Recognizing Planar Curves Using Curvature-Tuned Smoothing

Gregory Dudek John K. Tsotsos
Dept of Computer Science, University of Toronto
Toronto, Canada M5S 1A4
email: dudek@ai.toronto.edu

Abstract

In this paper the authors present a new technique for both the smoothing and decomposition of planar curves. This technique, dubbed "curvature-tuned smoothing" provides for robust rotation and translation invariant smoothing of planar curves. As the smoothing is performed, parts are extracted robustly and at multiple scales. These parts are subsequently usable for object recognition. The parts extracted correspond to regions of roughly uniform curvature and constitute a rich description of the original data. This representation describes some regions at multiple scales since multiple structures may co-occur.

1 Introduction

In this paper the authors present a technique for object description and representation based on the robust extraction of curvature data. The technique is based on the decomposition of the curve of interest into segments of roughly uniform curvature. From these segments a symbolic description of the input curve is constructed. The description allows the data to be dealt with either in terms of its coarse structure, or based on finer scale properties. Furthermore, because the representation is derived from local properties, it allows recognition to proceed despite the occlusion of parts of the object.

The recognition of planar curves is a significant problem for several reasons. Two-dimensional data can be directly applied (without 3-dimensional reconstruction) and its value has been long recognized by perceptual psychologists (Pomerantz, Sager and Stoever 1977). For example, there exist navigation tasks where two-dimensional world structure is a valuable cue to location (Brooks 1981; Dudek et al. 1988, are examples). Furthermore, it is a longstanding open problem and serves as a stepping stone to the related problem of shape and curvature description in three dimensions. The approach described here can, in fact, be generalized to three dimensions.

1.1 Related work

The utility of curvature information as a cue to object identification has been recognized for some time. Attneave illustrated the perceptual significance of curvature maxima over thirty years ago (Attneave 1954). More recently Lowe has shown that Attneave's perceptual results can be explained equally well by a representation based on curvature inflection points (Lowe 1985).

The physical relevance of curvature minima motivated the representation of contour using "codons", a curve description grammar that decomposes curves at local minima, and also uses curvature zero-crossing information (Hoffman and Richards 1984; Richards and Hoffman 1984). This decomposition parses curves into a vocabulary of standard parts, however, it depends on the curve being pre-smoothed appropriately so that specific structures of interest can be isolated from noise and other independent structures. The close relationship between curvature information and axes of symmetry has been examined by Leyton (Leyton 1987). The curvature primal sketch (Asada and Brady 1986) uses curvature information and tracks the derivative of curvature across scales to produce a syntactic description of curves. The fact that explicit parts are extracted by both these approaches is a very useful and appealing characteristic. Contour information has traditionally been easier to extract from images than volumetric information since it is often directly accessible. As a result, contour information, and curvature in particular, has been used as a practical recognition cue by several researchers (Connell and Brady 1985; Bhanu and Faugeras 1984; Brooks 1981; Kehtarnavaz and deFigueiredo 1988; Milios 1988; Pentland 1988).

Curvature derivatives, curvature extrema and, in fact, curvature zero-crossings are not, however, stable with respect to minor perturbations of the curve. A complete description of curve structure therefore depends on being able to deal with unstable characteristics of the data and structures at a variety of scales. In order to capture and make explicit important structure without excessive sensitivity to details, a multi-scale representation is required (Nishihara 1981).

The stable extraction and measurement of curvature information in the presence of noise has been dealt with in several ways (Besl 1988; Lowe 1988; Zucker et al. 1988). One characteristic of most existing curvature-measurement techniques is the assumption that there is a unique curvature that can be measured at each point. [1] While this is, of course, true in the analytic case, the assumption introduces significant difficulties for inverse problems involving noisy signals, such as those that occur in vision. Despite the appealing results that have been achieved by some researchers, the need for scale-specific operators (which also manifests itself as the need for the choice of a best smoothing scale, or the choice of an appropriate neighborhood for measurements) to deal with noise problems causes an *inherent preference* for certain ranges of curvature value and involves strong implicit assumptions about the underlying signal. The actual curvature of a signal depends on what we call noise and what we call signal, and consequently may take on differing values depending on our goals.

