

A CURVATURE SENSITIVE FILTER AND ITS APPLICATION IN MICROFOSSIL IMAGE CHARACTERISATION

J. P. Oakley and R. T. Shann

Department of Electrical Engineering
University of Manchester
Manchester M13 9PL U.K.

Abstract

A new class of oriented, curvature sensitive filters are introduced. These filters provide a low-level detection facility for noisy curves without a prior edge extraction stage. The application of these filters to the detection of Carboniferous Foraminifers (a type of microfossil found in plane rock sections) is described. A symbolic representation of the detected curves is stored in a database which is then queried to recover the required structures. We show that the curves identified by the filter correspond to salient features of the microfossil evidence in the image.

1 Introduction

We report here on a research project on techniques for content addressing of image databases. The aim is to provide the equivalent of "free-text" document retrieval for real image data. This is achieved by extracting some representation of image content using image analysis and then storing this information in a database together with the raw image data. The user interacts with the system to define a query set of interesting images which can be inspected on-screen.

One important application which is under investigation for this system is the dating of rock samples using microfossil evidence (a procedure known as biostratigraphy). Work has been reported on automatic or semi-automatic biostratigraphy of 3-D microfossil samples [1] but the systems described rely on graphical input from the user rather than image analysis. Here we are looking at a different type of microfossil (the Carboniferous Foraminifera), which because their hardness is similar to that

of the surrounding rock can only be viewed as 2-D sections from thin sections. A reasonably clear example of this type of microfossil is shown in figure 2 and a less clear example in figure 3. It can be seen that the actual microfossils can be distinguished by eye from the surrounding rock because of a pronounced geometrical structure. The level of visual noise is quite high, and this factor is inherent in the nature of the material, not an artifact of the imaging technique.

We have looked at the Canny-type edge maps for these images but the amount of data is huge and very few of the edge segments relate to significant object features. This has led us to look towards robust methods for detecting noisy curves, which are the predominant visual features in this application.

Low level schemes for curve detection have been reported [2] [3] but the work seems to have been motivated by physiological considerations rather than by a need for robust curve detection. We feel that any scheme which relies on bottom-up grouping of edge information is doomed to failure on images containing this level of visual noise. Instead we take a top-down approach and fit a local model of a curved boundary directly to the grey-level image itself. This is achieved by a process based on linear filtering of the image data with a specialised filter.

This filter is a variation of that reported in [4] as a method for precise localisation of curved boundaries, although the actual problem addressed here is quite different.

In an effort to quantify the performance of these filters we will show, using actual microfossil images:

- (a) In section 3.1, how the detection performance of the operator varies with σ_β (the filter arc length);
- (b) In section 3.2, that the curvature correction (i.e. the Newton step in κ) is self consistent;
- (c) In section 3.3, that the curvature of the filter has a strong influence on which curves are detected and that the curvature selectivity can be increased by placing a restriction on the size of the Newton step in κ ;
and
- (d) In section 3.4, that a simple Hough-like centre binning algorithm provides a reliable method for detecting microfossil structures.

2 Technical Approach

2.1 The Filter

The basic approach is to firstly filter the image with a specialised linear filter and then to detect local intensity extrema (minima and maxima) in the filtered image. There are typically about 200 of these local extrema (which we call "hits") in each image; this is a much more manageable amount of data than an edge map. The hits are then subject to further selection based on:

- (a) parametric information and
- (b) intensity ranking.

The parametric information quantifies the behaviour of the hit as the filter parameters are varied (this information is obtained from partial derivatives).

The filters are highly direction sensitive and so the whole process is repeated for several different orientations.

The filter used for detecting curved edges is illustrated in fig (1). The filter kernel is generated starting from a simple Gaussian shape which is elongated and then curved and finally differentiated in the direction perpendicular to the elongation so as to give an curved edge sensitive detector. The elongated curved Gaussian $G(x,y)$ is the product of two terms:

$$e^{-\frac{1}{2} \left[\frac{((\kappa x - \cos \theta)^2 + (\kappa y - \sin \theta)^2)^{\frac{1}{2}} - 1}{\sigma_{\alpha} \kappa} \right]^2}$$

