Shape Partitioning by Convexity

Paul L. Rosin
Department of Information Systems and Computing,
Brunel University,
Middlesex,
Middlesex UB8 3PH
UK.
email: Paul.Rosin@brunel.ac.uk

keywords: 2D shape analysis, decomposition, segmentation, visual perception

Abstract

The partitioning of 2D shapes into subparts is an important component of shape analysis. This paper defines a formulation of convexity as a criterion of good part decomposition. Its appropriateness is validated by applying it to some simple shapes as well as against showing its close correspondence with Hoffman and Singh's part saliency factors.

1 Introduction

A primary task in visual perception – for both biological and computer systems – is the analysis of shape. Despite its importance universal theories of shape have proven elusive, and much research continues to be carried out in a variety of disciplines including art [1], architecture [2], biological visual perception [3], psycholinguistics [4], qualitative reasoning [5], and computer vision [6]. One aspect of shape is the partitioning of a region into parts whose shapes are either simpler than the overall shape, or similar to an element from a predefined catalogue of primitive shapes [7]. Given the inherent difficulties of vision, particularly those related to the variability in the appearance of an object due to different viewpoints or articulation of parts, such a decomposition helps simplify the problem of perception. For instance, in many cases there will be a one-to-one correspondence between observable region parts and functional components of the viewed object.

Naturally image understanding involves a multitude of factors such as colour, texture, shading, and motion, as well as non-visual information such as contextual cues, prior expectations, etc. This paper is restricted to shape analysis, the importance and power of which was demonstrated by Biederman and Ju in experiments where both colour photographs and line drawings (i.e. only shape information was present) were recognised with comparable facility and speed [8].

Shape can be analysed by considering either a region's interior (i.e. the enclosed area) or exterior (i.e. its boundary) [9]. Since one can be constructed from the other interior and exterior representations of the region are equivalent, and may make no difference to the analysis (e.g. identical shape descriptors are calculated using either area or line moments). At other times explicitly representing the interior or exterior makes certain information easier to elicit. Some of the difficulty in determining a good method for decomposing shapes into parts is that the analysis needs to use both explicit boundary information (e.g. local concavities) and interior information (more global shape descriptions).

A popular shape representation since the 1960's, at least for extended ribbon-like objects, is the skeleton or axis. This can be defined as the points of local symmetry of the region. Labelling each axis point with the distance to the boundary, and connecting adjacent points, gives a curve in 3D space from which Sanniti di Baja and Thiel [10] determined part boundaries. First breaks were made at junctions in the skeleton so that it was decomposed into branches. Second, a polygonal approximation of each branch was made, so that segments represented parts with linear changes in width and orientation. Third, the result was tidied up, removing short and superfluous sections and merging others. Finally, from each remaining axis segment a subpart of the original shape could be reconstructed. A variation of this approach that makes the combination of interior and exterior information more explicit is presented by Abe *et al.* [11] who segment the axes based on dominant points detected along the boundary.

Based on psychophysical and ecological considerations Siddiqi and Kimia [12] described a partitioning scheme involving two types of parts: necks and limbs. Necks were determined by diameters of locally minimal inscribed circles in the region while limbs were lines through pairs of negative curvature minima having co-circular boundary tangents (i.e. they join smoothly). Competing candidate partitionings were resolved by computing salience values for parts. For necks these were defined as the product of the curvature disparity across the neck and the length of the part boundary line. A limb's salience was a function of the total curvature curvature across the limb and the extent of limb across the part line.

Recently, Singh *et al.* [13] criticised Siddiqi and Kimia's scheme, noting that the definitions for limbs and necks were too restrictive, and failed for a large class of shapes. They proposed an alternative method to partition shapes: the short-cut rule. Their definition of a part line, which they term a cut, is:

1. a straight line
2. crossing an axis
3. joining two boundary points
4. at least one of which has negative curvature;
5. if there are several possible competing cuts the shortest one is selected.

In addition, they use the minima rule [14] which states that negative minima of curvature provide points for cutting the shape.

This paper in turn points out some limitations of Singh *et al.*'s scheme, and proposes an alternative rule for partitioning shapes. All methods for segmenting shapes to date have had several drawbacks – either computational or perceptual – and the new approach is not perfect either.[1] However, it does have the advantages of appearing to provide perceptually reasonable results without requiring perfect line data while its guiding principles are straightforward and uncluttered.

2 Limitations of the Short-cut Rule

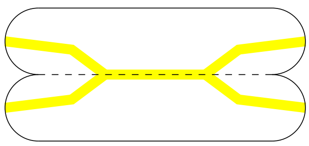

Figure 1: The cut (drawn dashed) shown will not be selected as it does not cross the axes (drawn grey)

A major weakness of Singh *et al.*'s short-cut rule is that it incorporates only very limited global shape information. Initial cues for cut locations (negative curvature minima) are exclusively determined from the boundary information. More global shape information is introduced in two ways.

1. The length of the cut, involving only minimal shape information.
2. Restricting of the cut to cross an axis. Using the object axes introduces several difficulties:
   (a) In the computer vision literature there have been many definitions of axes over the years [15, 16], and since they will often produce different axes from the same shape this variability will affect the generation of cuts. Singh *et al.* state that they use Brady and Asada's [17] definition of axes, but this appears to be an arbitrary choice since no justification is given.

[1] It may be that shape segmentation is as difficult as edge detection, for which there have been hundreds of algorithms proposed during the last 35 years without an entirely satisfactory one being found.

