurs at column 131 – exists between two peaks for each curve. Strictly speaking, in this case, the vectors do not disappear altogether at that point. However, the vector magnitudes at these points turn out to be much smaller than the two peaks that flank each of the points. Also as in figure 5.2, these magnitude

²The approximate location of row 135 is indicated by a black horizontal line across figure 5.7.

Figure 5.8: Vector magnitude curves of figure 5.7 along row 135 ($\sigma=4\sqrt{2}, 8\sqrt{2}, 16, 16\sqrt{2}$)



troughs all coincide at one location. However, the vectors magnitudes in figure 5.8 are in general less than those in figure 5.2 belonging to the same scales. This observation is, of course, expected, since the edge involved in figure 5.7 is much fuzzier than that in figure 5.1.

5.2.2 Phase profiles

If we look at the phase angle profile in the vicinity of the fuzzy edge along the same row, a sudden 90° switch of phase angle can still be seen occurring between columns 131 and 132, the former being the location of magnitude troughs of all the curves.


Figure 5.9 shows the phase angles of vectors in the vicinity of the edge. Readings at the same locations but belonging to different scales are marked differently. In addition, phase angles of vectors at different locations derived from the smallest filter ($\sigma = \sqrt{2}$) fluctuate somewhat. Those derived from progressively larger filters exhibit greater and greater consistency.

In figure 5.9, phase angles of vectors at column 132 and above are not shown. In general, the phase angles have readings that are in the neighborhood of 35° due to the 90° switch that occurs when one crosses from over a straight edge between a bright and dark region. Other than the actual phase angle values, the distribution of phase angles around the theoretically correct 35° is very similar to that around 125° as depicted in figure 5.9.

5.2.3 Orientation estimation

With all these different phase angle values, a question inevitably arises: which one of them represents the best estimate of the edge's orientation – 125° from the horizon? Since we have suggested that the longer an orientation vector is, the more reliable its computed phase angle will be, a natural way to find the best estimate is to select the computed phase angle of the longest orientation vector in the vicinity of the edge across all scales. According to the data, the best orientation estimate along row 135 is 125.0078° , which occurs at columns 128 and 127. This value is produced with the filters with $\sigma = 16\sqrt{2}$. The corresponding vector lengths at these locations are however not impressive, being only 42.85402 and 34.07756. The longest vector, also derived from the same filters, measuring 114.86032, occurs at column 111, and its computed phase angle is 125.0098° . Taking factors such as discretization and computer arithmetic errors into account, I find the estimate, which is only 0.0098° off the mark, to be very good indeed.

Figure 5.9: Vector phase angles of figure 5.7 along row 135 ($\sigma = \sqrt{2}, 4\sqrt{2}, 8\sqrt{2}, 16$)






Figure 5.10: A noisy and fuzzy edge

5.3 The Effect of Noise on Orientation Vector Lengths

We have seen that, except for minor differences, the configuration of the MOF of a straight but fuzzy edge largely mirrors that of an ideal step edge. Such a similarity is expected – for both images, grey level changes monotonically in the direction perpendicular to the edge’s orientations. It therefore will be worthwhile to investigate the configuration of the MOF of an image that violates that condition. To that end, white Gaussian noise ($\mu=0$, $\sigma^2=500$) has been introduced to the fuzzy edge image (figure 5.10).

Figures 5.11 and 5.12 depict vector magnitude profiles corresponding to row 135 of figure 5.10. As can be seen, filters with σ ’s equal to $\sqrt{2}$ and $2\sqrt{2}$ produce orientation vector fields that are almost random. However, orientation vector fields derived from filters with σ ’s equal to $8\sqrt{2}$ and $16\sqrt{2}$ preserve their general configurations. In particular, curves with σ ’s equal to $4\sqrt{2}$, $8\sqrt{2}$, and $16\sqrt{2}$ have distinct magnitude troughs, which are located at columns 131, 131, and 132 respectively.

A closer comparison between the curve in figure 5.11 derived from the filter with $\sigma = \sqrt{2}$ and the corresponding curve in figure 5.8 reveals that in many cases, the vector magnitude at a particular location in figure 5.11 is larger than the corresponding one at the same location




Figure 5.11: Vector magnitude curves of figure 5.10 along row 135 ($\sigma = \sqrt{2}, 2\sqrt{2}, 4\sqrt{2}$)




Figure 5.12: Vector magnitude curves of figure 5.10 along row 135 ($\sigma = 8\sqrt{2}, 16\sqrt{2}$)

in figure 5.8. The same holds for the curves derived from the filter with $\sigma = 2\sqrt{2}$, and to a lesser extent for those derived from the filter with $\sigma = 4\sqrt{2}$. However, this does not imply that these longer vectors in figure 5.11 reflect the underlying image structures more accurately than the shorter vectors in figure 5.8 do. We see that the curves in figure 5.8 corresponding to small σ 's are diminutive versions of those corresponding to larger σ 's. By contrast, the shapes of magnitude curves corresponding to smaller σ 's in figure 5.11 are scarcely repeated by those corresponding to larger σ 's. The implication is that the noise in figure 5.10 contributes to the dissimilarity between magnitude curves of large σ 's and those of small σ 's. It follows that magnitude profile patterns on the curves corresponding to small σ 's that are not mirrored by those corresponding to larger σ 's can be ascribed to noise and duly disregarded. Conversely, profile patterns corresponding to small σ 's that are repeated at higher scales reflect the presence of *bona fide* image structures, even though the vectors corresponding to small σ 's might be comparatively small.