and

$$e^{-\left[\frac{x^2 + y^2}{2\sigma_{\beta}^2} \right]}$$

Here θ is the angle of the filter and κ is the curvature. The first of these terms decays (via σ_{α}) with distance from the circle centred on $\cos(\theta)/\kappa, \sin(\theta)/\kappa$. The second term decays more gradually (via σ_{β}) with distance from the filter centre. The directional derivative $K(x,y)$ is given in terms of the derivatives with respect to the spatial directions x and y by

$$K(x,y) = - \{ G_x \cos \theta + G_y \sin \theta \}$$

where the minus signs cause the sense of differentiation to be radially outwards at the filter origin. This is our filter kernel (see fig 1). Derivatives of this expression with respect to x, y and κ up to second order were calculated using the symbolic algebra package Macsyma. By the usual argument [5] the derivatives of the filtered image with respect to these parameters can be obtained by filtering the image with the differentiated kernels. The filtered images are computed at every point on the pixel grid using an FFT convolution routine. At this stage we could detect local maxima by examining adjacent pixels. However, this is not a wholly reliable method, marking too few or too many points at places where the image is changing rapidly. Instead by computing the Newton step (a term from optimisation theory - see for example [6]) we are able to locate extrema of the filtered image and get a sub-pixel estimate for the position. The Newton step in x, y is given by

$$(\delta x, \delta y) = -H^{-1} (G_x, G_y)$$

where (G_x, G_y) is the vector of first derivatives (the gradient) and H is the matrix of second derivatives (the Hessian).

In the output of this filter, local maxima correspond to the edges of locally convex bright objects, while local minima correspond to the edges of locally convex dark objects. The Newton step in θ and κ can also be calculated and this gives an estimate for the optimum values for θ and κ .

For θ we used the step in (x, y, θ) space

$$(\delta x, \delta y, \delta \theta) = - \begin{bmatrix} G_{xx} & G_{xy} & G_{x\theta} \\ G_{xy} & G_{yy} & G_{y\theta} \\ G_{x\theta} & G_{y\theta} & G_{\theta\theta} \end{bmatrix}^{-1} (G_x, G_y, G_\theta)$$

and similarly for κ . It would be possible to compute a four dimensional Newton step in (x, y, θ, κ) space, but previous experience with this sort of filter has suggested there may be numerical problems doing that.

Obviously this estimate is short-range in that if the filter parameters are such that the filter is poorly matched to the visual object sought then the Newton estimate is unreliable.

Tests on phantom images of curved objects have been carried out and these experiments show the expected results and are not reported here. The key issue is the performance of the operators described when applied

to noisy image data. This issue has been addressed by the experiments reported in section 3.

2.2 The Database

The database infrastructure for this project is a simple prototype triple store, along the lines of the Essex Sierra IFS simulator [7], with each triple representing a relationship between two entities. Lexical items are either encoded directly into the identifiers (eg integers, floats) or stored separately (strings). Unlike a relational database where the access programs need to know the names of tables and fields in the database in order to access the data, in the triple store the lowest level of storage is more uniform (a triple of identifiers) and doesn't change when new relationships are introduced. The relational model prejudices (to some extent) the queries to be made (and in compensation offers efficient access), whereas the triple store allows the database schema to be represented in the database itself, so that it can easily evolve.

For the fossil database data items were computed and stored as follows: for a given orientation and curvature the filter output was computed, along with first and second derivatives with respect to position, angle and filter curvature. From these extrema of the filter output were computed together with estimations for the true values of the curvature and orientation of the curved edge.

Experiments can now be expressed as computations on the result of queries to the database: eg "what is the mean position of hits which have their centres within a given distance of a given point?". At present these queries are embedded in programs, but the plan is to implement them in a special-purpose language.

3 Experimental Results

An example of the 5 strongest hits from a single oriented filter is shown in figure 4. The curves are marked over a range $\pm 2\sigma_\beta$ and follow the centre line of the filter. In figure 5 we show the two strongest hits from each of 20 directions for the image in figure 2.

3.1 Choice of parameters

The best choice of σ_α and σ_β for a given domain depends on the noise in the image [8] which in this application precludes the use of very small values. The upper limit on σ_α is dictated by the proximity of clutter. The filter arc length parameter σ_β is more model-dependent: for example in most cases we are interested in longer curves of low curvature and shorter curves of high curvature, so σ_β should be chosen longer or shorter to match. Apart from this the expected amount of deviation from circularity will also influence σ_β . Table (1) shows the variation in the number of hits on the spiral fossil of image (fig 1) as σ_β is varied - in this case the best value is about 30 pixels. The next two tables (2,3) show the results for two rather more circular spiral fossil images extracted from published photographs. σ_α was kept constant at 4.4 pixels, which is about half the thickness of the fossil walls, and so avoids interference between hits arising from the inside and from the outside walls of the fossil.