(b) Robust computation of axes is difficult since from their very definition most axes are extremely sensitive to noise. Small perturbations of the boundary can radically alter the axis. In practice the results often need to be extensively post-processed to eliminate spurious axes. Thus, using region axes reduces the practical effectiveness of Singh *et al.*'s scheme although its inclusion was necessary to avoid short but undesirable cuts.

(c) Perceptually valid cuts need not properly cross an axis; figure 1 shows an example of a cut that instead crosses the junctions of the axis branches and for most of its length coincides with the central branch of the axis.

Figure 2: Two part objects, both with identical width cuts although shape (a) has more salient parts than object (b)

The second major weakness is that the sole determinant of salience is the length of the cut. Obviously (as demonstrated in figure 2) there is more to salience than cut length. In fact, in another paper Hoffman and Singh [18] isolate three factors affecting part salience: relative area, amount of protrusion, and normalised curvature across the part boundary, but they do not integrate these into the shape partitioning scheme.

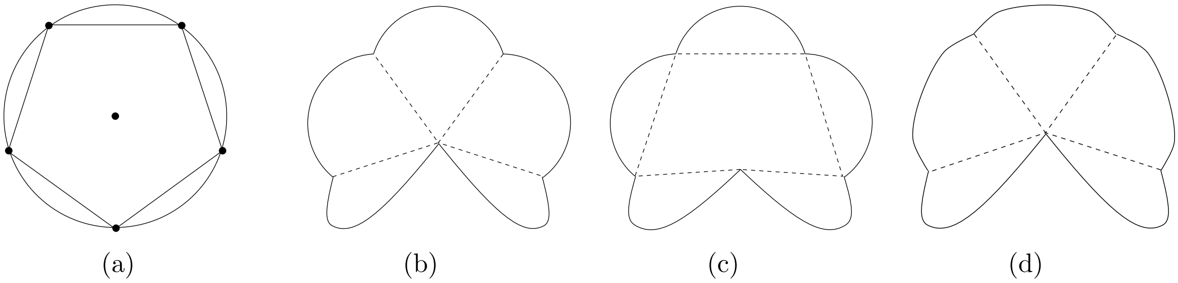

Figure 3: Constructing a shape prone to unstable partitioning; (a) circle with inscribed pentagon, (b) radial cuts, (c) cuts around core, (d) a version of (b) with less salient parts

Although in the last example the cut decomposed the shape correctly we will now construct the more problematic shape shown in figure 3. A pentagon inscribed in a circle of radius r has sides of length $\sqrt{\frac{1}{2}(5-\sqrt{5})}r \approx 1.18r$. Taking all the vertices of the pentagon except for one which is replaced by the centre of the circle, the vertices are connected by arcs. The central vertex is the cusp closest to the remaining cusps, and so Singh *et al.*'s cuts slice the region radially (figure 3b). If the central vertex is pushed slightly outwards, barely changing the shape's appearance, then the adjacent cusps become mutually closest to each other, changing the partioning to cutting protrusions off a central core (figure 3c). The unsatisfactory nature of the pie-slicing in figure 3b can be emphasised by reducing the bulging of the arcs which reduces their salience but leaves Singh *et al.*'s cuts unaltered (figure 3d).

Returning to Singh *et al.*'s list of conditions for a cut we find further difficulties. Restricting part boundaries to straight lines is obviously undesirable. A straight cut for the shape in figure 4a is not as appropriate as the curved one, while the straight cut in figure 4b is unacceptable since it is not contained within the shape.

Singh *et al.* only require one end of a cut to lie on a portion of boundary with negative curvature (i.e. an indent). This avoids unintuitive partitions such as figure 5a in which the only two cusps have been joined to each other. The shortest cut from each cusp as they propose does produce a more sensible partitioning (figure 5b), but the longer cuts in figure 5c look better yet. As Singh *et al.* point out in their concluding remarks their short-cut rule and minima rule do not incorporate

(a) (b)

Figure 4: The requirement for curved part boundaries; (a) a curved cut is more appropriate than a straight cut, (b) only the curved cut is valid.

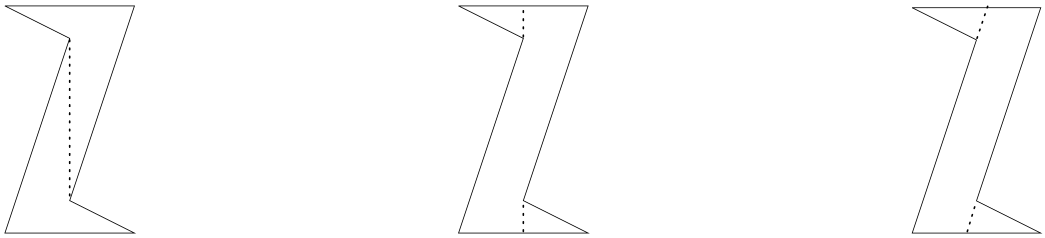

(a) (b) (c)

Figure 5: Should cuts be made to (a) the closest negative curvature minimum (b) the closest boundary point irrespective of its curvature, or (c) some other point?

all the important perceptual factors involved in shape perception (e.g. the Gestalt principles of symmetry and good continuation).

(a) (b)

Figure 6: Should cuts be made to (a) the closest negative curvature minimum (b) the closest boundary point irrespective of its curvature?

Further examples in which the shortest cuts are inferior are easy to find. For instance, in figure 6 the single cut joining the cusps looks preferable to the shortest cut.

