ELSEVIER

Pattern Recognition Letters 19 (1998) 63–75

Pattern Recognition Letters

Determining local natural scales of curves

Paul L. Rosin [1]

Department of Computer Science and Information Systems, Brunel University, Uxbridge, Middlesex UB8 3PH, United Kingdom

Received 16 May 1997; revised 24 September 1997

Abstract

An alternative to representing curves at a single scale or a fixed number of multiple scales is to represent them only at their *natural* (i.e. most significant) scales. This allows all the important information concerning the different sized structures contained in the curve to be explicitly represented without the overhead of redundant representations of the curve. This paper describes several approaches to determining the local natural scales of curves. That is, various possibly overlapping sections of the curve should be represented at certain scales depending on their shape. The merits and drawbacks of the techniques are described, and the results of implementing one of them are shown. © 1998 Elsevier Science B.V.

Keywords: Scale selection; Multi-scale; Smoothing; Curve

1. Introduction

Curves can contain a variety of structures at different scales. Moreover, these structures are often superimposed, e.g. fine detail upon medium scale detail, which in turn is superimposed on coarse detail. Most processing techniques applied to curves (e.g. feature detection, model matching) work best at the appropriate scales for each of these structures. Otherwise, the spurious detail and noise that is present in most real image curves is likely to produce undesirable side-effects. For instance, when segmenting a curve into primitive parts such as codons (Hoffman, 1983), noise, quantisation effects, and irrelevant detail will cause the curve to be over-segmented into many insignificant tiny parts (Rosin, 1993).

The most common solution to the problem of noise and unwanted detail is to smooth each point of the curve at a single scale. For instance, the determination of a region of support (i.e. appropriate spatial scale) for calculating the curvature along a curve goes back to the work of Rosenfeld and Johnson (1973) in the early 70s, and has been more recently developed by Teh and Chin (1989). Other approaches include cross validation (Shahraray and Anderson, 1989) and significance measures (Lowe, 1989; Rosin, 1994).

Subsequently more than one scale has been considered desirable (Kropatsch, 1987; Mokhatarian and Mackworth, 1992). In the extreme case, scale is considered as a continuous dimension, giving rise to scale-space (Witkin, 1984). The multi-scale approach enables the behaviour of curve evolution over scale to be studied, and also makes possible the representation of multiple overlapping scales of interest.

[1] E-mail: paul.rosin@brunel.ac.uk.

PII S0167-8655(97)00158-X

Since the standard multi-scale approach exhaustively represents the curves at a fixed sampling rate (e.g. octave separated) over a wide range of scales it produces a cumbersome representation; its advantage is that it is guaranteed to include every relevant scale (assuming a fine enough sampling of scale-space). On the other hand, the single-scale approach provides a more compact representation, but may be unable to represent the complete curve at the correct scale, particularly when several differently sized structures are superimposed. The single and multi-scale representations are the extremes of a range of possible trade-offs between conciseness and robustness. An intermediate approach is to represent the curve only at certain selected scales called ''natural scales'' (Bengtsson and Eklundh, 1991; Cesar and da Fontoura Costa, 1996; García and Fdez-Valdivia, 1994; Rosin, 1992). These scales are intended to capture all the significant structures in the curve. [2] Several scales may be necessary, particularly if several structures at different scales are superimposed. Thus, all the relevant information is retained and made explicit without the cost of redundant representations. However, since the technique is unlikely to be perfect, it is liable to be less robust than the fixed multi-scale approach.

The natural and single scale based approaches can be divided into global and local methods. Much of the previous work on determining natural scales calculated global scales for representing the curve (Bengtsson and Eklundh, 1991; Cesar and da Fontoura Costa, 1996; Rosin and Venkatesh, 1993; Rosin, 1992), However, if the curves contain different sizes of structures at different locations then the global approach is not entirely suitable, although if there are still significant amounts of each sized structure then the global approach can still work (Rosin, 1992). This problem can be overcome by calculating the natural scales locally rather than globally. A simple approach would be to segment the curve into a sequence of smaller sections and independently determine their natural scales. Unfortunately this has the drawback that coarse scale statistics concerning large structures cannot be effectively calculated from small curve sections. Therefore the spatial extent of analysis has to correspond to the current scale of analysis. In this paper we describe several approaches to determining the local natural scales of curves, and show results from the implementation of one of these methods.

In Section 2 we describe various techniques (some from the literature, some new) for determining local natural scales. This is followed in Section 3 by outlining a method for smoothing the map of estimated scales in order to produce more coherent (i.e. less fragmented) results. Finally, several examples are presented in Section 4 of applying one of the scale determination techniques on both synthetic and real data.

2. Techniques for determining local natural scales

A convenient way to represent the curve is by scale-space (Witkin, 1984). The problem of detecting the local natural scales then becomes the task of finding significant bands in scale-space as illustrated in Fig. 1. For simplicity, we assume the curve has approximately piecewise constant natural scales. Since the shape of a curve is often formed by relatively uniform processes the natural scales at points along the curve should also be relatively uniform (Hoffman, 1983). Minor variations in local natural scale did not pose a severe problem to the global natural scale algorithms since they were averaged out by calculating the significance measure over the whole curve. In contrast, many of the following local techniques assign a set of natural scales to each pixel. To be useful for high-level

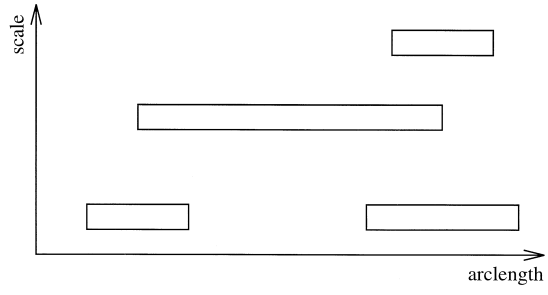

Fig. 1. Local natural scales in scale-space.

[2] In addition, there has also been recent interest in performing scale selection for intensity images (Jägersand, 1995; Kothe, 1996; Lindeberg, 1993).