A substantial body of work exists dealing with scale-space filtering, examining the properties of zero-crossings of the second derivative of a one-dimensional signal under blurring (Witkin 1983). In general, curve description by blurring techniques captures the structure of regions between zero-crossings by observing their stability over scale-space. It may be that explicitly extracting and representing the interesting structures at various scales, as in the technique described here, has advantages over observing implicit properties of the structures in the image. In fact, what we require is a representation that makes useful characteristics of the data explicit so that they can be readily accessed and used by subsequent processes. By using curvature information directly, we avoid several degeneracies that may occur with zero-crossing based representations and directly extract those features of the data that we are interested in using for recognition.

[1] This assumption is often only stated implicitly. Some techniques use multiple scales at an intermediate filtering stage, but then concentrate on selecting the single correct curvature at each point.

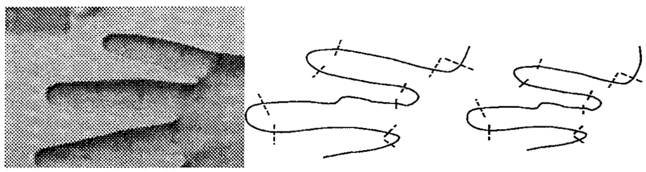

Figure 1: An image of a human hand and two boundary curves extracted by simple-minded segmentations at different scales. Regions of roughly uniform curvature, of the sort extracted by the method, are delimited by dashed lines. Note that although the extracted curves differ in scale, the *relative* curvature structure is the same. Observe also that each high curvature region has naturally limited spatial extent. (The bump of the middle finger is an artifact of the grey-level segmentation processes. Delimiting lines were abstracted and simplified manually to illustrate the concept; the actual regions extracted overlapped.)

1.2 Regularization

The term "regularization" (or "variational regularization") has been applied to techniques for converting ill-posed problems to well-posed ones; that is, guaranteeing that the problem being solved has a solution, preferably a unique one, and that it is continuous. Poggio and his colleagues first demonstrated the significance of ill-posedness to computational vision and the relevance of regularization as a general framework for dealing with a wide variety of vision problems[2] (Torre and Poggio 1984; Poggio, Torre and Koch 1985). As typically formulated, regularization takes the form of transforming an original problem of the form

$$A\,z = y \tag{1}$$

where z is an unknown solution matrix to be found and A is a (transformation) matrix, and y is the input data, to a new (minimization) problem of the form

$$||A\,z - y|| + \lambda||Pz|| \tag{2}$$

where $||Pz||$ is a "stabilizing functional" and λ is the regularization parameter. As usually applied in computational vision, the stabilizer is usually a low-order derivative of z enforcing smoothness of the final solution, and λ determines the tradeoff between the amount of smoothing and the closeness to the original problem. The explicit and practical application of regularizing techniques as smoothness constraints for surface reconstruction were examined by Grimson (Grimson 1981) and further pursued by Terzopoulos (Terzopoulos 1984). These authors made use of regularizers based on first and second derivatives; smoothing analogous to the minimization of bending energy in flexible membranes and thin solid plates. It was shown that regularization not only allowed a variety of surface interpolation and smoothing tasks to be conveniently and elegantly described in a formal setting, but that algorithms based on these techniques could be made reasonably efficient.

It is crucial to observe that the form of the stabilizing functional is of great significance in determining the final solution. This functional corresponds to an *a priori* assumption regarding the nature of acceptable world models; it coerces functions that are inconsistent with the assumption into other forms in a well-behaved fashion. If our goal in a vision problem is to detect and extract object parts with given characteristics, it is of great value to select a stabilizing functional that is consistent with this goal.

[2] In fact, a variety of early vision techniques from the early 1980s used constraints that are actually regularizing functions.