In figure 7a the shortest cuts make little sense, even though they do cross the region's axes. Actually, the true shortest cuts would be angled such as to be almost vertical, partitioning infinitesimal slivers off the region. On the other hand, if we consider cuts between pairs of cusps instead, the short-cut rule still leads to difficulties. The problem is that unlike our previous examples, the members of pairs of cusps forming cuts are not both the closest cusp to each other, as shown in figure 7b. Here the cusps closest to the ends are closest to the opposite central cusp leading to over-segmentation. Meanwhile the central cusps are closest to each other. The two triangular regions formed by the cuts have no perceptual relevance. Two alternative, more appropriate partitionings would be to keep either the two outer cuts or the inner one.

3 Convex Partitioning

Various formulations and approximations of convexity have been used as criteria by previous authors for object decomposition. For instance, some early work by Pavlidis [19] proposed segmenting polygons into convex subsets. However, the approach was computationally expensive, and a simpler implementation restricted to a decomposition into horizontal and vertical rectangles was shown.

Shapiro and Haralick [20] showed that dense clusters of internal line segments form at convex parts of regions. Thus they applied a clustering algorithm to identify local areas of high compactness which are then merged to form larger subparts. Unfortunately this process required specifying

(a) (b)

Figure 7: Problems with the short-cut rule; (a) selecting closest boundary points, (b) selecting closest cusps

many parameters, namely thresholds for cluster overlap, compactness, association, and size.

Held and Abe [21] defined an approximate measure of convexity based on the fraction of the region boundary that coincided with the region's convex hull. The initial stages of their algorithm was based on boundary dominant points and the skeleton, similar to that by Abe *et al.* [11]. A structuring element was applied to the segmented branches of the axes, and these were then merged dependent on their convexity value. Again various parameters were required to control the axes segmentation and merging stages.

Recently Latecki and Lakämper [22] avoided the many of the difficulties of the above approaches, using their so called "discrete evolution by digital linearization". Boundary points are iteratively deleted (or equivalently adjacent line pairs are merged) until the resulting shape is convex. At each iteration the line pair merge with the lowest cost (which is a function of its length and curvature) is selected. The iterations produce a hierarchy of maximally convex boundary arcs, each of which defines a cut by the straight line joining its endpoints. The advantage of the scheme is that it only requires one parameter to threshold the cuts according to their saliency. However, the disadvantages are threefold. First, the strict ordering of the line merging may restrict the formation of some salient cuts. Second, only boundary information is used even though region information is generally considered important. Third, Latecki and Lakämper state that for continuous data cuts would correspond to points of inflection. However, in practice they appear to be restricted to lie on indentations, i.e. near maxima of negative curvature. As we have previously discussed, this is over-restrictive, and causes poor results. For example, the L shape of the kangaroo's foot (in Latecki and Lakämper's figure 6) is not properly partitioned since it needs the cut to terminate at the maximum of positive curvature. Other examples of inappropriate cuts are shown in Latecki and Lakämper's figure 5, shapes 2, 5, and 7.

This paper proposes segmenting regions into roughly convex parts in a more direct manner that the above approaches, and avoids many of their complications. Only two components are required:

1. a measure of convexity, and
2. an optimisation scheme.

Convexity of a partitioned region is calculated as the weighted sum of the convexities of its parts

$$C_P = \frac{1}{A_R} \sum_{i=1}^{n} A_i C_i$$

where the region R is decomposed into n parts which individually have area A_i and convexity C_i, and the total area $A_R = \sum_{i=1}^{n} A_i$. A region's (or subpart's) convexity is calculated as the ratio of the area of the region to the area of its convex hull. Thus the calculation of convexity becomes

$$C_P = \sum_{i=1}^{n} \frac{A_i}{A_R} \frac{A_i}{H_i}$$

where H_i is the area of the convex hull of part i. The individual and combined convexity measures return a score of one for a perfect convex region and approach zero for shapes with extremely deep concavities. Given a specification of the number of desired cuts the aim of the optimisation stage is to find the best set of cuts to maximise C_P.

The advantage of this scheme is that it is extremely simple to define, not requiring many parameters such as Shapiro and Haralick's clustering method or Sanniti di Baja and Thiel's axis

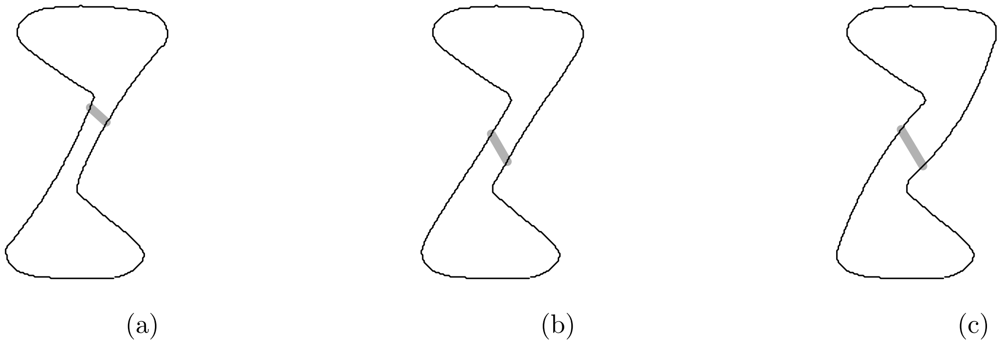

Figure 8: Convexity-based partitioning can occur at negative, zero, and positive curvature boundary points.

pruning/merging method. Moreover, convexity combines both interior and exterior aspects of shape, so that the salience of a segmentation is better reflected by convexity than by cut length.

As shown in figure 8, convexity-based partitioning does not restrict the cuts to terminate at concavities. However, as examples later in the paper demonstrate, partitioning does tends to occur at concavities of the shape, which is generally considered to be appropriate. However, in contrast to those schemes that explicitly identify the concavities by analysing the curvature of the boundary, these locations naturally fall out of the convexity scheme. This is preferable since curvature estimation is sensitive to noise whereas the convex hull is much more robust.