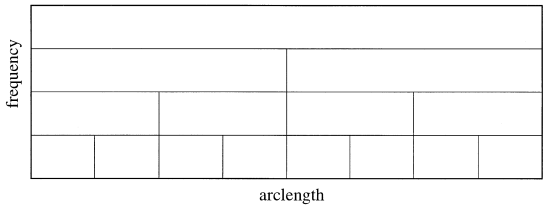

Fig. 2. Partitioning of scale-space.

processing adjacent pixels then need to be merged into curve segments. This can become problematic since both the deviation of scales and the increased length of the merged set of pixels should be considered. Since these are incommensurate quantities it is difficult to combine them in any principled manner. Thus, we require two parameters to specify the desired coherence in space and scale.

Below we outline various possible techniques (both old and new) and their limitations for determining the local natural scale of curves. It can be seen that most of them provide different means to produce a similar intermediate result – a scale-space significance map – which is then analysed to find natural scales. To prevent fragmentation or excessive duplication of the curve all the methods require some parameters effectively specifying a merging process as described above. Also, the methods permit the reconstruction of the curve at these natural scales although the methods vary, and include: subsampling followed by B-spline fitting; fitting parabolae; low pass filtering in the frequency domain; Gaussian smoothing; and regularisation.

2.1. Tangent variance

Hoffman described a method for detecting local natural scales that determined the degree of subsampling of the data to provide control points for a B-spline reconstruction of the curve at each scale (Hoffman, 1983). At each point p_t on the curve a pair of windows centred at $p_{t\pm f}$ and length w (where w is some function of the offset f) are considered. The line between each corresponding pair of points in the two windows (i.e. $p_{t+f+s} \rightarrow p_{t-f-s}$; $s=[-w/2,w/2]$) is taken as an estimate of the tangent at p_t. The variance of the tangents is calculated and this process is repeated over a range of offsets. Increasing the offset f increases the scale of the analysis. Natural scales are defined to be those offsets (scales) producing local minima of tangent variance.

One problem with this approach is that variance estimates will be unreliable at fine scales since the small windows will only provide a few tangent samples. Also, it is not reported how sensitive the algorithm is to the selection of the window length function which was chosen fairly arbitrarily. Moreover, Hoffman's single example only shows the detection of global natural scales on a synthetic curve. The problem of merging adjacent pixels with similar natural scales is not addressed.

2.2. Scale–space plot

Recently, Deguchi and Hontani (1994) described a method for calculating natural scales based on the scale-space plot of zero-crossings of curvature (Mokhatarian and Mackworth, 1992). The top of each arch specifies a natural scale. The complete curve is retained at every such scale, but it is smoothed adaptively to minimise distortion. At each point on the curve the selected scale lies on the largest scale closed arch whose peak is less than the natural scale. If there is no such arch then no smoothing is performed. This is shown in Fig. 3, where the dark line drawn in scale-space specifies the amount of smoothing applied to each point to generate the complete curve for the third natural scale. The two main weaknesses of this approach are first, since many parts of the final smoothed curves will be identical the representation has a great deal of redundancy. Second, there are likely to be a vast number of natural scales unless some thresholds are introduced. For instance, although not explicitly

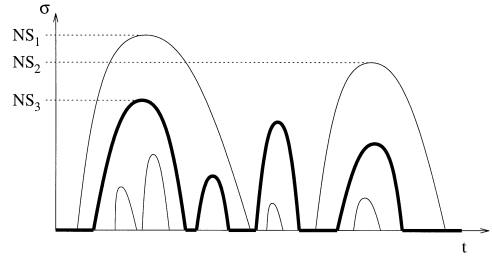

Fig. 3. Local scales selected by Deguchi and Hontani.

stated in the paper, a minimum level of smoothing is retained as a lower limit to eliminate noise.

2.3. Curvature-tuned smoothing

Dudek and Tsotsos (1990) employ a regularisation process (''curvature-tuned smoothing'') to smooth the curve where the stabilising functional defines a target curvature. This process is performed over a range of target curvatures to produce a multi-scale description. The energy functional being minimised indicates the appropriateness of the target curvature value. Discontinuities are inserted at points whose local energy is maximal and exceeds a threshold. The problem is non-linear, and so a suboptimal solution is found using a greedy algorithm. Curve sections whose average fitting energy is a local minimum over scale are retained. These sections are similar to the local natural scales of the curve. However, their approach is rather restrictive since it can only select curve sections with roughly constant curvature (i.e. approximately circular).

2.4. Complexity measures

Recently, Dubuc and Zucker (1995) describe two complexity measures (normal and tangential) which are based on the rate of growth of local orientated dilations of edges. They suggest that rapid changes in these complexity measures could be used to detect the emergence of natural scales, although few results are shown.

2.5. Interval tree

In his influential paper introducing scale-space Witkin described how zero-crossings could be linked over scale to form closed loops from which the interval tree is generated (Witkin, 1984). This is a ternary tree which partitions scale-space into rectangles. Each loop defines a rectangle whose upper scale bound is the maximum scale of the loop. The arc length positions of the two ends of the loop at the finest scale define the spatial bounds of the rectangle. The value of the maximum scale of the loop with the largest maximum scale that is contained within the rectangle defines the rectangle's lower scale bound. This is illustrated in Fig. 4; the thin

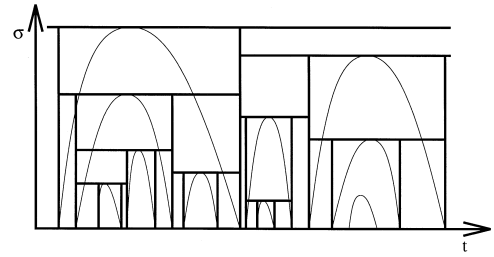

Fig. 4. Interval tree.

curves are the scale space plot of zero-crossings, and the interval tree is overlaid in bold. To generate a good single scale description of the curve Witkin used a stability measure based on a rectangle's persistence over scale which was simply equal to its height in the scale domain, although he also experimented with other (unspecified) more complex stability measures. Rectangles were selected from the interval tree by descending from the root until a rectangle's stability was greater than or equal to the mean stability of its children. We suggest that this approach could be extended to extract natural scales from the interval tree in a similar manner. All rectangles which locally maximise the stability measure over scale would be chosen as local natural scales.