It is not clear how the two measures of phase angle reliability (vector length and vector field consistency across scales) can be numerically combined. While vector length can be clearly expressed in number, a quantification for the degree of consistency between vector field configurations over a neighborhood across several scales is yet to be formulated. For instance, it is not obvious which are the discrete scales (σ 's) whose vector fields need to be taken into consideration in the computation of this "degree of consistency". Even if a reasonable quantification had been formulated, it is not clear how vector length and vector field consistency across scales can be effectively combined.

5.4 Summary

- We see that if we march along the magnitude profile of an orientation vector field in the direction perpendicular to the orientation of an edge, we would encounter a magnitude trough, which is preceded and followed by a local vector magnitude peak. In addition, orientation vectors located near a straight, noiseless step edge but on the two different sides of the edge have an angular difference of 90° . However, it is premature to regard these phenomena as an edge signature.
- The reliability of the phase angle of an orientation vector should always be considered in the light of the vector's magnitude.

- As described in section 5.1.3, an angular ambiguity exists for every orientation vector. Consequently, if an orientation vector is examined alone, it is impossible to determine with absolute certainty the orientation of the underlying image structure that gives rise to this orientation vector.
- The configuration of the MOF of a straight, fuzzy (Gaussian-blurred) but noiseless edge is largely similar to that of a straight, noiseless step edge, since the grey-level changes across the edge are monotonic.
- The orientation filters can be used to measure the orientation of a straight, fuzzy, but noiseless edge with very good accuracy, although the reading should be taken from the bright side of the edge. The orientation vectors are rotated by 90° if the reading is taken from the other side. This confirms the prediction made in section 4.4.1.
- When noise is introduced, the magnitudes of vectors that derive from small orientation filters can be quite misleading. Tracing the vector field configurations of a neighborhood across scales might help to counter the adverse effects of noise. Consistency across scale can be used complementarily to length to arrive at a reliability measure for an orientation vector. However, problems of quantification and scale selection remain to be overcome.

Chapter 6

Curved Edges and Elongated Structures

There are two objectives pertaining to this chapter:

- To introduce the notion of *macro-geometry* of an object and to identify and account for the influence of the object's macro-geometry on its MOF.
- To qualitatively observe and account for the behavior of the MOF in the presence of curved edges of varying curvatures. In particular, the differences in the MOF resulted from a change of edge curvature will be noted and also explained in terms of macro-geometry. These differences will also be linked to the lack of provision for dealing with curved structure on the part of Coggins' method (section 4.4.2).

To these ends the images of a highly eccentric elliptical disc and a circular disc are used. Observations are made on the MOF of the elliptical disc along its major axis as well as its minor axis. By doing so, one can readily see the differences reflected on the MOF by a difference in edge curvature. Finally, the MOF of a circular disc is computed so that one can qualitatively compare it with the MOF of the elliptical disc.

6.1 The macro-geometry of an object

The macro-geometry of an object (or a part of an object) refers to the overall geometrical shape of the object (or a part of the object) in an observation window. I refrain from using the term *large-scale geometry* because some geometrical structures of an object exist at small




Figure 6.1: An elliptical disc centred at (127,127)

scales, but they nevertheless could in some cases be crucial in defining the overall shape of the object (e.g. corners). In the context of the MOF of an object, "macro-geometry" is used in cases where the configurations of the orientation vector fields in an observation window are shaped by the overall geometrical shape, rather than the edge(s), of the object.

6.2 An Elliptical Disc

For a start, the MOF of an elliptical disc (major axis=50 pixels, minor axis=15 pixels, center at (127,127)) as shown in figure 6.1 is computed.

6.2.1 Vector profiles along the minor axis

Figure 6.2 shows the vector magnitude profiles of figure 6.1 along row 128 as produced by filters with $\sigma = 2, 2\sqrt{2}, 8$ and $8\sqrt{2}$. Table 6.1 shows the locations of vector magnitude troughs and peaks along row 128 of figure 6.1. Readings are taken for $\sigma = 1$ through $16\sqrt{2}$.

Magnitude peaks

Each of the curves corresponding to σ 's ranging from 1 to 4 has four magnitude peaks, with one corresponding to each side of the two edges. As σ increases, the two interior peaks move

vector length




Figure 6.2: Vector magnitude curves of figure 6.1 along row 127 ($\sigma = 2, 2\sqrt{2}, 8, 8\sqrt{2}$)

σ	Magnitude troughs	Magnitude peaks
1	122,132	121,123,131,133
$\sqrt{2}$	122,132	121,123,131,133
2	122,132	121,123,131,133
$2\sqrt{2}$	122,132	120,124,130,134
4	122,132	120,126,128,134
$4\sqrt{2}$	122,132	118,127,136
8	121,133	116,127,138
$8\sqrt{2}$	119,135	113,127,141
16	117,137	108,127,146
$16\sqrt{2}$	114,140	101,127,153

Table 6.1: Locations (columns) of magnitude troughs and peaks along row 127 of figure 6.1

closer and closer together until they merge at $\sigma = 4\sqrt{2}$. Meanwhile, the two external peaks move farther and farther away from the edges of the ellipse. Tracing across the magnitude peaks across scales, it can be observed that the value of external peaks assume a maximum at $\sigma=4$ (length: 132.6717), and the interior peaks assume a maximum at $\sigma=8$ (length: 254.9616).