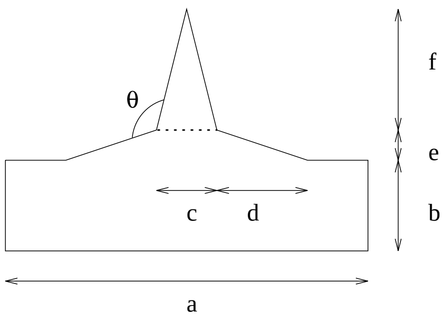

Figure 9: A simple parameterised shape for measuring salience

In fact, it can be seen that convexity is closely related to Hoffman and Singh's part salience factors. These consist of the size of the part relative to the whole object, the degree to which the part protrudes, and the strength of its boundaries (measurable as the turning angle). Using psychophysical experiments they showed that the factors exhibit high correlation with human vision behaviour. As a simple demonstration of the connection between convexity and Hoffman and Singh's part salience factors we examine the shape in figure 9 containing a block with one protruding part. The convexity of the total region is

$$C_1 = \frac{2(ab + ce + de) + cf}{a(2b + e + f)}.$$

After the cut the convexity of the resulting part decomposition becomes

$$C_2 = \frac{\frac{(ab+e(c+d))^2}{a(b+e)-de} + \frac{cf}{2}}{ab + de + c(e + \frac{f}{2})}$$

and so the improvement gained by partitioning is $S = C_2 - C_1$, which we will take as a measure of salience for the comparison. Hoffman and Singh's measure for part size is calculated as the relative area of the part

$$\frac{cf}{2(ab + ce + de) + cf},$$

the degree of protrusion is the ratio of the perimeter of the part (excluding the base) to the length of the base (i.e. the cut)

$$\frac{\sqrt{c^2 + 4f^2}}{c},$$

and the turning angle is

$$\pi - \theta = tan^{-1}\frac{d}{e} + tan^{-1}\frac{2f}{c} - \frac{\pi}{2}.$$

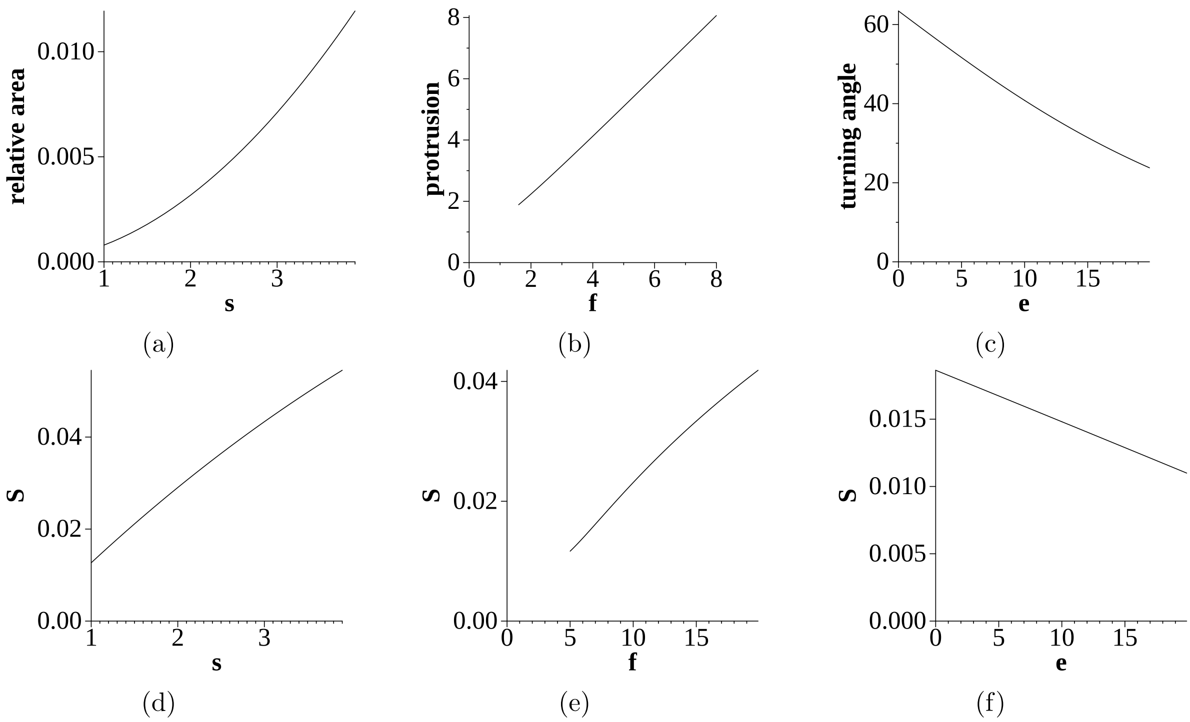

Figure 10: Saliency measures calculated over variations of the shape in figure 9; (a) & (d) increasing relative area (b) & (e) increasing protrusion; (b) & (f) decreasing turning angle (a-c) Hoffman and Singh's part saliency factors (d-f) corresponding convexity saliency factor

Figure 10 shows the effects that modifying the shape has on the saliency factors. The parameters are first set to $a = 50$, $b = 50$, $c = 2$, $d = 10$, $e = 1$, $f = 2$. Changing even one parameter can affect all the saliency factors; for instance, increasing c decreases the turning angle and the subpart's degree of protrusion and increases its relative area. Therefore to limit the changes to one factor at a time we modify the parameters as follows:

