

Organization of Smooth Image Curves at Multiple Scales

DAVID G. LOWE

Computer Science Department, University of British Columbia, Vancouver, B.C. V6T 1W5, Canada

Abstract

While edge detection is an important first step for many vision systems, the linked lists of edge points produced by most existing edge detectors lack the higher level of curve description needed for many visual tasks. For example, they do not specify the tangent direction or curvature of an edge or the locations of tangent discontinuities. In this paper, a method is presented for describing linked edge points at a range of scales by selecting intervals of the curve and scales of smoothing that are most likely to represent the underlying structure of the scene. This multiscale analysis of curves is complementary to any multiscale detection of the original edge points. A solution is presented for the problem of shrinkage of curves during Gaussian smoothing, which has been a significant impediment to the use of smoothing for practical curve description. The curve segmentation method is based on a measure of smoothness minimizing the third derivative of Gaussian convolution. The smoothness measure is used to identify discontinuities of curve tangents simultaneously with selecting the appropriate scale of smoothing. The averaging of point locations during smoothing provides for accurate subpixel curve localization. This curve-description method can be implemented efficiently and should prove practical for a wide range of applications including correspondence matching, perceptual grouping, and model-based recognition.

1. Introduction

Edge detection plays an important role in many computer vision systems (and apparently in biological vision) by identifying points of intensity discontinuity in an image. The locations of these intensity discontinuities usually reflect underlying discontinuities in the geometry or surface reflectance of a scene and thereby discount the effects of varying illumination and imaging parameters. For this reason, edges have proved to be one of the most reliable low-level features for bridging the gap between image intensities and scene properties.

Unfortunately, most existing edge detectors treat edges as essentially point properties. The edge points can be linked together on the basis of the image connectivity, but it is immediately apparent upon examining these linked sets that they do not correspond to geometric properties of the scene. When edges from objects that are widely separated in depth happen to intersect in the

image, they are just as likely to belong to the same edge list as two edges of the same object. The solution to this problem lies in the area of perceptual organization [7, 8, 20], in which higher-level groupings are created according to the likelihood that they arise from underlying properties of the scene rather than accidental properties of viewpoint or imaging. In the case of curve description, the most important property on which to base perceptual organization is smoothness or continuation. This is because the edges of most objects exhibit smooth continuation at some scales, whereas it is very unlikely that two objects separated in depth will happen to have edges aligning smoothly by accident. In addition, it is useful for many higher levels of analysis, such as further perceptual organization, correspondence matching or model-based recognition, to identify the larger-scale smooth structures in the edge data and thereby to obtain stable measurements of local orientation and curvature.

Unfortunately, smoothness is not encoded

directly in the original edge data, and any attempt at a simple definition must face the effects of noise and local scatter in the positions of edge points. The major approach to this problem in the past has been to fit straight line segments [2, 9, 16, 18] and circular or elliptic curves [4, 17, 19] to portions of the linked edge data, and to look for regions that satisfy various measures of goodness of fit. While this approach works well for certain industrial scenes that contain objects with only straight or circular edges, they force the introduction of arbitrary discontinuities in the description when faced with the more general classes of image curves found in most common scenes.

Another approach taken to curve segmentation is to look for the tangent discontinuities directly with local operators (often known as corner detectors). The problem with this approach is that corner detection in the presence of noise is contingent upon scale selection and is no longer a local problem. Figure 1 illustrates this with two curves, only one of which should be assigned a tangent discontinuity at its center. Yet both of these curves are identical over most of their length, and in particular in the region surrounding this potential tangent discontinuity. Therefore, corner detection requires that a description be chosen on the basis of global properties of the curve rather than simply a local neighborhood.

Local tangent direction and curvature of a sampled curve are defined only with respect to some scale of smoothing. Due to the variable effects of noise (which leads to edge-point scatter that is typically inversely proportional to intensity gradient), it is quite likely that different scales will be appropriate for different edges in the same image. The need for different scales of analysis is even more important when dealing with natural images which may contain small variations in the scene edges themselves (e.g., the bark of a tree trunk). For the sake of higher-level analysis and stability, we would like to be able to derive the larger-scale structure of curves even when the actual scene edge is not perfectly smooth.