Table 1 - Fossil 1

σ_β	hits on fossil	% total hits
11	16	8
22	20	10
29	22	11
45	20	10

Table 2 - Fossil 2

σ_β	hits on fossil	% total hits
11	13	11
22	18	13
29	17	13.5
45	21	15

Table 3 - Fossil 3

σ_β	hits on fossil	% total hits
11	10	5
22	14	7
29	15	7.5
45	15	7.5

3.2 Curvature correction

To get a feel for the extent to which we can predict the curvature of an edge from the Newton step in curvature we give three examples of points where filters of different radius have hits. (These points were chosen at

random from the spiral image fig 1). The best matching filters (the consensus of the estimated radii) are marked * in tables 4 to 6 below.

Table 4 - point 1

filter radius	estimated radius of curve
40	55.23
50	47.62*
60	47.09*
70	44.28
80	3.39
90	89.35

Table 5 - point 2

filter radius	estimated radius of curve
40	30.71
50	-131.12
60	1050.50
70	109.33
80	85.19*
90	84.97*

Table 6 - point 3

filter radius	estimated radius of curve
40	24.42
50	71.1
60	60.3
70	58.7*
80	51.2*
90	26.8

3.3 Curvature Selectivity

To show curvature selectivity, six filters with a radius ranging from 40 to 90 pixels were applied to the image in fig 2. The twenty strongest hits from each of twenty directions were taken. If the spiral were perfectly smooth then as we travel out from the centre along the curve of the spiral we would expect to encounter hits from filters of a given radius only in one section of the spiral.

In practice with the fossil shown (fig 1) the curvature is quite irregular, but a relatively uncluttered part of the spiral occurs at a distance of 50 to 70 pixels from the fossil centre and plotting the density of hits in this area against filter radius reveals a peak at a filter radius of 60 to 70 pixels, shown as the upper line in figure 6.

This shows that the geometrical form of the filter introduces a selectivity or bias towards specific curvatures in the image.

The curvature selectivity of the filter may be increased by insisting that the detected curves should "match" the filter in the sense that the Newton step should be less than some given threshold. The lower trace of figure 6

shows the result of discarding hits whose Newton step in curvature is greater than 0.01 pixel^{-1} . In principle the lower the threshold the greater the selectivity although in practice the effect of very small thresholds is to eliminate almost all the hits.

3.4 Microfossil detection

To show the usefulness of the hit evidence from these microfossil images we used a simple binning method for the centres of curvature of the hits on spiral fossil images. In a perfect spiral such centres would lie on a tight curve near the centre of the fossil. Thus dividing the image up into square bins and counting the number of centres in each bin should yield peaks centred on spiral fossils. To test this a series of 14 images containing spiral fossils and 4 images with no fossils were taken and the strongest ten hits from each of twenty angles was computed. For various possible bin sizes the highest occupancy bin was found, and the density of hits (per square pixel) recorded. For very large bin sizes this will tend to a constant value, the average density, while for very small sizes it will be inversely proportional to the bin area (with the bin holding just one hit). However, in between we expect images with spiral fossils to have higher densities, and this effect can be seen in fig 7 where the average density is plotted against the bin width for the fossil bearing and non-fossil bearing images. We can use this effect to detect microfossil evidence simply on the basis of a fixed bin count threshold.

5 Conclusions

The methods employed yield reliable cues for locating spiral microfossil structures in plane rock samples. The curve finding algorithm described here provides a rich source of visual evidence from quite noisy images. The ability of the algorithm to output points rather than curves is a marked advantage in our database application since this simplifies the construction of visual models/queries. We hope that the set of points resulting from a well directed query will be useful for characterisation of microfossils as well as detection. The operator has several novel aspects, in particular the exclusive processing of local extrema and the use of partial derivatives for parametric estimation. Much work remains to be done on how this new type of feature evidence can be used effectively.

We have found that a symbolic representation for the curve evidence gives us the ability to experiment quickly by simply changing a database query without re-coding a program. Using this tool we have been able to explore the area of visual querying in one application.