What distinguishes these two maxima is the difference of their sizes. This difference can be ascribed to the fact that the two maxima reflect underlying image structures that are quite different. The former indicates the existence of an edge in its vicinity; whereas the latter reflects the presence of an elongated structure centered at that location. The difference in magnitude can be intuitively explained from the observation that the latter maximum derives its value from the spatial support provided by the disc's interior between its two edges. This happens when the filters used are of the size ($\sigma = 8$) that can capture the overall geometrical shape of the disc. As such, the vector field configuration at this scale is shaped by the macro-geometry – not the edges – of the disc. In addition, in this instance the macro-geometry of the disc is such that it allows the response of one orientation filter to decisively prevail over that of its antagonistic filter (section 5.1). In contrast, the former maximum is shaped by the edge of the disc.

Since both edges and elongated image structures give rise to vector magnitude peaks, a

question naturally arises: how could one distinguish between a vector magnitude peak caused by an edge and one caused by the presence of an elongated image structure? Tracing vector magnitude peaks across scales seems to provide the answer. One can start from a vector magnitude peak of the lowest scale and try to search for the closest vector magnitude peak belonging to the next higher scale. The process is repeated until the highest scale is reached. At any one point, if a vector magnitude peak corresponds to two peaks belonging to the scale immediately below the one on which it resides, one can conclude that the magnitude peak reflects the presence of an elongated structure. Otherwise, we can conclude that the peak indicates the existence of an edge.

The method described above does not constitute an algorithm for distinguishing these two types of magnitude peaks. For instance, it is not yet known how well the method would perform in the presence of noise or arbitrarily complex image structures.

Magnitude troughs

Two magnitude troughs can be found at columns 122 and 132 for magnitude curves corresponding to $\sigma=1$ to $4\sqrt{2}$ (table 6.1). Magnitude troughs corresponding to σ 's larger than $4\sqrt{2}$ move away from the center of the elliptical disc as the value of σ increases.

The edges of the ellipse in figure 6.1 occur between columns 122 and 123, and between columns 132 and 133. As can be seen, the presence of the edges are accurately captured by the curves with corresponding σ values ranging from 1 to $4\sqrt{2}$ as magnitude troughs. Vector magnitude troughs for curves corresponding to larger σ 's match the edges of the ellipse less well. This is expected, since filters having larger σ 's are larger than the width of the entire ellipse. Thus, the filter responses are affected by the macro-geometry of the ellipse and not just the edge.

The phase angles

Table 6.2¹ shows the computed phase angles of the vectors whose magnitudes are magnitude peaks. There is really no surprise here – phase angles that correspond to local magnitude peaks that are outside the ellipse have readings that are close to 0° or 180° (the two quantities

¹Numbers inside parentheses indicate the columns (or rows) at which the peaks occur. For entries with four columns only the middle two correspond to positions inside the disc. For those with three columns only the middle ones correspond to positions inside the disc. The same conventions apply to tables 6.4.

σ	Phase angles (in $^{\circ}$)			
1	0.6960695(121)	90.71532(123)	89.28468(131)	179.3039(133)
$\sqrt{2}$	0.4458724(121)	90.47514(123)	89.52486(131)	179.5541(133)
2	0.2833135(121)	90.31390(123)	89.68610(131)	179.7167(133)
$2\sqrt{2}$	0.1439175(120)	90.13898(124)	89.86102(130)	179.8561(134)
4	0.09011207(120)	90.01866(126)	89.98134(128)	179.9099(134)
$4\sqrt{2}$	0.05985141(118)	90.00000(127)	179.9402(136)	N.A.
8	0.1014575(116)	90.00000(127)	179.8985(138)	N.A.
$8\sqrt{2}$	0.1877978(113)	90.00000(127)	179.8122(141)	N.A.
16	0.3264695(108)	90.00000(127)	179.6735(146)	N.A.
$16\sqrt{2}$	0.5417737(101)	90.00000(127)	179.4582(153)	N.A.

Table 6.2: Phase angles corresponding to magnitude peaks along row 127 of figure 6.1

being identical in our computations). As also expected, the interior peaks have readings equal or close to 90° (see section 5.1.2). These observations show that, in the presence of a slightly curved edge, the edge itself still constitutes a major factor in shaping the resulting vector field configurations in its vicinity. These vector field configurations are therefore little different from those of a straight edge.

6.2.2 Vector profiles along the major axis

In examining magnitude profiles along the ellipse's major axis, I have been able to identify two factors that give rise to local vector magnitude maxima – edge and macro-geometry. The angular profile is rather uninteresting – the data yield nothing more than what has been discovered in section 5.1.2. The angular profile of vectors along the major axis of the elliptical disc should be more interesting due to the high curvatures at points along the edge there.

Magnitude peaks and troughs

Figure 6.3 shows the vector magnitude curves of figure 6.1 along column 127 corresponding to four selected scales. A major difference between the curves in this figure as compared to those in figure 6.2 is that the vector magnitude troughs of the curves are very poorly defined.