Therefore, we need a technique for smoothing arbitrary curves at multiple scales. The most promising candidate would seem to be smoothing with a low-pass Gaussian filter, as has been

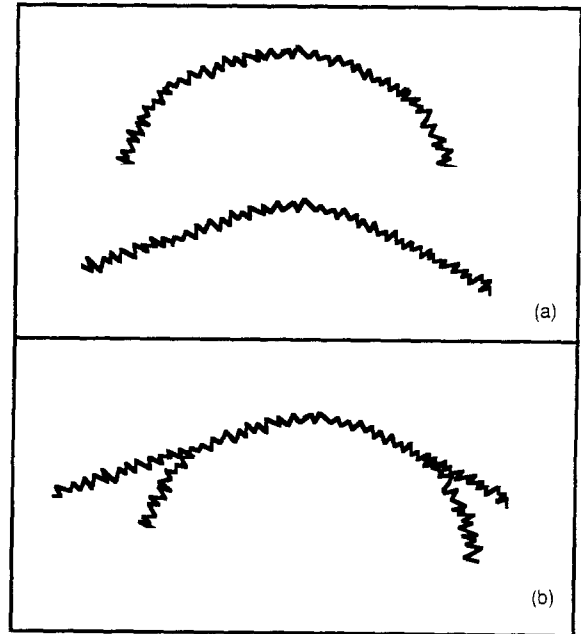


Fig. 1. Two noisy curves are shown in (a). The first curve would most naturally be described as a single circular arc, while the second would be described as two lines with a tangent discontinuity in the center. However, as is shown when the two curves are superimposed (b), they are identical over most of their length and in particular in the region surrounding the potential discontinuity. This suggests that corner detection can best be performed through a global search for smooth curve segments rather than as a local operation in the neighborhood of each potential corner.

proposed in many other areas of image analysis, since it allows for precise control in the frequencies that are filtered from the original data. Mackworth and Mokhtarian [10, 11, 14] have extensively studied the properties of smoothing two-dimensional parametric curves with Gaussians, and we will build upon their work in this paper. Similar approaches have been taken by Marimont [12] and Witkin [21]. In related method suggested by Asada and Brady [1], the parametric orientation function is smoothed rather than the coordinate functions. However, this method seems to be more suited to calculating discontinuities of curvature rather than recovering the underlying smoothed point coordinates.

One significant problem with Gaussian filter-

ing of the coordinate functions is that it results in a shrinking of the size of closed curves. The larger the curvature or the degree of smoothing, the greater is this amount of shrinkage. Horn and Weldon [6] rejected the use of a Cartesian parametric curve representation because of this problem, and instead suggested that curves be represented in a form they term the “extended circular image.” However, section 3 of this paper shows that this shrinkage effect can be compensated for in an efficient and effective manner using the standard parametric representation for curves.

Once the method of smoothing has been perfected, it can be applied at multiple scales and used to select smooth segments of the original curve. In section 4 we show that a measure of smoothness that maximizes the length of each curve segment while maintaining a low rate of change of curvature can be used for segmentation and selection of the scale of smoothing. This can be implemented in an efficient manner, and results are demonstrated for natural images.

A final issue is the relationship between smoothing of curves and smoothing of the original image. We believe that both forms of analysis must take place. It is quite possible that edge points or other feature tokens that can only be extracted from a fine-scale analysis of the original image will themselves have important larger-scale curve structure. Furthermore, it is likely that the multi-scale analysis of curves can be used to determine which scales of smoothing of the original image are most significant. We do not have a solution to the longstanding problem of combining different scales of image smoothing, but this paper does address this problem in the domain of image curves by providing a way to select from among multiple scales of smoothing for curve intervals. A biologically plausible implementation of these curve-smoothing techniques would be the use of a nonlinear operator to select and “mark” points of intensity discontinuity in an image. These marked points would then be low-pass filtered in a second stage and the resulting rate of change of curvature measured to select appropriate scales of analysis and the locations of tangent discontinuities.