- increasing relative area – c and f are both increased by scale factor s
- increasing protrusion – this is obtained by increasing f; the turning angle is fixed by setting

  $$e' = d \cot\left(tan^{-1}\frac{d}{e} + tan^{-1}\frac{2f}{c} - tan^{-1}\frac{2f'}{c}\right)$$

  where e' and f' are the new values of e and f; the relative area is then fixed by setting $a' = \sqrt{s}a$ and $b' = \sqrt{s}b$ where

  $$s = \frac{f}{f'} + \frac{f(c+d)(e-e')}{abf'}$$

- decreasing turning angle – e is increased; to maintain the same relative area b is modified to

  $$b' = b + \frac{(c+d)(e-e')}{a}$$

  where b' and e' are the new values of b and e; in addition, as increasing e causes the part to be pushed out relative to the horizontal surface adjacent to d, that section is removed by setting $d = \frac{a-c}{2}$

A good correspondence can be seen between Hoffman and Singh's part saliency factors and the convexity saliency factor. Nonetheless this does not necessarily always hold. Since Hoffman and Singh's factors do not uniquely determine the shape then even with such a simple shape there are alternative normalisations to those used in this paper which may behave differently.

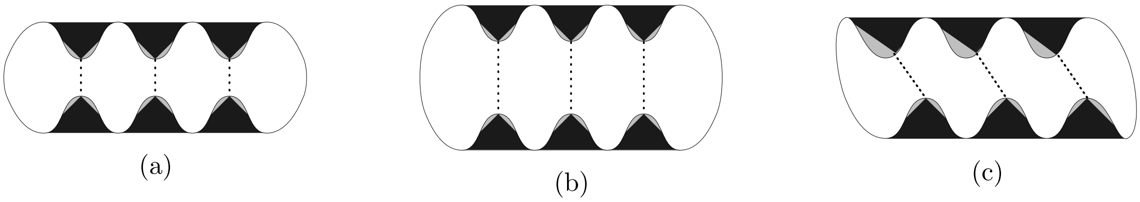

Figure 11: Effects of a cut on the convexity of (a) parts, (b) protrusions, (c) bends

Siddiqi and Kimia identify the major components of shape as a continuum they call the "shape triangle" between three categories: parts, bends, and protrusions. It is interesting to investigate the effect of the convexity measure on these shapes; three examples are shown in figure 11. The three optimum cuts that maximise convexity are drawn dotted; the gray areas show the convex hull outside the shape, while the black regions are the portions of the convex hull that are eliminated by the cut. In all cases similar cuts are formed – connecting two opposite concavities. Progressing from shapes with parts to protrusions makes no difference to the cut except that the relative area of the region to its convex hull increases so that the improvement in convexity caused by the cut decreases. A half cycle shift of the bottom of the part shape produces a bend, i.e. an undulating ribbon of uniform width. If parts and bends are analysed by their axes different properties would be needed for segmenting the two cases, namely: width and orientation respectively. Using convexity the differences do not need to be explicitly identified. Appropriate cuts are formed although the improvement in convexity for bends is less than for parts as seen by the smaller black areas.

4 Examples

The application of the convexity rule to some examples shapes drawn from previous papers in the field and other sources is shown in figures 12 and 13. The shapes have been roughly grouped into order according to the following characteristics: one concavity, one protrusion, two major indentations leading to two parts, three parts, four parts, and "L" shapes. Using a single cut it can be seen that the majority of the decompositions are sensible (figure 12). This includes shapes which are problematic for the algorithms of Latecki and Lakämper (figure 12b and 12t), Siddiqi and Kimia (figure 12c and 12k), and Singh *et al.* (figure 12m). In cases where a single cut is inappropriate the result providing maximal convexity is sometimes appropriate (figure 12p) while at other times less appropriate (figure 12q).

Further examples showing the decomposition resulting from pairs of cuts are shown in figure 13. Where there are three natural parts to the object then these have been found. In other situations the decomposition is also plausible. For instance, the parts in figure 13c have qualitatively different characteristics: elongated, tapering, and circular. In figures 13h–13j the third part that can now be detected using the addition cut is the connector between the two primary object parts.

5 Discussion

We have described a part decomposition scheme based on maximising convexity. Its advantages are first, that it is simple, and does not require many stages of processing with attendant parameters that require selection. Second, the convexity criterion appears perceptually valid, as tested on some simple shapes as well as against Hoffman and Singh's part saliency factors. Third, generalising the scheme to higher dimensions is straightforward. However, there remain some limitations with the proposed approach; these are listed with some possible solutions.

• *Efficiency.* A simple exhaustive search to find the optimal cuts has a time complexity of $O(n^{2p})$ for n boundary points and p cuts, which rapidly becomes impractical. A speed-up was obtained for the results shown in section 4 by using an exhaustive search at two scales. First a version subsampled by a factor between 5 and 10 was processed. This was subsequently refined on