Figure 6.3: Vector magnitude curves of figure 6.1 along column 127 ($\sigma = 2, 2\sqrt{2}, 8, 8\sqrt{2}$)

σ	troughs (col. 127)	peaks (col. 127)
1	78,176	77,79,175,177
$\sqrt{2}$	78,176	77,80,174,177
2	78,176	77,81,173,177
$2\sqrt{2}$	N.A.	84,170
4	N.A.	89,165
$4\sqrt{2}$	N.A.	95,159
8	N.A.	127
$8\sqrt{2}$	N.A.	127
16	N.A.	127
$16\sqrt{2}$	N.A.	127

Table 6.3: Locations (rows) of magnitude troughs and peaks along columns 127 of figure 6.1

It should be noted that the local magnitude minima located near row 127 of the magnitude profiles corresponding to $\sigma=2$ and $2\sqrt{2}$ are not magnitude troughs, since the magnitude peaks on their two sides are derived from different edges.

It can also be observed that pairs of local magnitude peaks and troughs corresponding to the edges of the ellipse move closer and closer to each other as scale increases until they finally merge at $\sigma = 2\sqrt{2}$. This phenomenon is clearly reflected on table 6.3. Magnitude troughs are absent from all entries following the entry corresponding to $\sigma=2$. At the same time, the number of magnitude peaks per entry is reduced from 4 to 2, implying that the outer magnitude peaks have merged with their corresponding magnitude troughs.

An explanation for the poor definition of magnitude peaks and troughs

Figure 6.4 shows three pairs of antagonistic filters of different sizes superimposed near the tip of a bright, elongated structure. If we compare figure 6.4 with figure 6.5, which depicts a straight edge, it is easy to deduce that the combined outputs of the antagonistic filter pairs in figure 6.4 are larger than that produced by the corresponding filter pairs in figure 6.5. This is because the responses to individual orientation filters in an antagonistic filter pair are more or less equal to each other in figure 6.5. However, the same can not be said of figure 6.4. The resultant orientation vectors at the centers of the filters in figure 6.4 are




Figure 6.4: Filters of different sizes superimposed on an elongated structure




Figure 6.5: Filters of different sizes superimposed on a straight edge

σ	Phase angles (in $^{\circ}$)			
1	90.00000(77)	90.00000(79)	90.00000(175)	90.00000(177)
$\sqrt{2}$	90.00000(77)	90.00000(80)	90.00000(174)	90.00000(177)
2	90.00000(77)	90.00000(81)	90.00000(173)	90.00000(177)
$2\sqrt{2}$	90.00000(84)	90.00000(170)	N.A.	N.A.
4	90.00000(89)	90.00000(165)	N.A.	N.A.
$4\sqrt{2}$	90.00000(95)	90.00000(159)	N.A.	N.A.
8	90.00000(127)	N.A.	N.A.	N.A.
$8\sqrt{2}$	90.00000(127)	N.A.	N.A.	N.A.
16	90.00000(127)	N.A.	N.A.	N.A.
$16\sqrt{2}$	90.00000(127)	N.A.	N.A.	N.A.

Table 6.4: Phase angles corresponding to magnitude peaks along column 127 of figure 6.1 (see footnote on page 37 for explanations)

therefore going to be longer than their counterparts in figure 6.5. The implication is that at a particular scale, the magnitude troughs are more well defined in the case of a straight edge than in that of a highly curved edge.

It is also evident in figure 6.4 that the combined responses of the antagonistic filter pairs increase as the size of the filters increases. This means that the corresponding magnitude troughs become more and more poorly defined as scale increases until they finally disappear. However, a magnitude trough can only disappear via a merger between itself and a magnitude peak. This accounts for the second observation described above.

The phase angles

If we compare table 6.2 with table 6.4, we can see that the switching of vectors by 90° that occurs when one passes from outside the elliptical disc to inside it does not take place at all when the view is taken along the major axis of the disc. Vector magnitude peaks situated both inside and outside the disc has phase angles of 90° . While vector magnitude peaks situated outside the disc are expected to have phase angles measuring 90° regardless of the curvature of its edge near which the peaks occur, the high curvature of the edge seems to be responsible for the magnitude peaks within the elliptical edge registering phase angles of




Figure 6.6: Filter superimposed on a highly curved edge

90°.

This observation is portrayed in figure 6.8². Each patch in figure 6.8 having uniform grey-level represents a pixel on the original vector field. The relative intensities of these patches are to be interpreted as relative lengths of the orientation vectors. The little thin rod within each patch indicates the phase angle of the vector at that location. Neighborhoods with low edge curvatures exhibits large angular differences between vectors inside and outside the ellipse. The angular differences become progressively ill-defined in neighborhoods having progressively high edge curvatures. The angular difference disappears altogether in the neighborhood that contains a curvature maximum of the edge.

An account for the lack of angular differences

Figures 6.6 and 6.7 are a couple of schematic diagrams showing the effect curvature has on the phase angles of magnitude peaks. In these figures, only X-orientation filters (section 3.4) are used, since the elliptical disc is symmetrical along the its major axis. Basically, there is a tug-of-war between the bright (positive) horizontal vanes and the dark (negative) vertical vanes of a X-orientation filter. If the underlying image has a highly curved edge with respect to the size of the filter applied, as in figure 6.6, then at locations inside but near the edge, the bright vanes might not cover enough area to offset the negative effect of the dark vanes, since they have considerable negative effect because the lower vane and a small part

²The black dotted line on the figure marks the approximate locations of the object's edge elements.