2. Curve Smoothing

This section will briefly present the basic methods and terminology for filtering a curve by Gaussian convolution. The reader is referred to Mackworth and Mokhtarian [11] for a more detailed development and the proof of a number of important properties of smoothed curves.

The curve to be smoothed is represented as two coordinate functions of a path parameter t :

$$x = x(t) \quad \text{and} \quad y = y(t)$$

In order to filter out high frequencies in this curve, we convolve these functions with a one-dimensional Gaussian $G_\sigma(t)$ of standard deviation σ :

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

We will also make use of convolutions with the higher derivatives of this kernel:

$$G'_\sigma(t) = \frac{-t}{\sigma^3\sqrt{2\pi}} e^{-t^2/2\sigma^2}$$

and

$$G''_\sigma(t) = \frac{1}{\sigma^3\sqrt{2\pi}} \left(\frac{t^2}{\sigma^2} - 1 \right) e^{-t^2/2\sigma^2}$$

Define $X(t)$ as the convolution $G_\sigma(t) \otimes x(t)$ for some selected value of σ . Since differentiation commutes with convolution, $X' = G'_\sigma \otimes x$ and $X'' = G''_\sigma \otimes x$, which provides a simple, numerically stable method for computing the derivatives.

The curvature $\kappa(t)$ of a Gaussian filtered curve can then be computed in terms of these derivatives of X and Y (note that κ is equal to $1/r$, where r is the local radius of curvature):

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

Although it is true that $X'^2 + Y'^2 = 1$ for a path-length parameterized curve, it should be noted that even if the original curve is parameterized by path length, the smoothed curve will not be in general. Therefore, it is not possible to drop the denominator in the above expression.

3. Smoothing Without Shrinkage

The major difficulty with the above methods for curve smoothing is that they will systematically shrink the size of a curve toward the center of curvature. The source of this shrinkage arises from the fact that each point on a curve is being averaged with its neighbors, which in both directions curve toward the local center of curvature. Therefore, even if the curve is entirely smooth to begin with, convolution with any averaging filter will cause each point to migrate toward the center as a monotonic function of curvature and degree of smoothing. For any application in which it is important to know the location of a curve in the image, which includes most aspects of higher-level vision, this variable migration would be a critical defect.

However, since this shrinkage is due to the amount of smoothing and the local curvature, we can use the known value of σ and the measured curvature of the smoothed curve to compensate for the degree of shrinkage that must have occurred. In fact the same argument can be applied to each coordinate function independently, as the shrinkage is a result of the underlying filtering process applied separately to $x(t)$ and $y(t)$.

Our goal then will be to predict the degree of shrinkage for each point of the smoothed curve $X(t)$ as a function of degree of smoothing σ and local curvature measure $X''(t)$. Consider a circle of radius r passing through the origin and centered at the point $(r, 0)$. The coordinate function $x(t)$ for this curve, for a path length parameter t , will be

$$x(t) = r \left(1 - \cos \frac{t}{r} \right)$$

Now consider the convolution of this function with $G_\sigma(t)$

$$\begin{aligned} X(t) &= G_\sigma(t) \otimes x(t) \\ &= \int_{-\infty}^{\infty} \frac{1}{\sigma\sqrt{2\pi}} \exp\left[\frac{-(t-u)^2}{2\sigma^2}\right] r \left(1 - \cos \frac{u}{r} \right) du \end{aligned}$$

We would like to compute the value of this convolution at the point $t = 0$. Since the original curve passes through the origin at this point, the value of the convolution represents the amount of

shrinkage as a function of σ and r . The following solution was obtained with the aid of the Macsyma system for symbolic algebra:

$$X(t) = r \left(1 - e^{-\sigma^2/2r^2} \right), \quad \text{at } t = 0 \quad (1)$$