In contrast to some of the other techniques the interval tree approach assigns natural scales to curve sections delimited by zero-crossings. For 2D curves the scale-space map is built up from the zero-crossings of the smoothed curvature values, so that the curve sections will be concave, convex, or straight sections. However, this still provides an inconvenient representation, particularly at fine scales which will contain vast numbers of sections. Thus, merging is still required, although at a curve section rather than pixel level.

2.6. Frequency analysis

The Fourier based technique that we have applied to determine global natural scales (Rosin and Venkatesh, 1993) can be extended to apply locally. Previously, any Fourier descriptors with large magnitudes were taken as indications of the presence of many structures at the corresponding scales. Whereas the Fourier transform is restricted to the global analysis of signals, local analysis requires that the curve

be decomposed instead by spatially localised basis functions such as Gabor filters or wavelets (Mallat, 1989). These partition the signal in both the spatial and frequency dimensions as shown in Fig. 2. Natural scales could be defined as clusters of basis functions (which are close both in scale and space) with locally maximal magnitudes.

A difficulty arises in representing the curve so that it can be decomposed by wavelets. In contrast to the Fourier transform which has both real and imaginary components the wavelet transform only has a real component. When using the Fourier transform the curve can be parameterised by its co-ordinates, allowing easy reconstruction of the curve at the natural scales. This is not possible with wavelets and instead a single valued parameterisation such as tangent angle or curvature must be used. This is less desirable since reconstructing the curve from its filtered (smoothed) tangents or curvatures can produce distortions. For instance, closed curves generally become open ones.

2.7. Zero-crossing density

Our original method for determining global natural scales smooths the curve $C=\{(x_t,y_t)\}$ with a Gaussian filter (corrected for curve shrinkage (Lowe, 1989)) at octave separated scales and calculates a significance measure S_σ at each scale. S_σ is defined as the sum of the number of zero-crossings of curvature at all points t on the curve normalised by the Gaussian smoothing scale σ:

$$S_\sigma=\sigma\sum_{t\in C} z_t,\quad \text{where } z_t=\begin{cases}1, & \kappa_t=0,\ \kappa_t'\neq 0,\\ 0, & \text{otherwise.}\end{cases} \tag{1}$$

Dropping subscripts for convenience, at each point curvature κ is calculated as

$$\kappa=\frac{X'Y''-Y'X''}{\left(X'^2+Y'^2\right)^{3/2}}, \tag{2}$$

where X', Y', X'' and Y'' are the results of independently convolving the X and Y coordinates of curve with the first and second derivatives of the Gaussian kernel $G_\sigma(t)$:

$$G_\sigma(t)=\frac{1}{\sigma\sqrt{2\pi}}\mathrm{e}^{-t^2/2\sigma^2}, \tag{3}$$

$$G_\sigma'(t)=\frac{-t}{\sigma^3\sqrt{2\pi}}\mathrm{e}^{-t^2/2\sigma^2}, \tag{4}$$

$$G_\sigma''(t)=\frac{-t}{\sigma^5\sqrt{2\pi}}\left(\frac{t^2}{\sigma^2}-1\right)\mathrm{e}^{-t^2/2\sigma^2}. \tag{5}$$

Natural scales are defined to be at scales producing local minima of S_σ.

If the significance value is considered as a normalised average zero-crossing density measure then it can be easily applied locally to sections of curve. We consider two approaches for calculating local zero-crossing density. Like the interval tree approach, the first treats each section of curve bounded by a zero-crossing of curvature as a primitive concave/convex/straight curve element. Within that section the density of zero-crossings of curvature is taken as the inverse of the length of the section. The significance value S_t^σ at each pixel in the curve section c of length l_c at scale σ is calculated as the normalised density of zero-crossings of curvature of the section in a similar manner to the global significance value:

$$S_t^\sigma=\frac{\sigma}{l_c},\quad t\in c. \tag{6}$$

We will call the above approach ''method A''. The second method, ''method B'', employs a Parzen estimator, a standard technique for estimating local densities (Fukunaga, 1992). The general form of the estimator at location t is

$$P(t)=\frac{1}{N}\sum_{i=1}^{N}\frac{1}{h}K\left(\frac{t-\boldsymbol{t}_i}{h}\right), \tag{7}$$

where N is the number of observations, h is a positive number and a function of N, $K(\cdot)$ is the kernel of the estimate (i.e. the point spread function), and $\boldsymbol{t}_i$ is the ith observation. We use a Gaussian kernel $G_\sigma(\cdot)$ with $h=\sigma$, so that the density of zero-crossings of curvature at scale σ is

$$P_t^\sigma=\frac{1}{\sum_i z_i\sigma}\sum_{i\in C} z_i G\left(\frac{t-i}{\sigma}\right) \tag{8}$$

and so the significance, calculated as the normalised zero-crossing density, is

$$S_t^\sigma=\sigma P_t^\sigma=\frac{1}{\sum_i z_i}\sum_{i\in C} z_i G\left(\frac{t-i}{\sigma}\right). \tag{9}$$

For both methods of zero-crossing density estimation the natural scales of pixels t are those σ's at which S_t^σ is a local minima over adjacent scales.

We note that an alternative approach involving zero-crossing density was taken by García and Fdez-Valdivia (1994) who first segmented the curve into sections, and then determined the most significant global natural scale of each section. However, their method is flawed since the initial segmentation process involved combining various arbitrary feature measures that were calculated at two arbitrarily chosen scales. To be correct, the segmentation process should also be performed at multiple scales.