1.3 Smoothing

In order to improve stability of image processing operations and deal with noisy data, it is usually necessary to apply smoothing to natural images. Furthermore, in many cases an assumption of smoothness provides not only robustness in the face of noise, but a crucial constraint that makes an otherwise underdetermined system soluble. The smoothing process, however, also alters the character of the data. Smooth data is, as often applied, that with minimized curvature. One conventional approach to smoothing to is solve a variational minimization problem such as the "thin-plate" model based on minimizing:

$$\int (u(t) - d(t))^2 + \lambda^2 (u''(t))^2 dt \tag{3}$$

where we wish to determine $u(t)$, the smoothed data, given $d(t)$, the input data. We do this by finding the function $u(t)$ that minimizes this functional over the region of interest. This function $u(t)$ is a compromise between the original data (the constraint of the first term) and our *a priori* smoothness criterion (the second term). Another common smoothing technique is Gaussian blurring, where the curve's function is convolved with a Gaussian.

These techniques and others like them typically smooth the image in a single uniform manner and then apply a second stage for picking models to describe the data. The degree of smoothing is also, as noted earlier, a goal-dependent problem. The structures of interest determine the required amount of smoothing. Even on already-ideal data, smoothing will usually have some deleterious effect, hence altering the character of subsequent computations.

What we propose is to perform a unitary operation which is specifically tuned to the model classes we wish to select. In short, we wish to smooth the data so that curvature information is not distorted. An appropriate smoother for application to planar curves should also be rotation and translation invariant. When our input is ideal data, for example a perfect circle, the effect of the smoothing is nil. In this way, potential model fits (segments of constant curvature) are not distorted by the results of the smoothing operation. This is in contrast to conventional smoothing, which can potentially distort the results of the model-search process and yield inappropriate measurements even with ideal data.

1.4 Curvature

We define a curve, $d(t)$, in the plane, parameterized by arc length t. It is given by the function:

$$d(t) = (d_x(t), d_y(t)) \tag{4}$$

and its curvature κ is given by

$$\kappa(t)^* = \frac{d'_x(t)d''_y(t) - d''_x(t)d'_y(t)}{\sqrt{{d'_x}^2 + {d'_y}^2}} \tag{5}$$

(where $'$ indicates differentiation with respect to t) which reduces to

$$\kappa(t) = d'_x(t)d''_y(t) - d''_x(t)d'_y(t). \tag{6}$$

due to the arc-length parameterization. It has been frequently observed that curvature properties provide an apparently powerful cue to the underlying structure of the curve. The structure of a plane curve is completely captured by it curvature function, which is, in turn, an invariant property of the curve itself and is hence insensitive to changes in the coordinate system (such as rotation). Furthermore, the curvature function is smoothly and monotonically related to the scale of the structure of the curve (i.e. the radius of curvature at any given point). As such, local relationships between the curvature of different parts of the curve are retained even when the curve changes size (and its absolute curvatures change, see fig. 1). This immediately suggests that the relative, "topological" structure of the curvature function may provide a scale invariant framework.

2 Technique

The technique we describe here is based on smoothing and subdividing an input curve by solving a set of variational optimization problems. We perform smoothing with discontinuities using a set of curvature-tuned minimization functions. The differing versions of the smoothed data provide a decomposition of the original curve.

The definition of smoothness has traditionally been closely related to the minimization of one or more of the derivatives of the curve. In fact, many approaches to curve description or decomposition presuppose that the data have been smoothed at an "appropriate" scale in advance. In general, a single "appropriate" degree of smoothing is very difficult to determine for any given set of data; too much smoothing tends to obliterate structure in the underlying data, too little fails to remove noise and instabilities.