However, we do not actually know the value of the original curve radius r , but rather must make use of the measured second derivative of the smoothed curve X'' :

$$\begin{aligned} X''(t) &= G''_\sigma(t) \otimes x(t) \\ &= \int_{-\infty}^{\infty} \frac{1}{\sigma^3\sqrt{2\pi}} \left[\frac{(t-u)^2}{\sigma^2} - 1 \right] \\ &\quad \times \exp\left[\frac{-(t-u)^2}{2\sigma^2}\right] r \left(1 - \cos \frac{u}{r} \right) du \\ &= \frac{e^{-\sigma^2/2r^2}}{r}, \quad \text{at } t = 0 \quad (2) \end{aligned}$$

This last result shows that $r \approx 1/X''$, as expected, for small values of σ . In fact, this approximation can be used to correct for most of the shrinkage error for typical values of σ , since it is correct to within 13% for $\sigma < r/2$. However, in practice the shrinkage correction will be implemented by table lookup and interpolation, so we can afford to solve (2) numerically for r .

This method has been implemented and tested on a wide range of examples, with results that indicate elimination of the shrinkage effect. A table is built giving the shrinkage error values (1) as a function of the second derivative of convolution X'' . Then, for each point coordinate of the smoothed curve, we interpolate the appropriate error value and subtract it from the original smoothed value.

The results of applying this method to a noisy circle are shown in figure 2. This circle was generated by adding uniformly distributed random noise in the radial direction to points lying on a circle. Following smoothing by a Gaussian with $\sigma = 8$, we recover the smooth circle shown in figure 2(a), but it has shrunk significantly in size. However, following the method for shrinkage correction given above, we instead get the results shown in figure 2(b). Here the smoothed circle maintains the same radius as the original curve.

An example of the application of this tech-

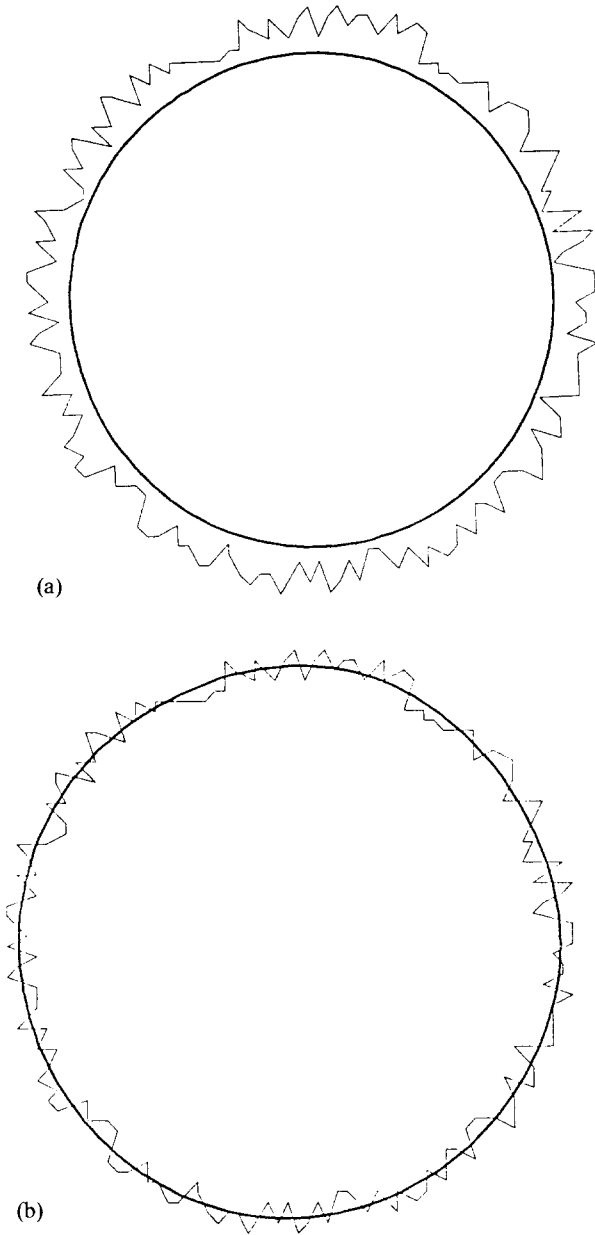


Fig. 2. A noisy circle can be smoothed with a Gaussian to recover the original smooth circle, but the radius of the circle will shrink as shown in (a). By applying the shrinkage correction technique described in this paper, see (b), it is possible to remove noise with any desired scale of smoothing while also retaining the original radius.