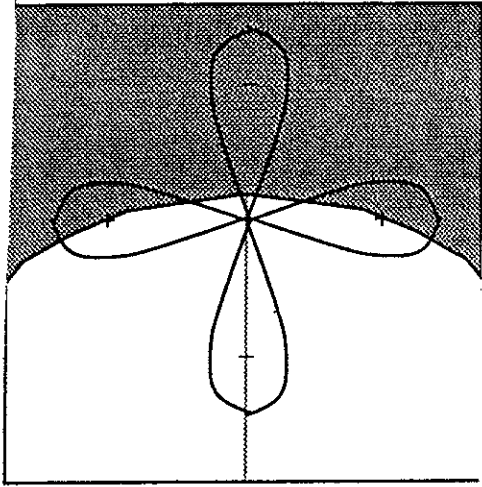


Figure 6.7: Filter superimposed on a slightly curved edge

of the upper vane are within the bright region of the underlying image. If that is indeed the case, then according to the way orientation vectors are computed (section 3.2), the resultant phase vector at that point would be 90° . If the underlying image has a flat curve relative to the size of the filters applied (figure 6.7), then the bright vanes would get enough support from the image's bright region and the computed phase angles along the major axis would be 0° or 180° . As larger filters are used, the edge concerned has to have smaller and smaller curvature for the computed phase angles to measure 0° or 180° .

Comments

The macro-geometry of the elliptical disc is responsible for the phenomena as observed in this section (6.2.2). In other words, the increased elongation, which goes in hand in hand with increased curvature, of the disc as one moves towards the edge curvature maximum corresponding to the disc's major axis primarily causes the poor definition of magnitude peaks and troughs, and the lack of angular differences between orientation vectors on opposite sides of the disc's edge.

6.2.3 Comments

Foregoing observations have shown that angular differences between vectors on one side of an edge and those on the other side and in the vicinity of the edge can vary considerably. In particular, such angular differences and magnitude troughs are ill-defined at locations with




Figure 6.8: Vector field configuration ($\sigma=1$) in the vicinity of an edge curvature maximum in figure 6.1

high edge curvatures. This observation eliminates the suggestion that a sudden switch of vector phase angles together with the presence of a sharp magnitude trough along a certain line of observation can be used as an indication of an edge.

It is important to note that the foregoing peculiar behavior of the orientation vector field corresponding to $\sigma = 1$ reflects the macro-geometry – and not the presence of an edge – of the elliptical disc in the neighborhood of the edge curvature maximum. There are two implications to this statement. First, the macro-geometry of an object (or a part of it) is not captured exclusively by large filters. It seems that the size of the filters needed to capture the macro-geometry of a part of an object decreases as the edge curvature of that part of the object increases. The second implication is that theoretically such a peculiarity in vector field configuration can be avoided by using filters that are small enough that the resulting orientation vector field would be shaped by the edge, rather than the macro-geometry in the neighborhood of the edge curvature maximum. As long as the edge curvature is not infinite, the suggestion above is valid. In real life, however, discretization errors prevent us from constructing arbitrarily small filters.

Finally, as observed in section 6.2.2, for an orientation vector field at a particular scale different edge vicinities in figure 6.8 are subject to different amounts of influence from the macro-geometry as well as the edge of the elliptical disc. This makes uniform interpretation of orientation vectors a difficult and error-prone process. This shortcoming reflects the lack of provision for handling curved image structures in Coggins' method, thus proving Zucker right that orientation determination of curved image structures has to be specially dealt with (section 4.4.2). A more desirable way of computing is to control the influence of the two factors by adaptively adjusting the size of the filters used according to the nature of macro-geometry at particular vicinities. At this point, it is not at all clear how this might be accomplished, nor is it known whether such an algorithm exists at all. The problem does not go away because the size of filters needed to escape the influence of macro-geometry may get arbitrarily small. The use of filters incorporating higher order derivatives may provide a new approach to solving this problem.

6.3 Summary

- In the presence of a slightly curved edge, the edge itself still constitutes a major factor





Figure 6.9: A circular disc (centred at (127,127))

σ	magnitude troughs	magnitude peaks
1	112,142	111,113,141,143
$\sqrt{2}$	112,142	111,113,141,143
2	112,142	111,114,140,143
$2\sqrt{2}$	112,142	110,114,140,144
4	113,141	110,115,139,144
$4\sqrt{2}$	114,140	110,117,137,144
8	116,138	109,119,135,145
$8\sqrt{2}$	N.A.	109,145
16	N.A.	108,146
$16\sqrt{2}$	N.A.	106,148

Table 6.5: Locations (columns) of magnitude troughs and peaks along row 127 of figure 6.9






Figure 6.11: Orientation vector field configuration near the edge of the circular disc in figure 6.9 at $\sigma = 8\sqrt{2}$ (Dotted line explained on page 42)

in shaping the resulting vector field configurations in its vicinity. These vector field configurations are therefore little different from those of a straight edge.