In general terms, we use a family of smoothing functions that define smoothness as the degree to which the function's curvature matches a target value. Specifically, we propose a family of smoothing measures of the form:

$$\epsilon(t, c_i) = ||u(t) - d(t)||^2 \; + \; \lambda^4(\kappa(t) - c_i)^2 \tag{7}$$

and

$$Energy, E(u, c_i) = \int \epsilon(t)\; dt \tag{8}$$

where $d(t)$ is the original data, $u(t)$ is a smoothed version of the data, and c_i is a constant that defines the target value for the curvature function $\kappa(t)$ (i.e. c_i is the *target curvature*). By minimizing this functional, E, over the curve while allowing the introduction of discontinuities, we obtain a segmented version of the curve. Repeating this process over the range of values for c_i a set of decompositions of the curve are obtained. Each value of c_i corresponds to selecting a different class of structures from the curve data; structures whose curvature can be approximated by the value of c_i. Observe that for c_i equal to zero, we have a smoothing operation very much like the aforementioned small-deflection thin-plate/thin-beam regularization. In fact, the $c_i = 0$ case corresponds to a more accurate description of thin-beam bending energy than the usual second derivative approximation. Similarly, the physical analogue for arbitrary c_i values is a bendable material with a natural degree of curvature and a propensity to shatter.

The value of ϵ over each segment along with the segment's length indicate the appropriateness of that curvature value as a description of the portion of the curve to which the segment applies. When the spacing of discontinuities is small relative to c_i, this suggests that the curve could not naturally be described using that degree of curvature. When long low-energy segments are extracted, this indicates that the corresponding value of c_i was well matched to the underlying data. For example, dense data from a circle of curvature q will be expressed as a single segment at target curvature q, while at other target curvatures substantially different from q the energy measure ϵ will induce a partitioning into many small segments with relatively high energy.

2.1 Local Interactions

As a variational problem, we can examine the Euler-Lagrange differential equation associated with this minimization and it's influence function (the Green's function for the locally linearized approximation). The form of the influence function expresses the sensitivity of the solution to local changes in the data. The rate of decay of the Green's function predicts the locality of interactions in terms of the final solution to the variational problem. The Euler-Lagrange equation for the functional being minimized is somewhat complex. We can approximate it by assuming that the the curve remains arc-length parameterized, which is the form of the input data. This leads to a a pair of (linearized) differential equations of the form:

$$\begin{aligned}\epsilon_x = \quad & 2(x(t) - d_x(t)) \; + \\ & 2\lambda^4(y''(t)A + A'y'(t)) \; - \\ & 2\lambda^4(2A'y'' + y'''(t)A + y'A) \; = \; 0\end{aligned} \tag{9}$$

$$\begin{aligned}\epsilon_y = \quad & 2(y(t) - d_y(t)) \; - \\ & 2\lambda^4(x''(t)A + A'x'(t)) \; + \\ & 2\lambda^4(2A'x'' + x'''(t)A + x'A) \; = \; 0\end{aligned} \tag{10}$$

where

$$A \; = \; u'(t) \times u''(t) - c_i = x'(t)y''(t) - y'(t)x''(t) - c_i \tag{11}$$

and hence

$$A' = u'(t) \times u'''(t) \tag{12}$$

$$A'' = u''(t) \times u'''(t) + u'(t) \times u''''(t). \tag{13}$$

Unfortunately, the associated Green's function for this problem has not been determined analytically at the time of this writing. We have, however, numerically estimated the response to local perturbations and found that the influence function can be approximated by an oscillatory function with an exponentially decaying envelope. This is consistent with expected behavior for Green's functions for similar physical systems, as well as the Green's function for conventional second-derivative "thin-beam" smoothers (Blake and Zisserman 1987).

2.2 Discontinuities

Discontinuities are inserted into the smoothed function iteratively when the local energy after convergence exceeds the discontinuity threshold, ξ. In principle, a discontinuity is inserted where the local energy function is maximal (and above ξ), then the minimization is re-solved, and the next discontinuity is inserted. In practice, because of the local nature of the interactions and the rapid decay of the linearized smoother's Green's function, multiple discontinuities can be inserted in a single pass and the function can be resolved rapidly.