Fig. 3. A map of Africa is shown with Gaussian smoothing at $\sigma = 8$ using the standard method (a) and with the shrinkage correction technique (b). It can be seen that the corrected curve tracks the original edge points much more closely than does the noncorrected curve. The two curves are shown superimposed in (c), with the shrinkage-corrected curve drawn with a darker line.

(continued)




Fig. 3 (continued)

nique to a more complex curve is shown in figures 3. In figure 3(a) a map of Africa is smoothed with a Gaussian filter with $\sigma = 8$. The shrinkage is apparent in that the smooth curve is systematically displaced toward the inside of each curved region. The result of shrinkage correction is shown in 3(b). Figure 3(c) shows the smoothed curves before and after correction overlaid upon one another. This illustrates the fact that the corrected curve has zeros of curvature at the same locations as the uncorrected curve. However, additional inflection points may occasionally be introduced when the original curve has local minima of curvature that are positive or maxima that are negative.

4. Identifying Smooth Curve Segments

Given the ability to smooth a curve at different scales, it is necessary to develop some way to determine which intervals of the curve and which scales of smoothing are most likely to reflect the underlying structure of the scene. As described earlier, when two independent objects project to

the same region of an image, their edges are unlikely to align smoothly; therefore, any list of edge points containing two independent object edges will likely contain a tangent discontinuity where the edges meet. This might lead one to believe that we should look for points of low curvature along an edge, as these are likely to be between the tangent discontinuities, which in theory have infinitely high curvature. However, the underlying image data is noisy and must be smoothed to recover a curvature estimate. Once this is done, a tangent discontinuity in the scene will often have a lower curvature than some genuinely curved edge in the image. Given that any curved object edge can project to arbitrarily high image curvature by simply becoming more distant from the camera, we can expect many instances of high image curvature in any natural scene.

Therefore, we have found that the third derivative, or rate of change of curvature, is a more useful measure of the underlying degree of smoothness of an edge. Edges that have a high curvature that is changing only slowly will still be considered smooth. Perhaps the reason that the rate of change of curvature has not been given more consideration in earlier work on segmentation is that it is often assumed that higher derivatives are very sensitive to noise. An underlying reason for the noise sensitivity is that higher derivatives tend to amplify high frequencies, which is where local forms of noise have their major impact. However, in the case of a Gaussian filtered curve, the high frequencies have all been removed and therefore do not influence the result of higher derivatives. Simple inspection of the shape of the third-derivative kernel will show that it is only slightly more responsive to higher frequencies than the first derivative. Interestingly, Binford [3] has argued for the importance of detecting discontinuities of curvature for high-level segmentation, which would naturally be detected by a third-derivative operator.

There are two criteria that must be balanced in selecting the smoothest segments to represent a curve. One is to minimize the rate of change of curvature, and the second is to maximize the lengths of the curve intervals that are described by a single segment. Surprisingly, perhaps, we found that the second criterion tends to override the

first. Even if some interval at some scale has a very low rate of change of curvature, it will tend to introduce a false discontinuity if we select this interval over an interval at another scale that covers a longer portion of the curve. We can see the reason behind this if we consider that the goal of segmentation is to uncover true tangent discontinuities such as occur at the intersection between two object edges, but to otherwise find some scale that will represent the underlying smooth structure of an edge. Although we have experimented with many much more complex criteria, the best method appeared to be one that simply sets a threshold on the size of the rate of curvature change that is sufficient to eliminate most tangent discontinuities, and to search across multiple scales of smoothing for intervals that cover the maximum length of the curve. Once one interval is selected, that portion of the curve is removed from consideration and the same method is applied to the remaining portions.

