Determining Local Natural Scales of Curves

Paul L. Rosin
Institute for Remote Sensing Applications
Joint Research Centre
21020 Ispra (VA), Italy
email: paul.rosin@jrc.it

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.*

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 [3], noise, quantisation effects, and irrelevant detail will cause the curve to be over-segmented into many insignificant tiny parts [11]. This implies that the curve should be analysed at multiple scales. A number of scale-based approaches exist, and can be ordered by two categories based on 1) the size of the retained set of scales; and 2) the spatial extent of analysis.

The standard multi-scale approach exhaustively represents the curves at a fixed sampling rate (e.g. octave separated) over a wide range of scales [6, 7]. At the other end of the spectrum the curve can be represented at a single scale [13]. An intermediate approach is to represent the curve only at certain selected scales called "natural scales" [9, 10]. These scales are intended to capture all the significant structures in the curve. Several scales may be necessary, particularly if several structures at different scales are superimposed. Each of these approaches provides a different trade-off between conciseness and robustness. The fixed multi-scale approach is the most cumbersome representation, but also the most robust – it is guaranteed to include every relevant scale (assuming a fine enough sampling of scale-space). The single scale approach provides the most compact representation. However, it may be unable to represent the complete curve at the correct scale, particularly when several differently sized structures are superimposed. Natural scales provide a compromise between the above two approaches. 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 also be divided into global and local methods. Our previous work on determining natural scales calculated global scales for representing the curve [9, 10]. 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 [10]. To overcome such problems the natural scales should be calculated 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. Previously we have described a method for locally determining single scale curve descriptions [12]. 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.

2 Techniques for determining local natural scales

A convenient way to represent the curve is by scale-space [14]. The problem of detecting the local natural scales then becomes the task of finding significant bands in scale-space as illustrated in figure 1. For

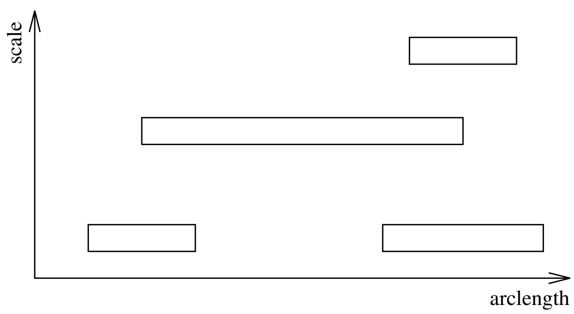

Figure 1: Local natural scales in scale-space

simplicity, we are assuming here a constant natural scale over a section of curve. Although examples containing varying scales are possible (e.g. figure 2) they can usually be approximated by piecewise constant scales. In addition, 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 [3]. Whereas our algorithm for determining global natural scales was non-parametric [10] it is difficult to design a non-parametric algorithm for determining local natural scales. Even if a real world (i.e. non-synthetic) curve appears to contain regular structures of similar size they are unlikely to be precisely identical. Therefore their local natural scales will vary slightly. This did not pose a severe problem to the global natural scale algorithm since such irregularities were averaged out by calculating the significance measure over the whole curve. In contrast, all the following local techniques assign a set of natural scales to each pixel. However, individual pixels are not suitable for high-level processing such as feature detection or model matching. Therefore, adjacent pixels with similar natural scales must be merged to form contiguous sections of curve. If only pixels with identical natural scales were merged the resulting representation would still be too fragmented to be useful. The fragmentation can be overcome by allowing some amount of deviation between adjacent natural scales during merging. Ideally a cost function should incorporate both the deviation of scales and the increased length of the merged set of pixels. But since the dimensions of scale space are incommensurate it is difficult to decide in any principled manner how they should be combined.

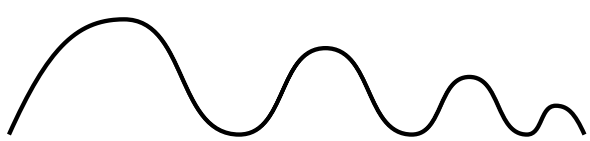

Figure 2: Curve containing continually varying scales

Below we outline various possible techniques 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 map – which is then analysed to find natural scales. Also, all permit the reconstruction of the curve at these scales although the methods include subsampling and B-spline fitting, parabola fitting, low pass filtering in the frequency domain, Gaussian smoothing, and regularization.

1/ 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 (although his single example only shows global natural scales) [3]. 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. $p_{t+f+s} \rightarrow p_{t-f-s}$; $s = -\frac{w}{2} \ldots \frac{w}{2}$) 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 increases the scale of the analysis. Natural scales are defined to be those offsets (scales) producing local minima of tangent variance. One potential 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.

2/ Witkin described how zero-crossings could be linked over scale to form closed loops from which the interval tree is generated [14]. 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. Witkin defined a stability measure based on a rectangle's persistence over scale which was simply equal to its height in the scale domain. Rectangles were selected from the interval tree by descending from the top until a rectangle's stability was greater than or equal to the mean stability of its children. These provided a good single scale description for each section of the curve which was then reconstructed by independently fitting parabolas to the original data between the zero-crossings localised by the selected rectangles. This approach could be extended to extract natural scales from the interval tree in a similar manner by choosing all local maximally stable rectangles rather than just the top-most ones.