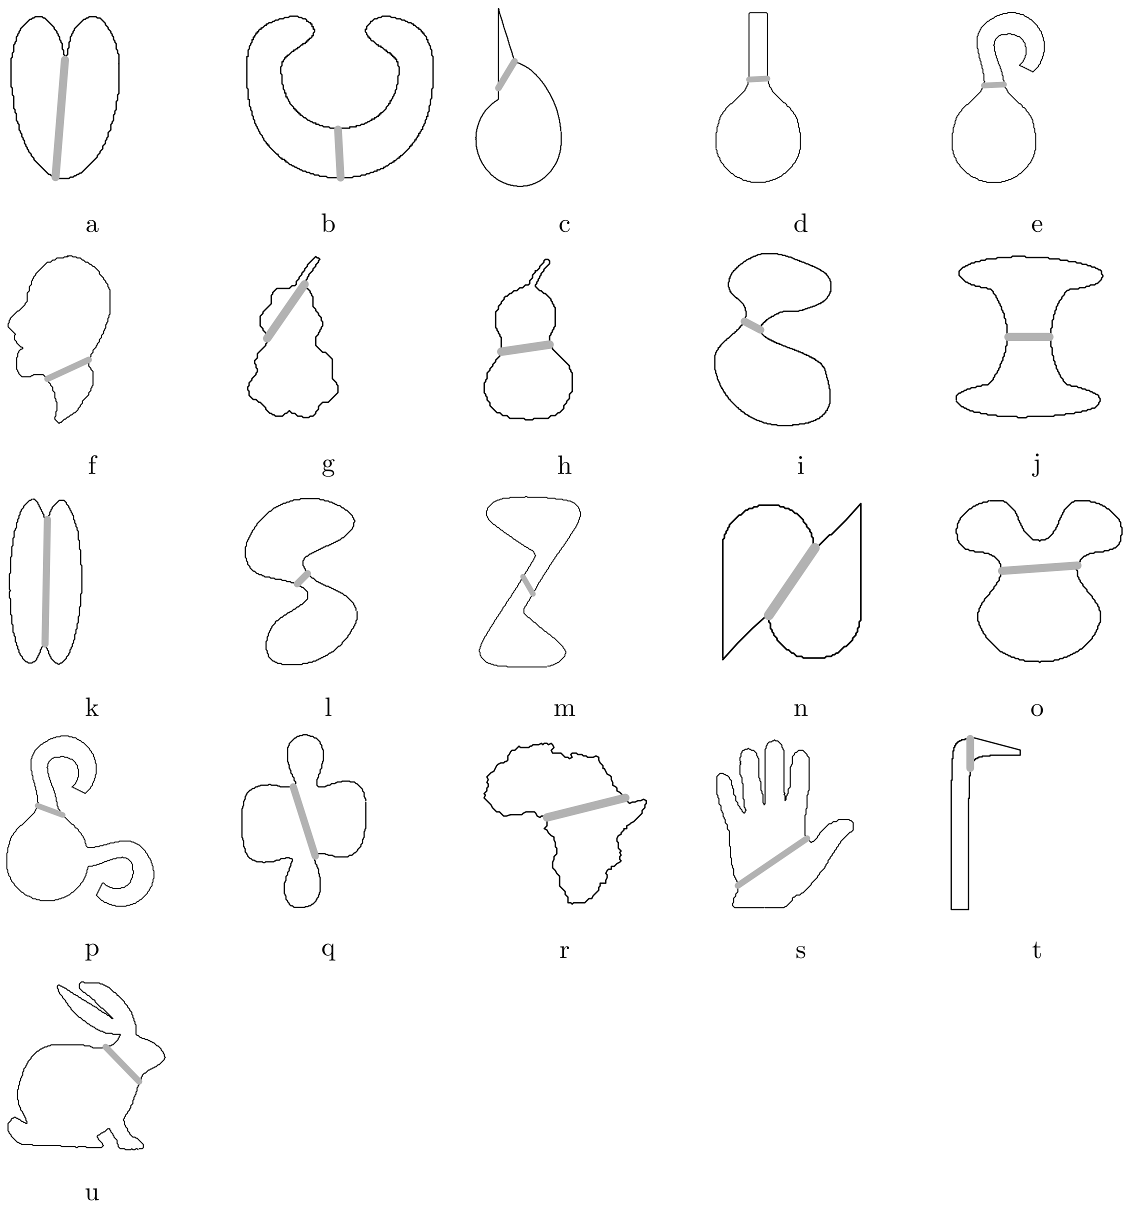

Figure 12: Example shapes with best single cut

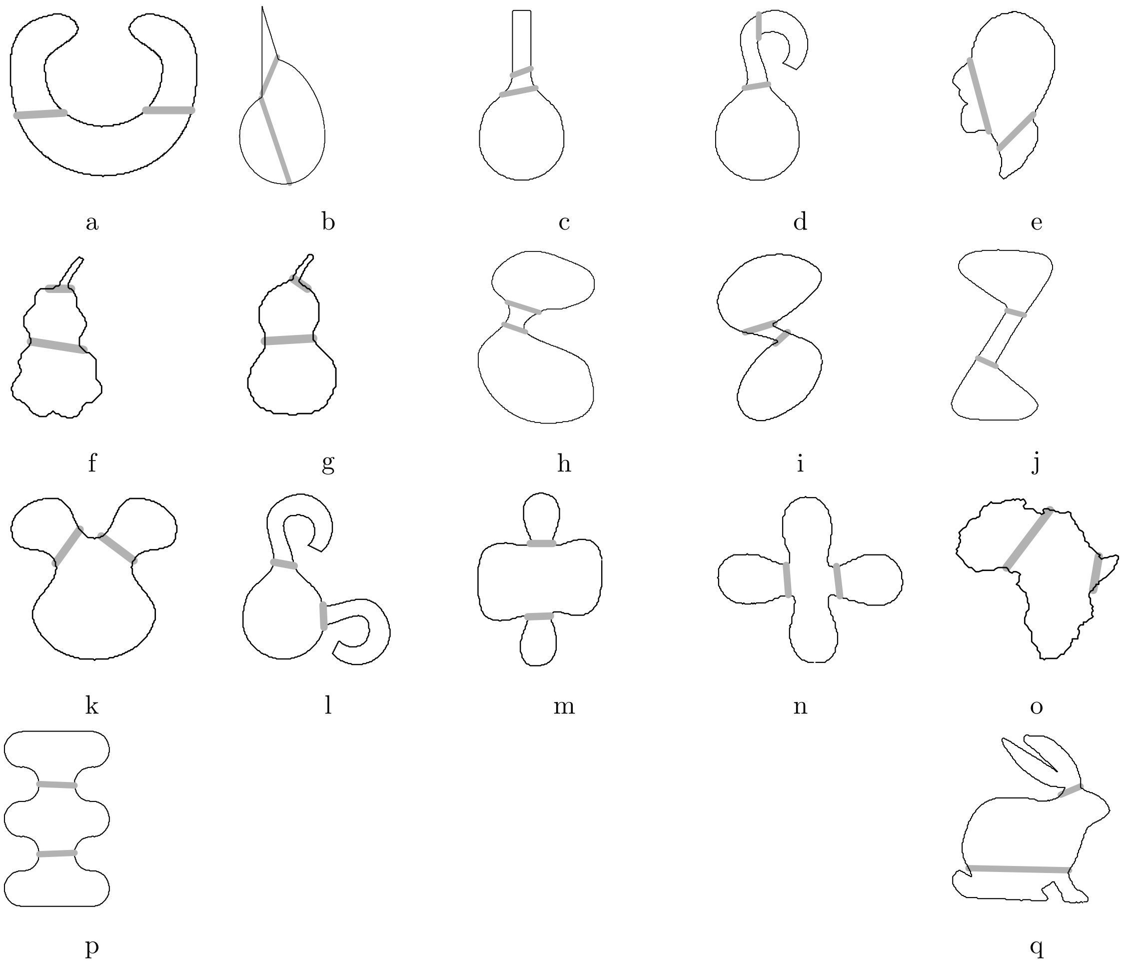

Figure 13: Example shapes with best pair of cuts

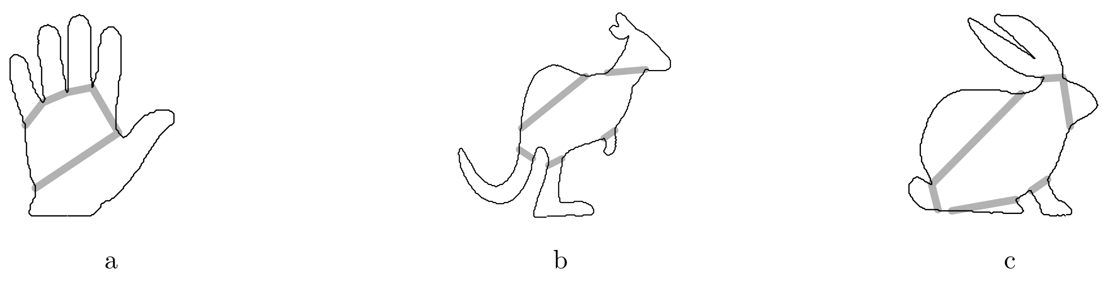

Figure 14: Example shapes with five cuts

the full resolution version. However, the time complexity remains exponential, and so for a larger number of cuts a faster algorithm is necessary. As one possible solution we have experimented with applying a simple random optimisation approach. Dominant points are found on the curve using a standard algorithm (Ramer's polygonisation). These are used as seeds for initial endpoints of cuts. The threshold used for detecting the seed points is not crucial; there is a tradeoff between subsequent efficiency and accuracy/correctness. All valid cuts formed by pairs of cuts are determined. The constraint is that the line formed must lie within the shape. To determine a good set of p cuts many sets of p randomly selected cuts are tested, refined by shifting their endpoints, and the best (i.e. producing the most convex partitioning) is retained. The time complexity is reduced to a component for the initialisation which is $O(n \log n)$ for Ramer's algorithm (although many alternative linear time polygonisation algorithms are available), the computation of the convex hull which for an n-gon is $O(n)$ [23], the calculation of the areas of the regions which is $O(n)$, and the test that the cut is contained in the polygon which is also $O(n)$. Thus, by replacing Ramer's algorithm an overall linear algorithm is possible. Figure 14 shows some reasonably good results obtained for partitioning with five cuts. However, in order to produce acceptable results this simple scheme requires many iterations, involving large amounts of processing time. A better approach would be to use a genetic algorithm to direct the optimisation since this would enable partial solutions to be reused unlike the current scheme in which each random set is generated and tested in isolation from all the others.

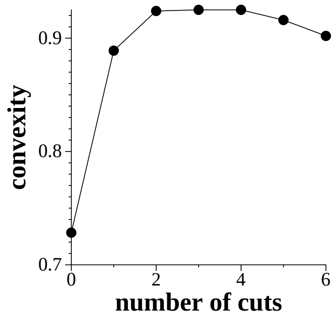

Figure 15: Convexity as a function of the number of cuts can indicate the appropriate number of cuts