6 Acknowledgements

This work is supported by the SERC. The fossil samples were provided by Fiona White of the Department of Geology, and we would like to thank her for sharing her expertise in foraminifera. The Sierra database simulator software was kindly made available by Prof. Lavington and his group at the University of Essex.

7 References

- [1] Swaby P. A. "A Graphical Expert system for Microfossil Identification", BCS Expert Systems Conference, London, 20-22 September 1989.
- [2] Parent P. and Zucker S.W. "Trace Inference, curvature consistency and curve detection", IEEE PAMI, Vol II, No 8 1989 pp 823-839.
- [3] Zucker S.W. David C. Dobbins A. and Iverson L. "The organisation of Curve Detection: Coarse tangent fields and fine spline coverings", Proceedings of 2nd ICCV, pp 568-577, 1988.
- [4] Oakley, Robinson & Shann, "An Efficient and Robust Local Boundary Operator", BMVC 91.
- [5] Shann & Oakley "Novel Approach to Boundary Finding" IVC 8 p32 (1990).
- [6] Murray Gill and Wright "Practical Optimisation".
- [7] Lavington & Wang "The External Procedural Interface for the IFS/2" Dept Computer Science, University of Essex report CSM 164, June 1991.
- [8] J.P. Oakley and R. T. Shann "An efficient method for finding the position of object boundaries to sub-pixel precision", Image and Vision Computing, Volume 9, no 4, August 1991.




Figure 1 FILTER KERNEL



Figure 2 SPIRAL FOSSIL



Figure 3 NOISIER FOSSIL




Figure 4 HITS FROM ONE DIRECTION




Figure 5 HITS FROM SEVERAL DIRECTIONS

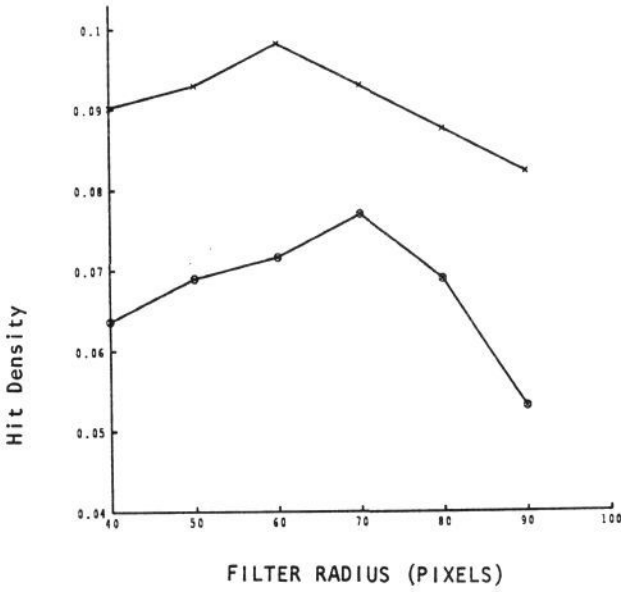


Figure 6 EFFECT OF FILTER CURVATURE

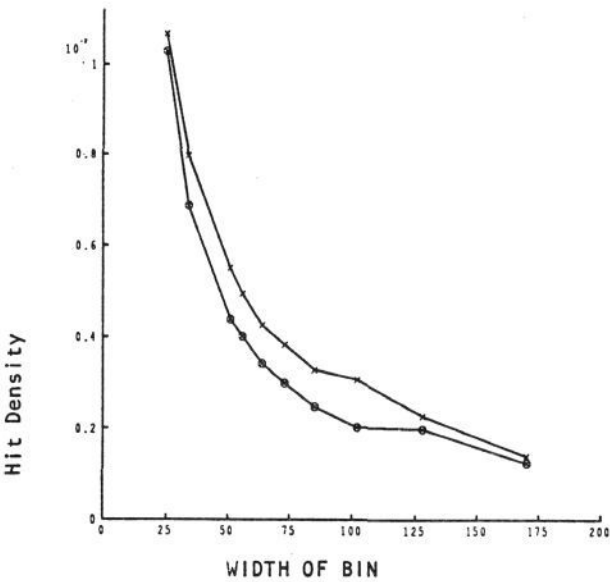


Figure 7 MICROFOSSIL DETECTION PERFORMANCE