3. Merging natural scales

As stated in Section 2 pixel-based techniques for determining local natural scale are prone to noise and minor variations in the data. The resulting fragmentation is undesirable since most applications will assume the data contains a certain amount of spatial coherence. We employ a straightforward merging method that is carried out in the following stages:

(1) Minima over scale are detected in the significance map which is then binarised into minima (set to 1) and non-minima (set to 0).

(2) Information is integrated across scales by applying 1D Gaussian smoothing at each point at each scale. For the ideal case the scales adjacent to a significance minimum will be non-minima. Therefore, to prevent the smoothing operation degrading the likelihood of the existence of a minimum an unnormalised Gaussian filter is used: $G_\sigma^u(t) = e^{-t^2/2\sigma^2}$. The parameter $\sigma = \Sigma$ specifies the required degree of coherence over scale.

(3) Smoothing along the curve is performed at each scale independently by applying a standard 1D Gaussian filter at each point, where the smoothing parameter $\sigma = T$ determines spatial coherence.

(4) Non-maximal suppression is performed at each point across adjacent scales; non-maxima are reset to zero. Local scales are then retained at points which are more likely than not to be a local minimum over scale by thresholding the smoothed minima map at 0.5.

(5) Finally, to ensure there is no fragmentation, horizontal gaps smaller than T in the binary smoothed minima map are removed by relabelling them with their opposite label (i.e. minima → non-minima and non-minima → minima).

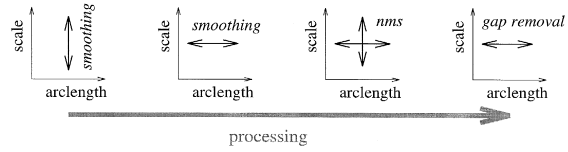

Fig. 5. Sequence of processes for merging local natural scales.

The sequence of processes is illustrated by the diagram in Fig. 5. Of course, the two 1D Gaussian filters in the spatial and scale domains could be combined into a single 2D Gaussian filter. However, the separable version described is more efficient. It

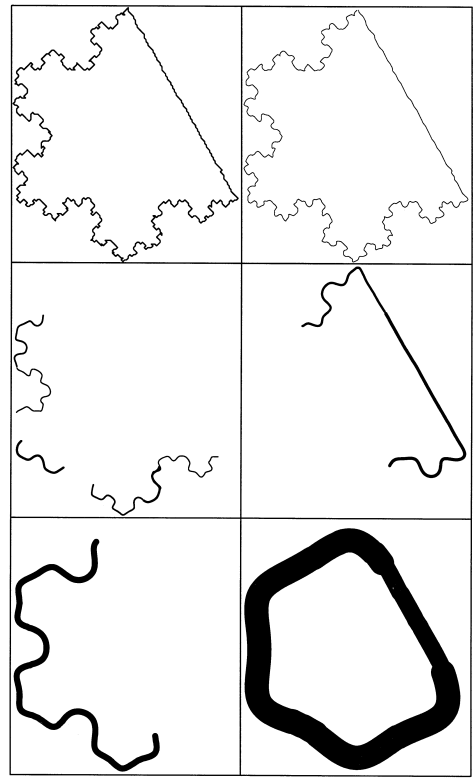

Fig. 6. (a) original stochastic Koch curve; (b)–(f) local natural scales; (g) zero-crossings of curvature; (h) significance measure – log mapped for display purposes; (i) minima of significance measure; (j) minima map after smoothing and non-maximal suppression; (k) minima map after gap removal; (l) log mapped significance measure using Parzen estimator.

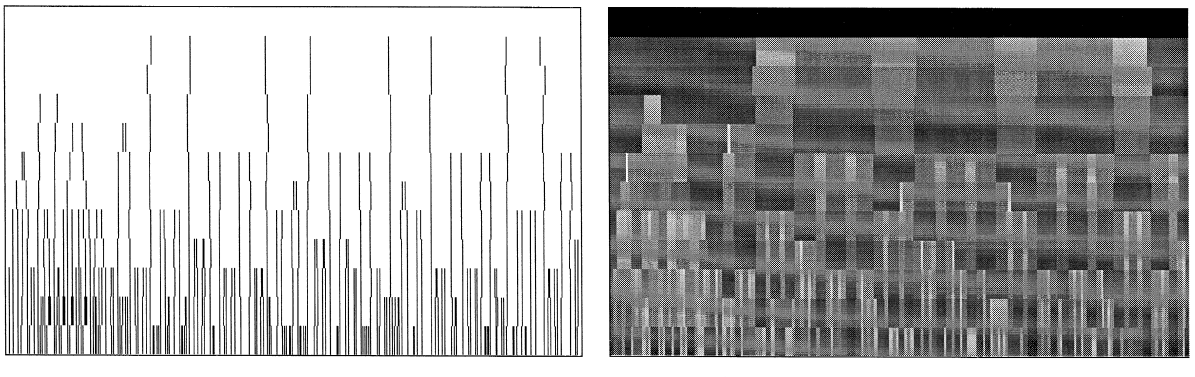

(g) (h)

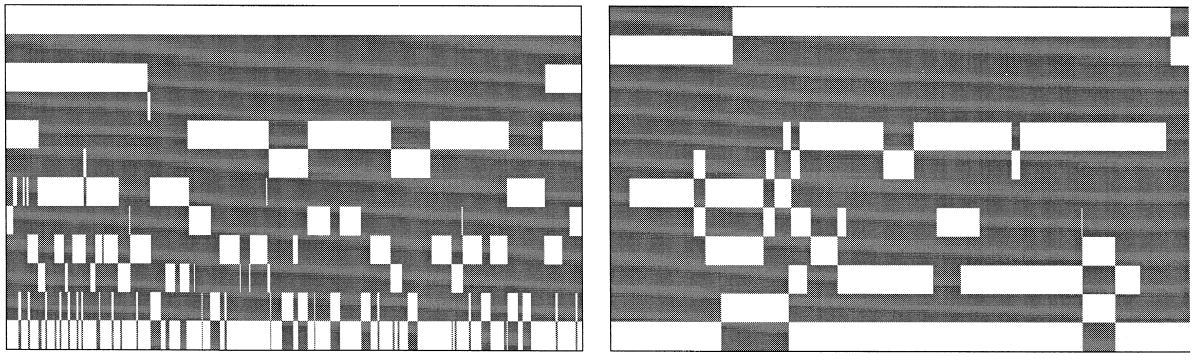

(i) (j)