One other issue that must be addressed is the smoothing of curves out to the termination of the underlying list of edge points. The convolution kernels are defined over an infinite range, but can be safely truncated at a distance of 3σ from their center. Nevertheless, without some special method for handling termination, this would leave an undefined region of 3σ at the end of each curve, which would be a serious loss of data for most practical applications. There appears to be no ideal solution to this problem, but good results were achieved by estimating the endpoint tangent by extrapolating from the measured tangent and curvature of the closest curve point for which a reliable estimate is available. The curve is reflected about an axis normal to this extrapolated tangent to provide data for smoothing up to the endpoint. This method of extrapolating from local tangent and curvature measurements is similar to that used to define co-circularity by Parent and Zucker [15].

5 Implementation and Results

All of the methods described above have been implemented in Sun/Lucid Common Lisp. Edges are first detected by the Canny [5] edge finder and

are linked on the basis of image connectivity to yield lists containing edge points that are one pixel apart. While this does not produce exact path-length parameterization, it is close enough that any deviation will only have a minor effect on the local scale of smoothing. These lists are then used as input to the following sequence of operations:

Initial Smoothing. Each edge is smoothed by Gaussians at a range of scales with σ increasing by a factor of $\sqrt{2}$ from one scale to the next. In the examples to be presented, 7 scales of smoothing were used with σ ranging in value from $\sqrt{2}$ pixels to $8\sqrt{2}$ pixels. The first and second derivatives are also calculated at each point by convolution with the appropriate kernels, and the shrinkage compensation is applied. Curvature κ' is computed by using the finite difference of points that are σ units apart. Curvature is scaled by a factor of σ and change of curvature by a factor of σ^2 to make them scale invariant, so that a single threshold can be used across all scales.

Interval Formation. The linked lists of edge points at each scale are broken into intervals in which all points in an interval have change of curvature below some scale invariant threshold. For the following examples we chose a threshold value of $\sigma^2 \kappa' < 0.2$. Reducing this threshold forces curves to be smoother, at the cost of introducing more discontinuities into the description. There is also a minimum-length threshold of 2σ required for each interval, which prevents zero-crossings of κ' near corners from being considered as short smooth intervals.

Interval Selection. We consider all of the intervals at all scales for a given edge, and select the interval that covers the greatest length of the original edge list. This interval is extended at each end using the method for handling terminations described above. Then the portions of all other intervals that overlap this selected interval are removed from consideration, and the selection process is repeated. This results in the final set of selected intervals covering as much of the original curve as possible. The output curves are represented as a sequence of linked points that are spaced at a separation proportional to the

measured local radius of curvature (this is a user-defined parameter that trades off the degree of data reduction with the maximum angular separation of adjacent points). The specification of each output point includes the smoothed location of subpixel accuracy, the tangent direction, the curvature, and the scale of smoothing.

The output of this smoothing and segmentation process for some realistic examples are shown in figures 4 and 5. The image shown in figure 4 is of a totem pole digitized from a grainy photograph originally taken in 1896. The Canny edge finder was used to produce the linked edge points shown in figure 4(b). The results of applying the smoothing and segmentation methods described above are shown in 4(c). The displayed width of each smoothed output curve is proportional to the σ of smoothing used for that curve. In general, the method has been successful at selecting scales of smoothing that remove large



(b)



(a)




(c)

Fig. 4. (a) A noisy image of a totem pole digitized from an 1896 nitrate negative (reproduced by permission of the B.C. Provincial Museum), (b) Edges detected by the Canny edge finder, (c) The final curve segments output by the multiscale smoothing and segmentation algorithm. The thickness of each curve is proportional to the scale of smoothing selected for that curve.



(a)



(c)



(b)

Fig. 5. Similar results as for figure 4, but for a higher quality image of an industrial-type scene.

amounts of noise and yet correctly identify locations of tangent discontinuity. The noisy totem pole image illustrates the capability for recovering the underlying scene curvatures from degraded edges without a high degree of initial image blurring that would otherwise merge edges and lead to a loss of data. Figure 5 contains an example of the results for an image of the type found

in current robotics applications. In most cases, the tangents and curvatures of these smoothed edges seem to provide good estimates for the projected values of the underlying scene curves. The accuracy of the curve smoothing can be judged from the greatly enlarged examples of figure 6, showing the smoothed curves superimposed on the original linked edge points.