Unfortunately, the discontinuity detection process introduces unavoidable non-linearities in the process of finding the optimum solution. Pairs of discontinuities, for example, may lead to a single discontinuity being detected where two should actually occur, or the placement of the adjacent discontinuities may be incorrect. Fortunately, several techniques for extracting the discontinuities in similar problems in a reasonably stable manner exist (Terzopoulos 1983; Terzopoulos 1986; Blake and Zisserman 1987). Even more important, however, is the fact that discontinuities interact with one another primarily only when they are closely spaced. This is the case because the influence function associated with this minimization problem decays rapidly with distance. When discontinuities are well separated, their interaction and hence the severity of the non-linearity they introduce is greatly reduced. Since we are interested in finding areas where an underlying curve is well described at a given curvature, and hence where the size of the segments that serve as "good" descriptions is comparatively large, the interaction between discontinuities for the segments that serve as the most appropriate parts will generally not be a major difficulty.

As a result of the non-linearity nature of the system being solved, it can, in principle, admit local minima. In practice this are not of great concern for two reasons: local minima in this type of system do not appear to be physically plausible (Tauchert and Lu 1987) and we are interested in solutions to the minization problem that occur close to the initial data, hence drastically constraining the solution space for the minimization.

2.3 Signals to symbols

Fitting the curve with smoothing functions having different curvature tunings, c_i, produces several different sets of discontinuities; one for each curvature tuning. The regions between discontinuities at a given tuning are those segments of the original curve that can be described well as a curve of curvature roughly c_i. These segments are analogous to the straight-line segments that are produced by a linear approximation to a curve. Unlike straight-line segments, however, there may be multiple segments for any given region of the curve corresponding to the multiple curvatures at which it can be described.

The set of segments taken together form a multi-scale structure, with curvature being the scale parameter. Curvature serves naturally as a scale parameter since the maximum extent of a region with a given curvature is proportional to its radius of curvature (i.e. a circle of circumference $2\pi/k$).

We can also compute the mean fitting energy of each segment – that is, how well it serves as a description of the underlying data. By extracting the largest and lowest-energy elements from this structure, we can determine the alternative descriptive structures that are most salient for the given curve. It appears that by using only the set of segments that have locally minimal energy over scale, a powerful description of the curve can be extracted. As a result, smooth data can be described in a compact form since it is composed of rather large segments, each one capturing the structure of a substantial region of the data.

3 Empirical constraints

An implementation of the above decomposition was carried out and the results of its application are described below. The implementation iteratively inserts discontinuities at local energy maxima using a greedy algorithm. This suffices to demonstrate the basic characteristics of the representation, particularly since precise discontinuity location is not as crucial as the detection of large discontinuity-free regions.

An crucial parameter in the application of the technique is the sampling density in curvature space. That is, the particular curvature tunings, c_i values, at which the smoothing and decomposition is performed. The maximum curvature used for sampling reflects the maximum curvature for which structure will be detected in the original data. Any variations of the data at higher curvatures will fail to produce a good match at any scale and thus will not be faithfully represented. The spacing between curvature space samplings is related to the ability of the representation to discriminate between curves with similar curvatures. If the objects of interest are known to be discriminable from specific ranges of curvature, then the resolving power of the representation can be focused there, while sampling other parts of curvature space more sparsely.

The constant that determines the amount of "smoothing", λ, must be chosen to give an appropriate trade-off between the local measurements and the global model. The choice of this constant is determined by observing that it is directly related to the interaction range of the resulting smoothing. An noted above, the influence of local perturbations falls off exponentially with distance. The λ parameter determines the rate of decay. Further, since we are dealing with curves, we can conclude that we do not expect regions of a given curvature to have an extent greater than the circumference of a circle with a corresponding curvature. Hence, the maximum extent of a region with curvature c_i is $2\pi/c_i$. This implies that at any point, the region of support for the smoothing should be no greater than this. We can define the region of support for the exponential Green's function as the region over with is exceeds some arbitrary constant (eg. 0.5). At the edge of the region of support, which has extent $2\pi/c_i$ the Green's function envelope gives us

$$e^{-\frac{t}{\lambda^4}} = e^{-\frac{1}{\lambda^4 c_i}} = \frac{1}{A} \tag{14}$$

where A is a constant. Hence the relationship between the smoothing constants at each scale is given by:

$$\lambda^4 = \frac{ln(A)}{c_i} = \frac{\lambda_0^4}{c_i} \tag{15}$$

where λ_0 is the value chosen for the base case of $c_i = 1$. [3]