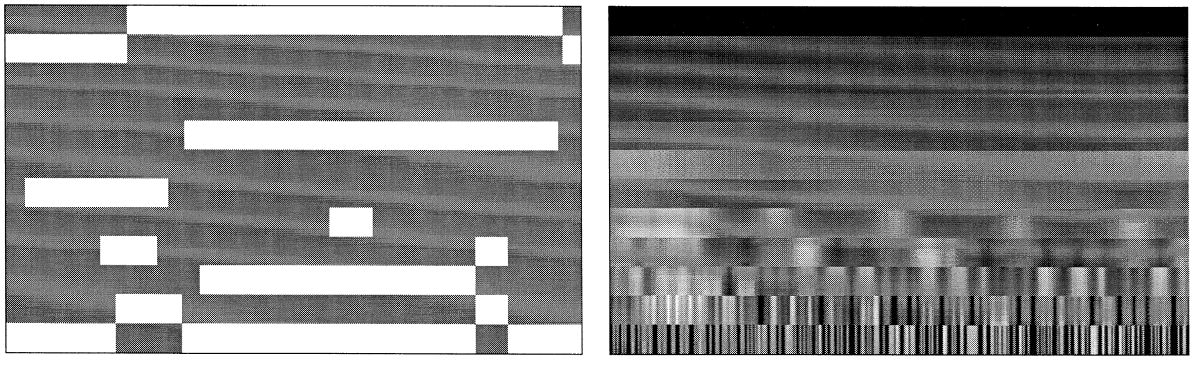

(k) (l)

Fig. 6 (continued).

should be noted that although there is a relationship between the smoothing parameters Σ and T the exact formulation depends on the nature of the data. Of particular relevance are: the magnitude and type of the noise; the length of sections of (ideally) identical local natural scale; and the separation in scale of overlapping (in the spatial dimension) sections of local natural scale. For instance, if the ideal local natural scales are well separated in scale, then substantial smoothing over scale can be performed to merge similar scales without fear of incorrectly over-merging distinct natural scales. A similar argument applies to spatial smoothing. However, in practise it may be difficult to determine the values of Σ and T

except through experimentation. In our examples we generally used $T = 40$ and $\Sigma = T/60$.

4. Examples

In this section we show some examples of applying the method of Section 2.7 (the zero-crossing density technique) to both synthetic and real data. We previously showed the results of successfully determining the global natural scales of the Koch fractal curve (Rosin, 1992). Since the curve has by definition well defined scales it was easy to verify the correct natural scales. Here we modify the curve in two ways to produce a more challenging test. First a stochastic version is made by randomly shifting each interior vertex during its generation. Perturbation is done by adding zero mean Gaussian noise with variance linearly decreasing at each level of recursion. Second, one side of the curve is replaced by a straight line with added noise, so that the natural scales are no longer global. Fig. 6(a) shows the stochastic Koch curve, and Fig. 6(b)–(f) show the curve sections at their detected natural scales drawn with thickness proportional to the scale. It can be seen that each level of structure has been reasonably isolated although there are gaps in the second and third natural scales (Fig. 6(c),(d)). The intermediate processes can be seen by examining the generation of the minima map. Fig. 6(g) shows the zero-crossings of curvature from which the significance map (Fig. 6(h)) is formed using the first method for

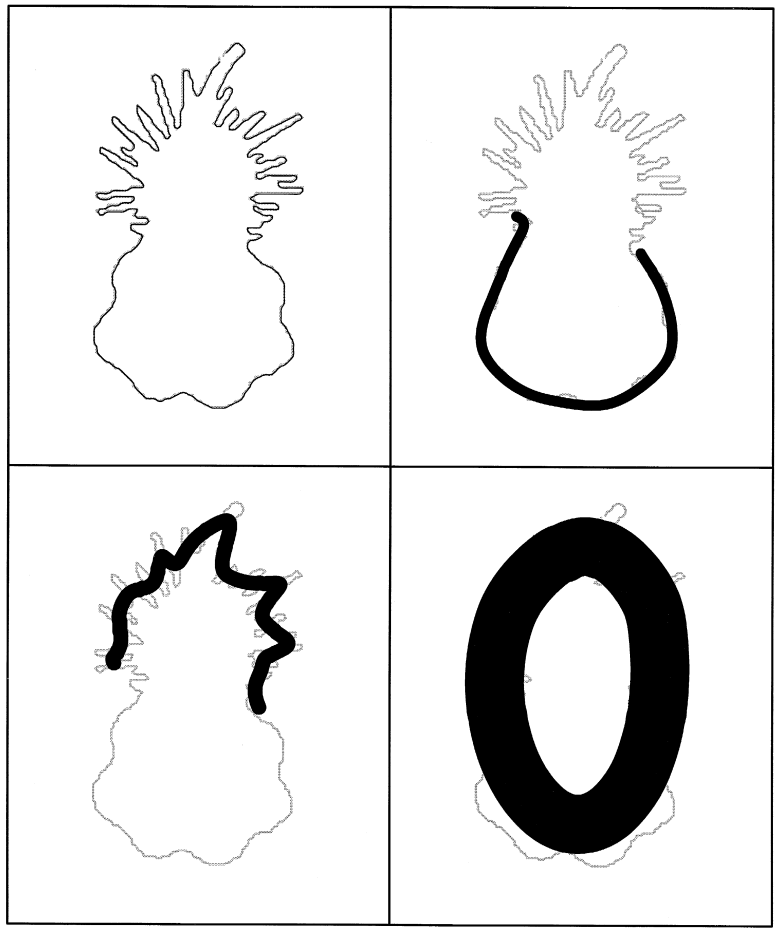

Fig. 7. Local natural scales of the half-spiny pear.