The current implementation of this system is in Common Lisp and was not designed with a great concern for efficiency. Running time is about 2 minutes on a Sun 3/60 for these examples. However, there is every reason to believe that it can be implemented as efficiently as any other curve segmentation method. In the current Common Lisp implementation, each scale of Gaussian smoothing is calculated independently, yet much greater efficiency could be achieved by an incremental method that bases the results at larger scales of smoothing on the results of each previous scale.

6 Conclusions and Future Research

The ability to combine smoothing and segmentation at multiple scales is an important capability for many applications of computer vision. In most cases, the final output of the method de-





Fig. 6. (a) An enlarged view of the linked edge points from around the eye of figure 4, showing the pixel quantization

error. (b) The smoothed curves superimposed on these edges, demonstrating the accuracy of interpolation.

scribed in this paper can be expected to reflect the underlying structure of the scene in terms of segmentation, scale, tangent direction, and curvature. This means that these descriptions will tend to remain stable across changes in viewpoint and imaging conditions, and therefore could play an important role in correspondence matching in stereo or motion. The increased measurement accuracy in the locations of edges, resulting from the local averaging, should also be of value in stereo or motion interpretation. The specific application that we intend to develop is in model-based vision, in which smooth curve segments can be matched to models with arbitrarily curved surfaces and markings. The stability of tangent and curvature measurements should allow these measurements to play an important role in model-based matching.

Possibly an even more important application will be in the area of perceptual organization. The smoothing and segmentation process is itself an aspect of perceptual organization, as it involves identifying higher-level structures in the linked edge data on the basis that such smooth curves are unlikely to arise by accident from independent scene edges. But these smooth curves can also play an important role in later stages of grouping, which are based upon curvilinearity, parallelism, proximity of terminations, and other relationships [8]. Since these forms of grouping require local tangent and curvature estimates as well as segmentation at tangent discontinuities, they could not be applied to the original linked edges without this higher level of smoothing.

There are a number of areas in which the methods described in this paper could be improved. One straightforward improvement would be to use an edge finder that interpolates the position of edge points to subpixel accuracy. This should provide a substantial improvement to the segmentation results at the finest scales of smoothing, as these are currently being dominated by pixel quantization effects. Another useful improvement would be to allow for retention of multiple scales of description for a single curve segment when they are qualitatively different. While this would tend to clutter and detract from the appearance of output to a human observer, these multiple scales of description could prove

useful for many higher-level matching and organization processes. There is also clearly a need for further study of the underlying theory of optimal detection of tangent discontinuities in the presence of varying noise and curvature.

An important problem for further research is in combining these techniques for curve description with multiscale methods for the underlying edge detection. The use of a curve smoothness criterion allows a second dimension of analysis to be used to select among multiple scales of description, in addition to the scale-space behavior of edge points as suggested by Marr and Hildreth [3] and Witkin [21]. While it is true that an edge will tend to have a stable position across a range of scales, our own empirical examination of images shows that many nonedges appear to also have this behavior. Thus we hypothesize that the use of smoothness criteria along the length of an edge will prove necessary for selecting among multiple scales of image smoothing. Some recent work by Zucker et al. [22] provides a biologically plausible model that combines early vision with the inference of curve properties.

Availability

In order to facilitate the further development and use of these methods, the original Common Lisp code is being made available to any researcher who would like to experiment with these techniques. This code can be requested through electronic mail to "lowe@vision.ubc.ca."

Acknowledgements

I would like to thank Farzin Mokhtarian, Alan Mackworth, and Bob Woodham for valuable discussions and assistance with this work.

The image in figure 4 shows the Nan-wa-ka-wie (Wise One) figure from the Wakas Totem Pole carved in 1893 in the Nimpkish village of Yilis at Alert Bay, British Columbia. It is from an 1896 nitrate negative, obtained by Ann Howatt Krahn, Art Conservator, courtesy of the B.C. Provincial Museum. Digitization was done in the UBC

Laboratory for Computational Vision by Nedenia Krajci.