- High edge curvature causes the magnitude trough on the magnitude profile curve to be less well defined. Similarly, it diminishes the angular differences between orientation vectors located on the opposite sides of an edge. This observation rules out the suggestion that a sudden switch of vector phase angles together with the presence of a sharp magnitude trough along a certain line of observation can be used as an indication of an edge. These observations indicate that the macro-geometry of the elliptical disc now plays an important part in the dispositions of the relevant orientation vector fields.
- Besides the existence of an edge, a peak on a magnitude profile curve can also arise when the orientation vector field concerned captures the macro-geometry of an object.
- The size of the filters needed to capture the macro-geometry of a part of an object decrease as the edge curvature of that part of the object increases. This implies that the above quantity is variable in an image where different parts have different edge curvatures. This means that uniform interpretation of the phase angles and magnitudes of orientation vectors on a orientation vector field is error-prone. In these cases, uniform processing, which is an important attribute of Coggins' method, becomes less applicable. However, at this point, the possibility of formulating an effective adaptive algorithm is still a research question.

Chapter 7

Corners

In the foregoing experiments it seems that angular differences between vectors that are on the two different sides of an edge in the neighborhood of that edge get smaller as the edge curvature progressively increases. It will be interesting to observe the configuration of the orientation vector field in the neighborhood of a point whose edge curvature is infinite, i.e. a corner. More important, it will be educational to see whether the prediction made in section 4.4.3 can in fact be verified.

In this experiment, corners with openings measuring 30° and 150° are used. They are shown in figures 7.1 and 7.2 respectively. The opening of the former forms an acute angle while the latter forms an obtuse angle. Because of this the image structure from which the corner in figure 7.1 derives is more elongated than the corresponding image structure in figure 7.2. Figures 7.3 and 7.4 respectively show the vector field configurations at $\sigma=2$ in the neighborhoods of the corners in figures 7.1 and 7.2. Vector fields belonging to other σ 's are not shown, since they do not offer additional insights.

7.1 The 30° corner

In figure 7.3 vectors belonging to border pixels that are relatively far away from the corner exhibit the now familiar configuration that suggests the presence of a (relatively) straight edge. As in the case of figure 6.8 the angular differences between vectors located on the two sides of the edge essentially disappear at the few pixels where the corner can be said to be located. The vector magnitudes at these pixels are relatively large, which seems to contradict one's intuitive prediction that the orientation preference at this locality is ill-defined.




Figure 7.1: A 30° corner




Figure 7.2: A 150° corner




Figure 7.3: Orientation vector field configuration for figure 7.1 at $\sigma=2$ (Dotted line explained on page 42)




Figure 7.4: Orientation vector field configuration for figure 7.2 at $\sigma=2$ (Dotted line explained on page 42)




Figure 7.5: Thought experiment on an acute corner

Pixel location	Magnitude	Phase angle
(126,128)	117.172195	174.862885°
(126,129)	150.207001	6.529789°
(126,130)	193.218887	13.328004°
(126,131)	216.458984	15.303263°
(126,132)	215.584061	15.430973°
(127,128)	128.582443	22.607244°
(127,129)	144.372299	25.519289°
(127,130)	140.869064	27.067478°
(127,131)	121.525391	27.554371°
(127,132)	96.097267	27.075390°
(128,128)	100.564995	46.707012°
(128,129)	105.356613	55.033375°
(128,130)	101.833786	63.129478°
(128,131)	98.585449	70.896576°
(128,132)	98.990242	77.600044°

Table 7.1: Orientation vectors in the immediate neighborhood of the corner in figure 7.1

Pixel location	Magnitude	Phase angle
(126,125)	38.956779	61.071079°
(126,126)	24.129566	146.732773°
(126,127)	52.344913	158.478195°
(126,128)	73.119606	171.759460°
(126,129)	99.072830	177.547119°
(127,125)	135.296265	61.964439°
(127,126)	118.647255	64.859787°
(127,127)	84.846024	70.551071°
(127,128)	46.269466	75.744041°
(127,129)	14.590264	74.742706°
(128,125)	125.460617	63.606365°
(128,126)	147.371719	68.202316°
(128,127)	154.719498	74.731964°
(128,128)	145.083893	81.659805°
(128,129)	130.141205	86.347122°

Table 7.2: Orientation vectors in the immediate neighborhood of the corner in figure 7.2




Figure 7.6: Thought experiment on an obtuse corner

Table 7.1 shows the magnitudes and phase angles of orientation vectors over a 5x3 (pixels) region containing the corner, which I designate to be located at (127,130). Note that in a discrete (and by implication aliased), grey-level image such as figure 7.1, it is not always easy to pin-point the exact location of a fine image structure such as a corner without any ambiguity. With this disclaimer thus made, I have failed to observe any distinguishing quality pertaining to pixel (127,130) that sets it apart from the others.

Among the vectors recorded in table 7.1, the one at pixel (126,131) is the longest. Since I have attached some importance to the length of a vector (section 5.1), it will be worthwhile to account for the length of the vector at (126,131). Figure 7.5 shows an antagonistic pair of orientation filters with their centers superimposed upon the corner. This schematic diagram is necessarily a simplification because the filter pair depicted are not the only antagonistic pair that has influence on the final computed orientation vector at the corner. This qualification notwithstanding, they are the most influential pair in determining the final disposition of the vector. It can be easily inferred from figure 7.5 that the outcome of the competition between the antagonistic filter pair heavily favors the "horizontal" orientation filter, since the "vertical" orientation filter does not get much spatial support from the bright region in the figure. Consequently, the phase angle of the final orientation vector is more or less in the same direction as the preferred orientation of the "horizontal" orientation filter. This accounts for the lengths and phase angles of the orientation vectors in the vicinity of the corner.