As the value of the discontinuity threshold, ξ, increases, the number of segments used to represent the data decreases, since energies that might have caused a discontinuity to be inserted no longer do so. On the other hand, the amount of change in the data needed to produce a change in the representation will also increase. Thus, ξ is analogous to an error or resolution parameter for the fitting process, in the sense that increasing values of ξ produce sparser sets of segments (each being larger), which fit the data with less fidelity. Since both λ at c_i are scale factors, the value of ξ can be selected as a function of λ on the difference between c_i and c_{i+1}.

The following figure (fig. 2) illustrates the result of applying this decomposition to a curve with structures at several scales. The portions of the decomposed curves marked with darker lines are the components that are locally minimal in energy and used to make up the representation.

It can be observed that the marked sections of the curve appear to be major structures of the curve. Any section of a curve

[3] The properties of the data and model allow us to determine the appropriate value for λ_0 in a principled manner.

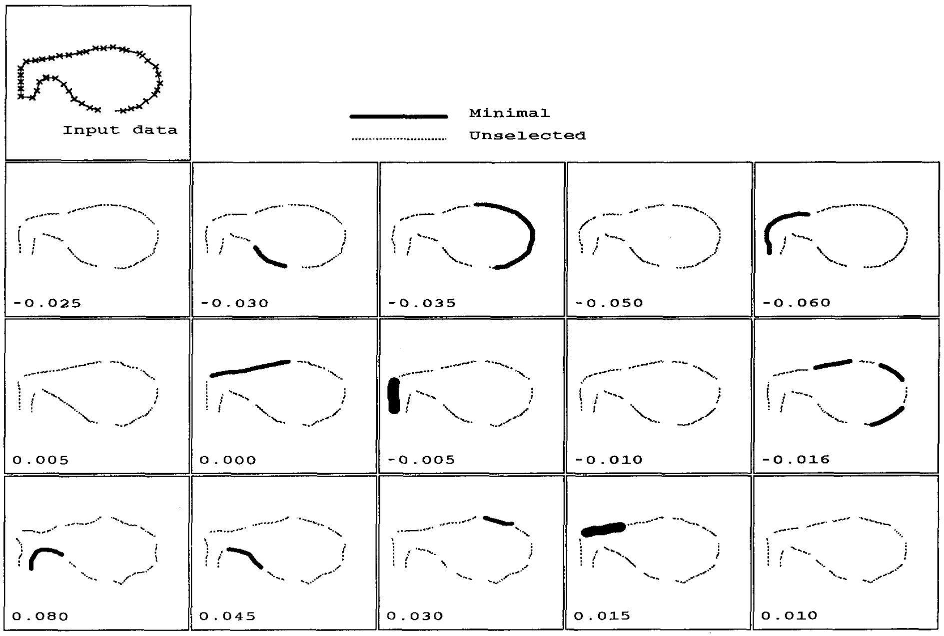

Figure 2: A curve decomposed at a variety of curvature scales. Each box is the decomposition (and smoothing) at a different curvature tuning. The small number in each box indicates the curvature value for the tuning. Low-energy (i.e. minimal) segments are marked by heavy lines.

having relatively uniform curvature will produce a low energy at some curvature scale. If this region is of a different curvature from that of its neighbors, it will be bounded by discontinuities. The threshold determining how different the neighboring curvatures must be is controlled by the parameter ξ. Note, however, that merely concatenating the major segments does not produce a reconstruction of the original curve. Segments from different scales correspond to decompositions of the curves which are not directly compatible since, among other things, they are based on incompatible partitionings of the underlying data. Although we can relate the positions of the segments from different scales to one another, inferring a single structure is a more complex, and potentially unnecessary, process. Although it appears that the multi-scale representation can be shown to be complete under appropriate sampling conditions, the existence of a reconstruction procedure is neither necessary to establish this result, nor is it even a particularly critical requirement for a useful representation.