3/ The Fourier based technique that we have applied to determine global natural scales [9] can be extended to apply locally. Previously, any Fourier descriptors with large magnitudes were taken as indications of the presence of many structures at similar 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 [5]. These partition the signal in both the spatial and frequency dimensions as shown in figure 3. Natural scales could be defined as clusters of basis

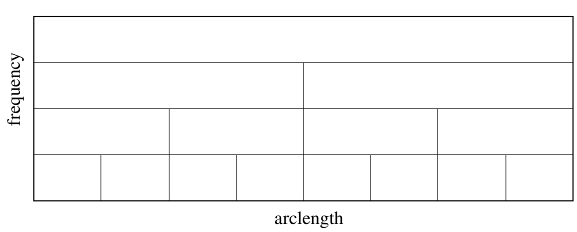

Figure 3: Partitioning of scale-space

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.

4/ Our original method for determining natural scales calculated a significance measure S_σ 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:

$$S_\sigma = \sigma \sum_t z_t \text{ where } z_t = \begin{cases} 1 & \kappa_t = 0, \kappa'_t \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

In fact the zero-crossing function z_t also has to account for straight sections of curve where the curvature is exactly zero, although this rarely occurs with real data. 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. Each section of curve bounded by a zero-crossing of curvature is treated as a primitive concave/convex curve element. Within that section the density of zero-crossings of curvature is taken as the inverse of the length of the concave/convex section. The significance value S_t^σ at each pixel in the concave/convex 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:

$$S_t^\sigma = \frac{\sigma}{l_c}; t \in c$$

The natural scales of pixels are those σ's at which S_t^σ is a local minima over adjacent scales.

5/ Recently, Deguchi and Hontani [1] described a method for calculating natural scales based on the scale-space plot of zero-crossings of curvature [7]. 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 figure 4 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 that the final smoothed

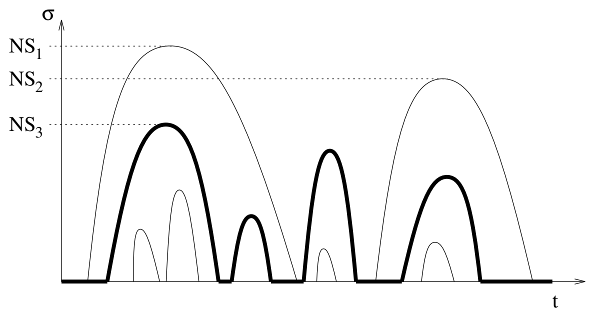

Figure 4: Local scales selected by Deguchi

curves contain much redundancy since many parts of them will be identical. Second, there are likely to be a vast number of natural scales unless some thresholds are introduced. For instance, although not explicitly mentioned in the paper, a minimum level of smoothing is retained as a lower limit to eliminate noise.

6/ Related work is described by Dudek & Tsotsos [2]. A regularization process ("curvature-tuned smoothing") is employed 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. Sections of curve whose energy exceeds a threshold are segmented and those section with locally minimal average energy over scale are retained. These sections are similar to the local natural scales of the curve. However, their approach is more restrictive since it can only select curve sections with roughly constant curvature.

3 Example of the concave/convex length method

In this section we show an example of method 4 (based on the zero-crossing counting global natural scale technique [10]) and describe its operation in further detail. It is applied to the spiny/rippled pear from Richards *et al.* [8]. The curve is smoothed by a Gaussian filter at octave separations with Lowe's method for correcting curve shrinkage applied [4]. The scale-space plot of the zero crossings of curvature are shown in figure 5a (the scale dimension has been stretched for display purposes). Note that it is not necessary to link the zero-crossings over scale. This is especially useful when there are artefacts like the dense triangle of zero-crossings. These occur on diagonal straight lines and are a product of the quantisation of the original pixel

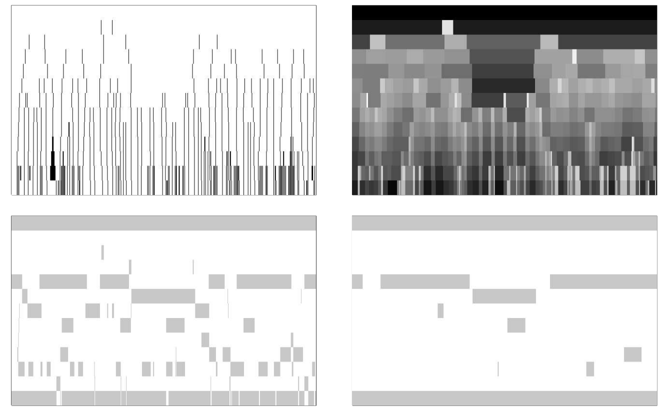

Figure 5: (a) zero-crossings of spiny/rippled pear; (b) significance measure; (c) minima of significance measure; (d) majority filtered significance minima

co-ordinate values. The significance measure at each pixel over scale is shown in figure 5b; the image has been log mapped for display purposes. Dark points correspond to low values of the significance measure. The dark patch extending from the coarsest scale to a medium scale in the middle of the curve corresponds to the bottom of the pear which has relatively little structure at coarse scales compared to the spiny top half.