7.2 The 150° corner

Table 7.2 shows the magnitudes and phase angles of vectors computed with $\sigma = 2$ (figure 7.4) in a 5x3 neighborhood surrounding the corner, which is located at (127,127). As in the case of figure 7.3, there is nothing distinguishing about the orientation vector at the corner. Moreover, the vectors depicted in table 7.2 are in general shorter than those depicted in table 7.1, although in both cases, the vectors are taken from small neighborhoods containing the corners. Figure 7.6 is a self-explanatory diagram that accounts for the foregoing observation. Unlike figure 7.5, the vertical filter in figure 7.6 derives considerable spatial support from the bright part of the figure 7.2 and therefore has been more effective in canceling the response of the horizontal filter. As a result, the orientation vectors in figure 7.4 are in general shorter

than those figure 7.3.

7.3 Comments

The final disposition of the orientation vectors at the two corners, like the one at the edge curvature maximum in section 6.2, is shaped by the macro-geometry of the bright triangles in the neighborhoods of the two corners. I suggested there that, at least theoretically, the influence of the macro-geometry of the elliptical disc in the neighborhood of the edge curvature maximum can be effectively diminished by using small enough filters. I later pointed out that this is infeasible, as discretization errors effectively put a lower bound on the size of the filters that can be meaningfully constructed. In the case of a corner, however, it is not just a practical, but a theoretical impossibility to construct filters that are small enough to avoid the influence of the macro-geometry surrounding the corner, since the spatial extent at the corner is infinitely small. As a result, even orientation filters that are infinitely small can capture the macro-geometry around the corner.

Because the vector field configuration (as produced by the smallest orientation filters that can be meaningfully constructed) in the neighborhood of a corner is largely similar to that of a point of very high (but not infinite) edge curvature (compare figure 7.3 with figure 6.8), it follows that a corner is indistinguishable from a point of high edge curvature on an orientation vector field. In other words, the MOF is not very good in preserving corner information. This verifies the prediction made in section 4.4.3.

Chapter 8

Summary and Evaluation

The objective of this thesis has been to conduct a preliminary investigation on the strengths and limitations of the MOF, which was proposed by Coggins as a multiscale representation of image structures. In order to provide a basis on which the MOF (and the computational method leading to the MOF) can be objectively assessed, a number of related computational frameworks and algorithms were summarized. After that, observations were made on the MOF of several image structures through 1-D as well as 2-D windows. I also attempted to account for all phenomena whose origins were not immediately obvious. In the process I introduced a number of concepts, which will be summarized below.

8.1 A final summary

8.1.1 Observations

1. If one marches along the magnitude profile of an orientation vector field in the direction perpendicular to the orientation of an edge, we would encounter a magnitude trough, which is preceded and followed by a local vector magnitude peak. In addition, orientation vectors located near a straight, noiseless step edge but on the two different sides of the edge have an angular difference of 90° . A similar configuration can also be observed on the MOF of a straight, Gaussian-blurred, but noiseless edge. In the case of a noisy image, some randomness is introduced onto the orientation vector fields. However, the degree of randomness decreases dramatically as the size of the orientation filters used increases.

2. An interpretive ambiguity exists for the phase angle of every orientation vector. Consequently, if an orientation vector is examined alone, it is impossible to determine with absolute certainty the orientation of the underlying image structure that gives rise to this orientation vector.
3. Whereas the orientation vectors of a slightly curved edge are little different from those of a straight edge, the magnitude trough on the the magnitude profile curve of a highly curved edge are rather poorly defined. The angular differences between orientation vectors located on the opposite sides of an edge are similarly diminished.
4. One would expect the preferred orientations of orientation vectors at a corner and its immediate vicinity are very poorly defined or even undefined. However, it had been observed on the MOF that the quantities described above are rather well defined.

8.1.2 Theoretical findings and evaluations

1. Other factors remaining constant, the phase angle of a long orientation vector is considered to be more *reliable* (see section 5.1.2) than that of a short one.
2. One factor that affects the validity of the above statement is the noise level in an image. When noise is introduced, the orientation vector magnitudes that derive from small orientation filters can be quite misleading. Tracing the vector field configurations of a neighborhood across scales might help countering decreased validity of vector length as a reliability measure. It was proposed that consistency across scale, on top of its length, be used to arrive at a reliability measure for an orientation vector. However, problems of quantification and scale selection remain to be overcome.
3. The MOF can be used to measure the orientation of a straight, fuzzy, but noiseless edge with very good accuracy, although the reading should be taken from the bright side of the edge.
4. The idea of the influence of *macro-geometry* was introduced to account for the second and third observations noted in point 3 of section 8.1.1. The observations indicate that the macro-geometry of the elliptical disc played an important part in the dispositions of the relevant orientation vector fields.

5. Besides the existence of an edge, a peak on a magnitude profile curve can also arise when the orientation vector field concerned captures the macro-geometry of an object.
6. The size of the filters needed to capture the macro-geometry of a part of an object decreases as the edge curvature of that part of the object increases. Thus the above quantity is variable in an image where different parts have different edge curvatures. Therefore different parts of the orientation vector field derived from a particular σ are subject to different degrees of influence from the object's macro-geometry. This fact undermines the effectiveness of uniformly interpreting the phase angles and magnitudes of orientation vectors. As a result, the MOF's ability to perform orientation estimation at highly curved edges is compromised.
7. Due to the influence of macro-geometry, orientation vectors at corners have relatively large magnitudes. Common sense would be violated if the magnitudes of these orientation vectors were to be interpreted as point 1 in this section. Worse, designing a filter that would avoid the influence of the corner's macro-geometry is a theoretical impossibility, as a corner has by definition an infinitely small spatial extent. I am therefore forced to conclude that the MOF is not good at extracting corners.