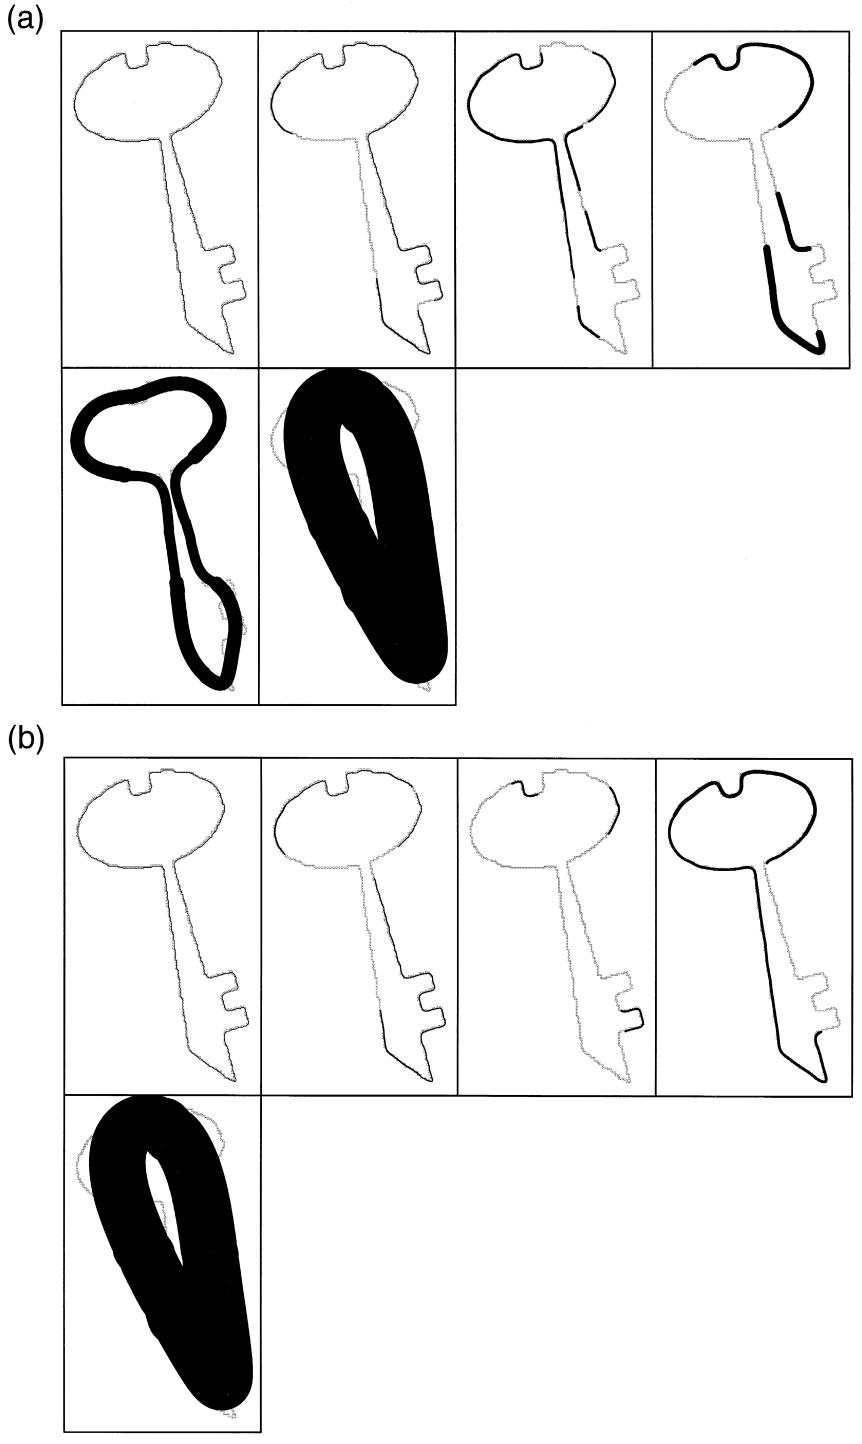

Fig. 8. (a) Local natural scales of Deguchi's key. (b) Method B.

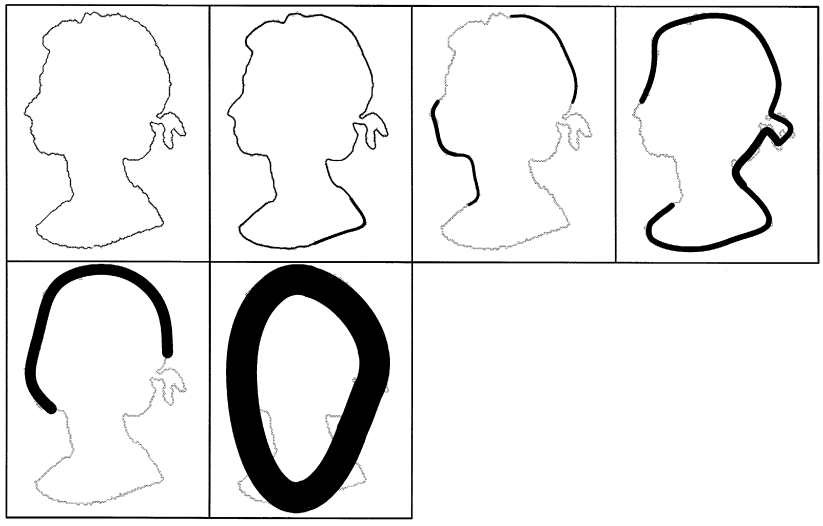

Fig. 9. Local natural scales of the Queen's head.

calculating zero-crossing density. Dark points correspond to low values of the significance measure. The repetitive structure of the Koch curve is evident in the right two thirds of the significance map, while the remaining third on the left corresponds to the straight segment. The minima displayed in white (Fig. 6(i)) are fragmented. After smoothing ($T = 32$ and $\Sigma = T/35$) and non-maximal suppression fragmentation is reduced (Fig. 6(j)), and gap removal further reduces fragmentation (Fig. 6(k)). The result of using the Parzen estimator to generate the significance map is shown in Fig. 6(l). Some natural scales have been missed at the coarse level, suggesting the zero-crossing density estimator is oversmoothing, otherwise the remainder are similar to the previous results.