Local minima of significance over scale are plotted in grey in figure 5c. Since the measure is calculated locally at every pixel it is not surprising that although there is a certain spatial coherence it is also prone to noise and minor variations. As discussed earlier this fragmentation is undesirable to most applications, but overcoming it unfortunately requires some threshold. Currently we clean up the set of minima by independently applying spatial majority filtering at each scale. The neighbourhood of the filtering operation can be varied to alter the amount of cleaning up. Figure 5d shows the result of applying the majority filtering with a window size of 32. It can be seen that small spatial gaps in adjacent natural scales have been closed and most small isolated stretches of natural scales have been eliminated. The remaining sections of natural scale shorter than the filtering window size (i.e. 32) are further deleted. The remaining sections of curve at their natural scales are plotted in black over the original curve in figure 6. Both the finest and coarsest scales for the complete curve have been retained as describing important aspects of the curve, namely the fine spiny and ripply detail. The shorter sections of curve shown in the middle row describe the underlying smooth base of the pear and the underlying

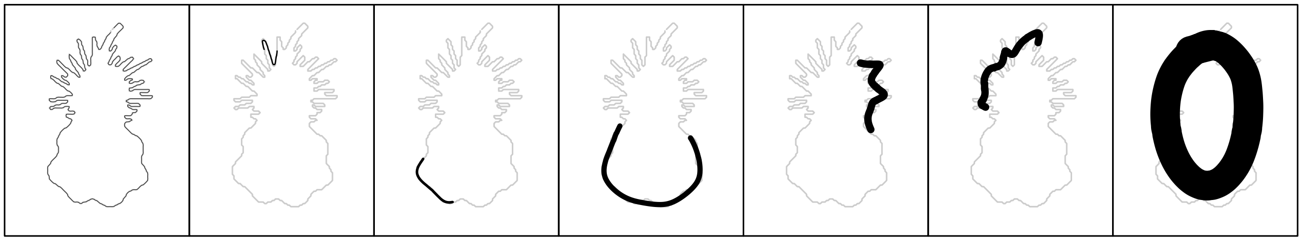

Figure 6: Selected local natural scales (black) overlaid on original curve (grey)

ripples of the top half of the pear. The remaining two curve sections do not appear to describe such obvious features.

4 Conclusions

We have described several techniques for determining the local natural scales of curves. Most are based on creating a significance value over scale space which is then used to partition scale-space, thereby generating local natural scales. An unfortunate by-product of extending the global analysis of natural scales to a more local basis is that in real curves the natural scales calculated at a pixel level tend to be fragmented. This can be overcome by merging fragments, but requires the introduction of a parameter which had been previously avoided in the global analysis. When high level knowledge is available it may be possible to automatically determine a suitable merging threshold. Currently we merge at each scale independently. However, this ignores information from adjacent scales that could help in the merging decision, and we are currently investigating alternative methods for merging natural scales. We are also currently implementing the other methods for determining natural scales so that their performance can be compared.

References

[1] K. Deguchi and H. Hontani. Multiscale contour approximation based on scale space analysis with a stable Gaussian smoothing. In *Proc. 2nd Int. Workshop on Visual Form*, pages 139–148, 1994.

[2] G. Dudek and J. K. Tsotsos. Recognizing planar curves using curvature-tuned smoothing. In *Proc. 10th ICPR*, pages 130–135, 1990.

[3] D. Hoffman. *Representing shapes for visual recognition.* PhD thesis, MIT, 1983.

[4] D. G. Lowe. Organization of smooth image curves at multiple scales. *Int. J. Computer Vision*, 3:119–130, 1989.

[5] S. G. Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans. PAMI*, 11:674–693, 1989.

[6] D. H. Marimont. A representation for image curves. In *Proc. AAAI*, pages 237–242, 1984.

[7] F. Mokhatarian and A. Mackworth. A theory of multiscale curvature-based shape representation for planar curves. *IEEE Trans. PAMI*, 14:789–805, 1992.

[8] W. Richards, B. Dawson, and D. Whittington. Encoding contour shape by curvature extrema. *J. Opt. Soc. Am. A*, 3:1483–1491, 1986.

[9] P. Rosin and S. Venkatesh. Extracting natural scales using Fourier descriptors. *Patt. Recog.*, 26:1383–1393, 1993.

[10] P. L. Rosin. Representing curves at their natural scales. *Patt. Recog.*, 25:1315–1325, 1992.

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

[12] P. L. Rosin. Non-parametric multi-scale curve smoothing. *Int. J. Pattern Recognition and AI*, 8:1381–1406, 1994.

[13] B. Shahraray and D. J. Anderson. Optimal estimation of contour properties by cross-validated regularization. *IEEE Trans. PAMI*, 11:600–610, 1989.

[14] A. P. Witkin. Scale-space filtering: a new approach to multi-scale description. In *Image Understanding 1984*, pages 79–95. Ablex, 1984.