4 Discussion

The representation presented above would appear to have a variety of desirable characteristics. It is multi-scale, and hence deals with noise naturally. It allows arbitrary degrees of precision in describing the underlying data. It has intuitively appealing characteristics, decomposing curves in readily comprehensible manner. By selecting the appropriate segments from the curvature-space description of the curve, it produces a compact representation for smooth curves.

When any structure of significant size is added to a plane curve, or existing part of the curve is deformed, this changes the curvature function in the region. If the curvature of the changes introduced is below the maximum curvature being used for the decomposition, then these changes will be manifested in the representation. Furthermore, the extent of the perturbation to the representation of the curve will be local to the region modified. As such, this representation should produce similar representations for similar curves. Likewise, the topology of the representation will be preserved under rotation and scaling of the underlying curve. In other words, as the underlying data is rotated or scaled, the part decomposition produced by this technique will retain its structure. This suggests that this representation is appropriate to curve matching and recognition since its structure is sensitive to the underlying curve data rather than the viewing conditions.

By correctly sampling curvature space and appropriately setting the discontinuity threshold, ξ, it appears that this representation can be made complete. That is, it can be made arbitrarily precise in discriminating two curves and can, in principle, be used to fully reconstruct the original curve data.

The representation has been used in a simple curve matching algorithm and is readily applicable. Typical objects that have been discriminated are shown in figure 3. Low curvature areas and large segments are useful for matching gross characteristics of objects and narrowing the search space. Small segments and those with high curvature provide additional information for discrimination between qualitatively similar objects.

5 Summary

The approach presented in this paper serves as a new robust technique for the description of curved objects. Since the method

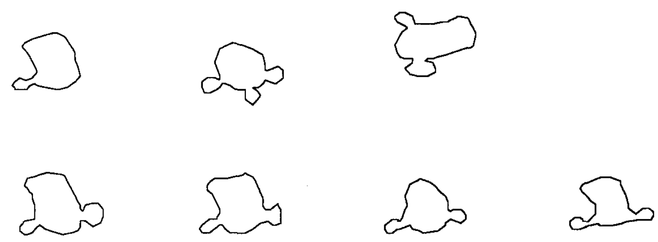

Figure 3: Several of the types of curved objects that have been discriminated using a simple matching algorithm based on the representation described here. Examples are shown of objects that differ both in "qualitative" and "quantative" properties.

5 Summary

The approach presented in this paper serves as a new robust technique for the description of curved objects. Since the method includes a technique for discontinuity detection, it can be applied to piecewise smooth data such as polygonal objects, as well as genuinely curved objects. The strength of the technique derives from the expressive power of curvature information, and the fact that curvature primitives are extracted using a robust operation largely insensitive to noise, rotation and scaling. It can also be extended in a natural manner to deal with three-dimensional surface data (?).

The examples of object recognition in this paper are based on a very simple matching technique. Since the representation implicitly creates are scale hierarchy, matching could readily exploit the coarse-to-fine structure of the curvature-tuning space. Matching could also be based on structures at specific scales (for example only-coarse or only-fine level data) in a straightforward manner. This might be useful if objects were distinguished by specific substructures rather that their overall coarse structure. More elaborated matching techniques such as this are currently the objects of ongoing research.

The discontinuity detection procedure presented above is a comparatively simple one. The simplicity of the technique is made feasible by the desire to accurately describe only regions which are comparatively free of discontinuities. As such, the precise placement of discontinuities near one another is not crucial to the results.

6 Acknowledgements

The authors would like to thank Allan Jepson and Andrew Zisserman for their insightful comments.