8.2 Future research directions

The main problems associated with the MOF are summarized as follows:

1. There is an inherent interpretive ambiguity of the phase angle of every orientation vector.
2. A method of scale-tracing for the MOF so that the consistency of the vector field configurations in a neighborhood can be quantitatively measured is yet to be devised.
3. Information present in the MOF still can not used adaptively across several scales so that the influence of macro-geometry can be held at a constant level.
4. The MOF is unsuitable for extracting corners.

These four problems will have to be solved or compensated for by other means before the MOF can be effectively used to perform multiscale image segmentation. Problem 1 can be

avoided altogether by using single-lobe instead of double-lobe orientation filters (see figure 3.1). This modification has in fact been undertaken by Fritsch ([3]).

Scale-tracing had been attempted by Lifshitz and Gauch ([4], [10]). At this point, it is not clear how their methods can be adapted for use in conjunction with the MOF. A future research direction clearly involves a thorough investigation of these two methods so that their relevance to the MOF can at least be ascertained.

Appropriate adaptive use of information across several scales is a very difficult problem for which one can provide a satisfactory solution. The first step towards the realization of a solution is the determination of the "appropriate" scale(s) at which an image structure can be meaningfully analyzed. Back *et al* ([1]) attempted to solve this problem for 1-D edges. A promising future research direction is to generalize this approach to encompass other 2-D image structures.

Point 4 forces me to agree with Hsieh ([6]) that corners have to be explicitly extracted in MOF. Future work therefore includes finding a way to integrate the output of a corner detecting algorithm with the MOF.

BIBLIOGRAPHY

- [1] S. Back, H. Neumann, H.S. Stiehl, "On Scale-space Edge Detection in Computed Tomograms", *Mustererkennung 1989 (Proc. 11th DAGMS Symposium, October 1989)*, Springer-Verlag, 1989.
- [2] J.M. Coggins, "A Multiscale Description of Image Structure for Segmentation of Biomedical Images", *Proc. First Conference on Visualization in Biomedical Computing*, Atlanta, Georgia, May 22-25, 1990.
- [3] Dan Fritsch, private communications, Biomedical Engineering, University of North Carolina at Chapel Hill.
- [4] J.M. Gauch, "The Multiresolution Intensity Axis of Symmetry and its Application to Image Segmentation", *Tech. Report TR89-047, Computer Science Department, University of North Carolina at Chapel Hill*.
- [5] S. Grossberg and E. Mingolla, "Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations", *Perception and Psychophysics*, vol. **38**(2):141-171, 1985.
- [6] C.-H. Hsieh, "A Connectionist Algorithm for Image Segmentation", *Tech. Report TR89-008, Computer Science Department, University of North Carolina at Chapel Hill*, February 1989.
- [7] D.H. Hubel, *Eye, Brain, and Vision*, Scientific American Library Series #22, W.H. Freeman & Company, New York, 1988.
- [8] R.A. Hummel and S.W. Zucker, "On the foundation of relaxation labelling processes", *IEEE Trans. Pattern Analysis and Machine Intelligence*, PAMI-5:267-287, 1983.

- [9] M. Kass and A.P. Witkin, "Analyzing Oriented Patterns", *Proc. 9th IJCAI*, pp. 944-952, Los Angeles, California, August 18-23, 1985.
- [10] L.M. Lifshitz, "Image Segmentation via Multiresolution Extrema Following", *Tech. Report TR87-012, Computer Science Department, University of North Carolina at Chapel Hill*.
- [11] K.-C. Low and J.M. Coggins, "Multiscale Vector Fields for Image Pattern Recognition", *SPIE Vol. 1192 Intelligent Robots and Computer Vision VIII: Algorithms and Techniques*, 1989.
- [12] H. Neumann, "Extraction of Image Domain Primitives with a Network of Competitive/Cooperative Processes", *Proc. German Workshop on Artificial Intelligence 1988 (12.Fachtagung Künstliche Intelligenz)*, Schloß Ehringerfeld, 19.9-23.9.1988.
- [13] H. Neumann and H.S. Stiehl, "A Competitive/Cooperative (Artificial Neural) Network Approach to the Extraction of N-th Order Junctions", *Mustererkennung 1989 (Proc. 11th DAGMS Symposium, October 1989)*, Springer-Verlag, 1989.
- [14] P. Parent and S.W. Zucker, "Curvature Consistency and Curvature Detection", *Technical Report 85-12R, Computer Vision and Robotics Lab., McGill University, Montreal*, May 1985.
- [15] A. Ravishankar Rao and R. Jain, "The analysis of oriented textures through phase portraits", *Proc. First Conference on Visualization in Biomedical Computing*, Atlanta, Georgia, May 22-25, 1990.
- [16] S.W. Zucker, "Early Orientation Selection: Tangent Fields and the Dimensionality of Their Support", *Computer Vision, Graphics, and Image Processing*, vol. 32:74-103, 1985.