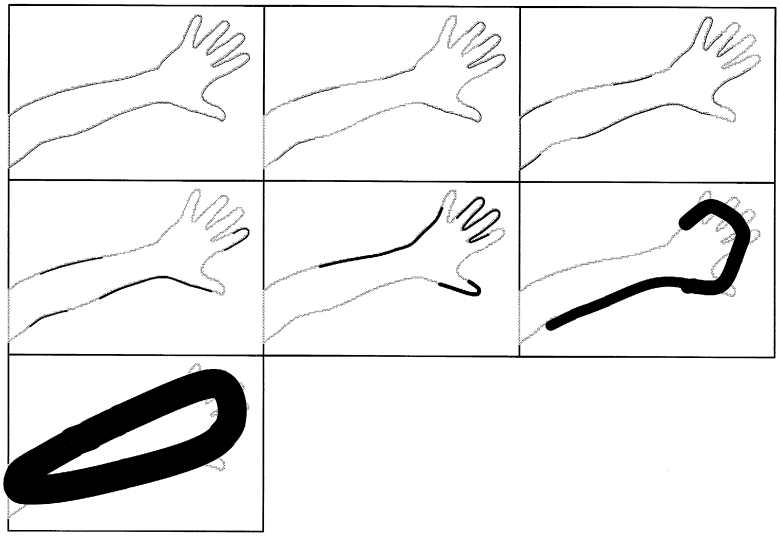

Fig. 10. Local natural scales of an arm.

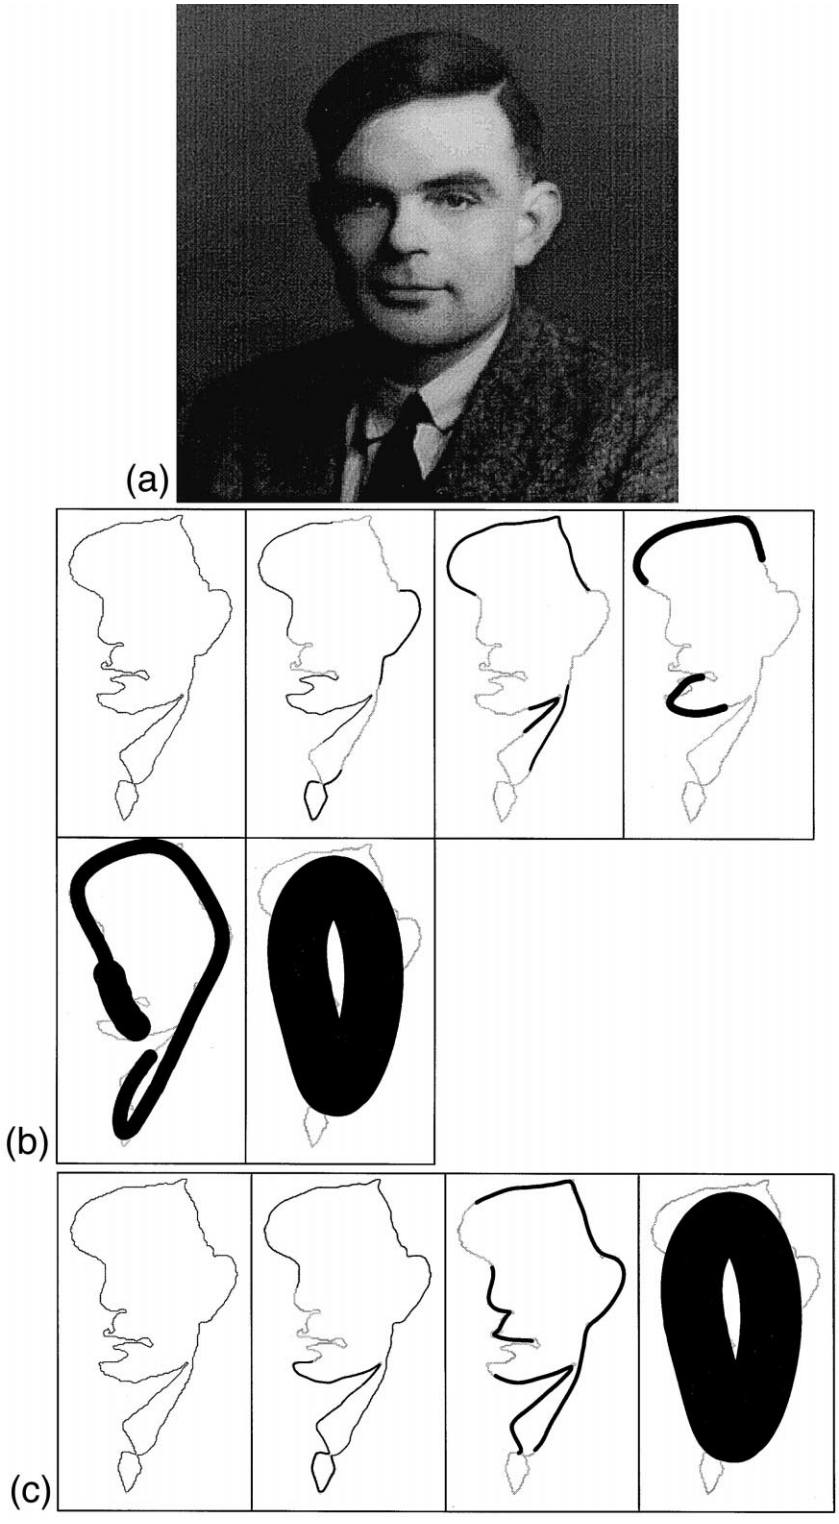

Fig. 11. (a) Image of Alan Turing. (b) Method A. (c) Method B.