This work was supported by the Natural Sciences and Engineering Research Council and through the Information Technology Research Center, a Province of Ontario Center of Excellence. The second author is a Fellow of the Canadian Institute for Advanced Research.

References

Asada, H. and Brady, M. 1986. The curvature primal sketch. *Pattern Analysis and Machine Intelligence*, 8(1):2–14.

Attneave, F. 1954. Some informational aspects of visual perception. *Psychological Review*, (61):183–193.

Besl, P. J. 1988. *Surfaces in Range Image Understanding*. Springer-Verlag, New York, N.Y.

Bhanu, B. and Faugeras, O. D. 1984. Shape matching of two-dimensional objects. *PAMI*, PAMI-6(2).

Blake, A. and Zisserman, A. 1987. *Visual Reconstruction*. MIT Press, Cambridge, Mass.

Brooks, R. A. 1981. Symbolic reasoning among 3-d models and 2-d images. *AI*, 17:285–348.

Connell, J. H. and Brady, M. 1985. Generating and generalizing models of visual objects. AI memo 823, MIT AI lab.

Dudek, G., Jenkin, M., Milios, E., and Wilkes, D. 1988. Robot exploration as graph construction. RBCV-TR-88-23, Dept. of computer science, University of Toronto.

Grimson, E. 1981. *From images to surfaces: a computational study of the human earlu visual system*. MIT press, Cambridge, Mass.

Hoffman, D. D. and Richards, W. A. 1984. Parts of recognition. *Cognition*, (18):65–96.

Kehtarnavaz, N. and deFigueiredo, R. J. P. 1988. A 3-d contour segmentation scheme based on curvature and torsion. *Pattern Analysis and Machine Intelligence*, 10(5).

Leyton, M. 1987. Symmetry-curvature duality. *Computer Vision, Graphics and Image Processing*, (38):327–341.

Lowe, D. G. 1985. *Perceptual Organization and Visual Recognition*. Kluwer Academic Publishers, Boston, Mass.

Lowe, D. G. 1988. Organization of smooth image curves at multiple scales. *Proceedings of the 2nd ICCV*, pages 558–567.

Milios, E. 1988. Recovering shape deformation by an extended circular image representation. *Proceedings of the 2nd ICCV*, pages 20–29.

Nishihara, H. K. 1981. Intensity, visible-surface, and volumetric representations. *AI*, 17:265–284.

Pentland, A. 1988. Automatic extraction of deformable part models. Vision Sciences TR-104, MIT Media Lab.

Poggio, T., Torre, V., and Koch, C. 1985. Computational vision and regularization theory. *Nature*, 317:314–319.

Pomerantz, J. R., Sager, L. C., and Stoever, R. J. 1977. Perception of wholes and of their component parts: some configural superiority effects. *Journal of experimental psychology*, 3(3):422–435.

Richards, W. and Hoffman, D. D. 1984. Codon constraints on closed 2d shapes. AI memo 769, MIT AI lab.

Tauchert, T. R. and Lu, W. Y. 1987. Large deformation and postbuckling behavior of an initially deformed rod. *Journal of Non-Linear Mechanics*, 22(6):511–520.

Terzopoulos, D. 1983. The role of constraints and discontinuities in visible surface reconstruction. Technical Report, MIT AI lab.

Terzopoulos, D. 1984. *Multiresolution computation of visible-surface representations*. PhD, Dept. EE and CS, Massachusetts Institute of Technology.

Terzopoulos, D. 1986. Regularization of inverse visual problems involving discontinuities. *Pattern Analysis and Machine Intelligence*, 8(4):413–424.

Torre, V. and Poggio, T. 1984. On edge detection. AI Memo 768, MIT AI Laboratory.

Witkin, A. P. 1983. Scale-space filtering. *ICJAI 1983*, 2.

Zucker, S. W., David, C., Dobbins, A., and Iverson, L. 1988. The organization of curve detection: Coarse tangent fields and fine spline coverings. *Proceedings of the 2nd ICCV*, pages 568–577.