The author is a Scholar of the Canadian Institute for Advanced Research. This work was also supported by the Natural Sciences and Engineering Research Council and by the National Science Foundation under grant DCR-8502009.

References

1. Haruo Asada and Michael Brady, "The curvature primal sketch," *IEEE Trans. PAMI* 8(1):2-14, 1986.
2. D.H. Ballard, "Strip trees: A hierarchical representation for curves," *Comm. ACM* 24(5):310-321, 1981.
3. Thomas O. Binford, "Inferring surfaces from images," *Artificial Intelligence* 17:205-244, 1981.
4. R.C. Bolles, P. Horaud, and M.J. Hannah, "3DPO: A three-dimensional part orientation system," *Proc. 8th Intern. Joint Conf. Artif. Intell.*, (Karlsruhe, West Germany), pp. 1116-1120, 1983.
5. John Canny, "A computational approach to edge detection," *IEEE Trans. PAMI* 8(6):679-698, 1986.
6. B.K.P. Horn, and E.J. Weldon, "Filtering closed curves," *IEEE Trans. PAMI* 8(5):665-668, 1986.
7. David G. Lowe, and Thomas O. Binford, "Perceptual organization as a basis for visual recognition," *Proc. 3rd Nat. Conf. Artif. Intell.*, pp. 255-260, August 1983.
8. David G. Lowe, *Perceptual Organization and Visual Recognition*. Kluwer Academic: Boston, 1985.
9. David G. Lowe, "Three-dimensional object recognition from single two-dimensional images," *Artificial Intelligence* 31(3):355-395, 1987.
10. A.K. Mackworth, and F. Mokhtarian, "Scale-based description of planar curves," *Proc. 5th Canadian Soc. Comp. Studies Intell.* London, Ontario 114-119, May 1984.
11. A.K. Mackworth, and F. Mokhtarian, "The renormalized curvature scale space and the evolution properties of planar curves," *Proc. IEEE Conf. Comput. Vision Pattern Recognition*, Ann Arbor, June 1988.
12. David H. Marimont, "A representation for image curves," *Proc. 4th Nat. Conf. Artif. Intell.* Austin, Texas, pp. 237-242, August 1984.
13. David Marr, and Ellen Hildreth, "Theory of edge detection," *Proc. Roy. Soc. (London)* B 207:187-217, 1980.
14. F. Mokhtarian, and A.K. Mackworth, "Scale-based description and recognition of planar curves and two-dimensional shapes," *IEEE Trans. PAMI* 8(1):34-43, 1986.
15. Pierre Parent, and Steven W. Zucker, "Trace inference, curvature consistency, and curve detection," Tech. Rep. CIM-86-3, Computer Vision and Robotics Lab, McGill University, 1985.
16. T. Pavlidis, *Structural Pattern Recognition*. Springer-Verlag: New York, 1977.
17. A.P. Pridmore, J. Porrill, and J.E.W. Mayhew, "Segmentation and description of binocularly viewed contours," *Image and Vision Computing* 5(2):132-138, 1987.
18. L.G. Roberts, "Machine perception of three-dimensional solids." In *Optical and Electro-optical Information Processing*. J. Tippet (ed.) MIT Press: Cambridge, MA, 159-197, 1966.
19. Yoshiaki Shirai, "Edge finding, segmentation of edges and recognition of complex objects," *Proc. Intern. Joint Conf. Artif. Intell.* 75:674-681, 1985.
20. Andrew P. Witkin, and Jay M. Tenenbaum, "On the role of structure in vision." In *Human and Machine Vision*. J. Beck et al. (eds.) Academic Press: New York, 481-543, 1983.
21. Andrew P. Witkin, "Scale-space filtering," *Proc. 8th Intern. Joint Conf. Artif. Intell.*, Karlsruhe, West Germany, pp. 1019-1022, August 1983.
22. Steven W. Zucker, Chantal David, Allan Dobbins, and Lee Iverson, "The organization of curve detection: Course tangent fields and fine spline coverings," *Proc. 2nd Intern. Conf. Computer Vision* pp. 568-577, Tampa, Florida, Dec. 1988.