The following examples all use the smoothing parameters set to $T = 40$ and $\Sigma = T/60$. Unless other otherwise stated method A is used rather than method B.

A second synthetic example that clearly demonstrates varied local scales is the spiny/rippled pear from Richards et al. (1986). The first natural scale (Fig. 7(a)) describes the complete curve at the original scale. Thereafter, the scales describe the bottom and top halves at appropriate intermediate scales (Fig. 7(b),(c)) and the complete figure again as a single blob (Fig. 7(d)).

The key shown in Fig. 8 has been redigitised from Deguchi's paper. Method A has correctly extracted the main features (the top of the key, the prongs, the overall shape) although there are some additional redundant sections. The results using method B are less successful.

The head and shoulders of the Queen, digitised from a British stamp, is shown in Fig. 9. Extracted features include the ribbon, the head, as well as the overall shape.

The next example, this time an arm, is shown in Fig. 10. Although there is some redundancy the hand has been successfully located as a single feature.

Finally, the last example is more demanding. since the scales are not as obviously demarcated as in the previous examples.

The image of Alan Turing (Fig. 11(a)) was thresholded, the curve bounding part of his face and clothes extracted, and the results using methods A and B are shown in Fig. 11(b) and Fig. 11(c). Both methods produce similar descriptions at the fine and coarse levels: the complete curve is represented at both extremes of scale, thereby capturing the finest detail and the global shape. The methods differ at the intermediate scales, but generally extract significant features such as the chin, ear, hair line, shirt collar, etc., at the appropriate scales.

5. Conclusions

We have described several techniques for determining the local natural scales of curves. Most are based on creating a significance value over scale which is then used to partition scale-space, thereby generating local natural scales. To produce practical results the significance values must be filtered to prevent excessive fragmentation of the curve into small adjacent sections with similar but non-identical natural scales.

In particular we showed results of applying the zero-crossing of curvature density method followed by smoothing of the minima map (which was derived from the significance map) to merge similar curve sections. Relatively concise multi-scale representations were obtained, even for fairly complex curves. Method A for smoothing the significance values tended to produce better results than method B in many of our examples (including those not shown here). Both methods produced some redundancy, i.e. additional spurious local natural scales. This is probably due to the use of a single set of fixed smoothing parameters. Although individual tuning of the parameters values would have improved the results a better approach would be to develop a method to automatically determine these values.

References

Bengtsson, A., Eklundh, J.O., 1991. Shape representation by multiscale contour approximation. IEEE Trans. Pattern Anal. Machine Intell. 13 (1), 85–93.

Cesar, R.M., da Fontoura Costa, L., 1996. Shape characterization in natural scales using the multiscale bending energy. In: Proc. 13th ICPR, pp. 735–739.

Deguchi, K., Hontani, H., 1994. Multiscale contour approximation based on scale space analysis with a stable Gaussian smoothing. In: Proc. 2nd Internat. Workshop on Visual Form, pp. 139–148.

Dubuc, B., Zucker, S.W., 1995. Indexing visual representations through the complexity map. In: Proc. ICCV, pp. 142–148.

Dudek, G., Tsotsos, J.K., 1990. Recognizing planar curves using curvature-tuned smoothing. In: Proc. 10th ICPR, pp. 130–135.

Fukunaga, K., 1992. Introduction to Statistical Pattern Recognition. Academic Press, New York.

García, J.A., Fdez-Valdivia, J., 1994. Representing planar curves by using a scale vector. Pattern Recognition Letters 15 (9), 937–942.

Hoffman, D., 1983. Representing shapes for visual recognition. Ph.D. Thesis, MIT.

Jägersand, M., 1995. Saliency maps and attentional selection in scale and spatial coordinates: An information theoretic approach. In: Proc. ICCV, pp. 195–202.

Kothe, U., 1996. Local appropriate scale in morphological scale-space. In: Proc. ECCV96, Vol. I, pp. 219–228.

Kropatsch, W.G., 1987. Curve representations in multiple resolutions. Pattern Recognition Letters 6, 179–184.

Lindeberg, T., 1993. On scale selection for differential operators. In: Proc. SCIA, pp. 857–866.

Lowe, D.G., 1989. Organization of smooth image curves at multiple scales. Internat. J. Comput. Vision 3, 119–130.

Mallat, S.G., 1989. A theory for multiresolution signal decomposition: The wavelet representation. IEEE Trans. Pattern Anal. Machine Intell. 11, 674–693.

Mokhatarian, F., Mackworth, A., 1992. A theory of multiscale curvature-based shape representation for planar curves. IEEE Trans. Pattern Anal. Machine Intell. 14, 789–805.

Richards, W., Dawson, B., Whittington, D., 1986. Encoding contour shape by curvature extrema. J. Opt. Soc. Am. Ser. A 3, 1483–1491.

Rosenfeld, A., Johnson, E., 1973. Angle detection on digital curves. IEEE Trans. Comput. 2 (2), 875–878.

Rosin, P.L., 1992. Representing curves at their natural scales. Pattern Recognition 25, 1315–1325.

Rosin, P.L., 1993. Multi-scale representation and matching of curves using codons. CVGIP: Graphical Models and Image Process. 55, 286–310.

Rosin, P.L., 1994. Non-parametric multi-scale curve smoothing. Internat J. Pattern Recognition and AI 8, 1381–1406.

Rosin, P., Venkatesh, S., 1993. Extracting natural scales using Fourier descriptors. Pattern Recognition 26, 1383–1393.

Shahraray, B., Anderson, D.J., 1989. Optimal estimation of contour properties by cross-validated regularization. IEEE Trans. Pattern Anal. Machine Intell. 11, 600–610.

Teh, C.-H., Chin, R.T., 1989. On the detection of dominant points in digital curves. IEEE Trans. Pattern Anal. Machine Intell. 11 (8), 859–872.

Witkin, A.P., 1984. Scale-space filtering: A new approach to multi-scale description. In: Image Understanding 1984, Ablex, pp. 79–95.