- *Number of cuts.* A means is required for specifying the number of parts to decompose the region into. This is the same problem present with the segmentation of curves into straight lines, and the same solutions can be applied. In some situations the expected number of parts is known in advance and therefore can be directly applied. Otherwise a threshold on the quality of the output can be applied, i.e. increasing numbers of cuts are tested until the resulting convexity is acceptable. Alternatively attempts can be made to automate threshold selection by looking for a discontinuity in the convexity versus number of parts graph, using Rosin's unimodal threshold selection method [24]. A flattening of the graph indicates that additional cuts are not significantly improving the quality of the output, and are therefore not cost-effective and undesirable. This is demonstrated by the plot in figure 15 which is generated from the shape in figure 12e (also seen in figure 13d) which correctly indicates that two cuts are most appropriate. The drop-off in the graph for greater numbers of cuts is due to the increased difficulty in performing the optimisation as the number of sets of valid cuts increases exponentially.
- *Straight versus curved cuts.* Like most previous algorithms for part decomposition for algorithmic simplicity the cut is restricted to a straight line even though we showed that this is not always appropriate. In fact, even if curved cuts were allowed they would not necessarily be chosen by our scheme. For example, figure 16a shows a simplified version of figure 4 containing a semi-circular disk which is partitioned along its length by two sections of a semi-circular cut connected by a straight cut of length $2s$. Without loss of generality we can set $b = 1$, which simplifies the expression for the convexity of the partitioning to

$$\frac{1}{c^2\pi}\left(\frac{(p+\pi-a^2\pi)^2}{p+\pi}+\frac{(p+\pi-c^2\pi)^2}{c^2\pi}\right)$$

Figure 16: Effect on convexity of straight or curved cuts

where

$$p = 2s\sqrt{1-s^2} - 2\tan^{-1}\left[\frac{s}{\sqrt{1-s^2}}\right]$$

and $0 \leq s \leq \sqrt{1-a^2}$ determines the extent of the straight cut. For any values of $0 < a < 1$ and $c > 1$ the graph of convexity versus s has the shape shown in figure 16b. As s is increased from zero (i.e. only a curved cut is used), increasing the extent of the straight cut, the convexity increases. That is, the convexity measure favours a straight rather than a curved cut.

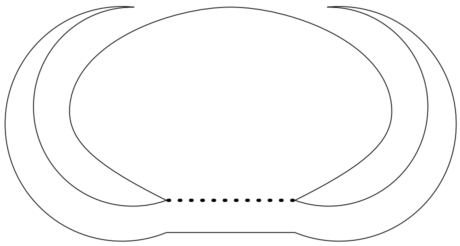

Figure 17: Convexity based partitioning can produce poor results

- *Inappropriateness of convexity.* Although the convexity criterion has been shown to work well on many examples, it is not always appropriate. For instance, the crab-like shape in figure 17 provides a counter-example. The most likely intuitive cut would be to separate the body and legs as indicated by the dotted line. However, the complete shape (before partitioning) is already fairly convex despite the gaps between the legs and body, and the convexity can be further increased to approach unity by pulling the legs in towards the body. If the shape is segmented at the dotted line the outer leg-pair region contains a large concavity, lowering the overall convexity score. Thus, in this instance the convexity criterion favours no segmentation over the more appropriate partitioning in figure 17.

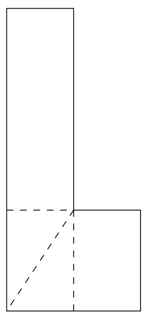

Figure 18: Convexity provides no indication of which is the best of the three possible cuts

- *Saliency.* Convexity does not always provide complete saliency information. For instance, in figure 18 all three cuts indicated would produce decompositions into perfectly convex parts with identical scores although Singh *et al.* demonstrated a perceptual preference for the shortest cut. One way to overcome this and other deficiencies of the convexity measure would be to augment it with other saliency factors such as length of cut, size of segmented regions, goodness of boundary continuation, etc. The difficulty would be to learn how to combine them appropriately. Otherwise the same difficulties arise as with snakes, which are often formulated to minimise some weighted sum of error factors; although few guidelines are given for setting the weights, their values are often critical to the final result.

References

[1] J. Willats, *Art and Representation: New Principles in the Analysis of Pictures*, Princeton University Press, 1997.

[2] M.Y. Cha and J.S. Gero, "Shape pattern recognition using a computable pattern representation", in *Artificial Intelligence in Design '98*, J.S. Gero and F. Sudweeks, Eds. 1998, pp. 169–187, Kluwer Academic Publishers.

[3] S. Pinker, Ed., *Visual cognition: An introduction*, MIT Press, 1985.

[4] J.P. Denny, "The 'extendedness' variable in classifier semantics: universal features and cultural variation", in *Ethnolinguistics: Boas, Sapir, and Whorf revisited*, M. Mathiot, Ed. 1979, pp. 97–119, Mouton.

[5] A.G. Cohn, "Qualitative spatial representation and reasoning techniques", in *Proceedings KI-97*, G. Brewka, C. Habel, and B. Nebel, Eds. 1997, vol. LNAI 1303, pp. 1–30, Springer Verlag.

[6] S. Loncaric, "A survey of shape analysis techniques", *Pattern Recognition*, vol. 31, no. 8, pp. 983–1001, 1998.

[7] I. Biederman, "Recognition by components: A theory of human image understanding", *Psychological Review*, vol. 94, no. 2, pp. 115–147, 1987.

[8] I. Biederman and G. Ju., "Surface versus edge-based determinants of visual recognition", *Cognitive Psychology*, vol. 20, pp. 38–64, 1988.

[9] T. Pavlidis, "A review of algorithms for shape analysis", *Computer Graphics and Image Processing*, vol. 7, no. 2, pp. 243–258, 1978.

[10] G. Sanniti di Baja and E. Thiel, "(3,4)-weighted skeleton decomposition for pattern representation and description", *Pattern Recognition*, vol. 27, no. 8, pp. 1039–1049, 1994.

[11] K. Abe, C. Arcelli, T. Hisajima, and T. Ibaraki, "Parts of planar shapes", *Pattern Recognition*, vol. 29, no. 10, pp. 1703–1711, 1996.

[12] K. Siddiqi and B.B. Kimia, "Parts of visual form: Computational aspects", *IEEE Trans. PAMI*, vol. 17, no. 3, pp. 239–251, 1995.

[13] M. Singh, G.D. Seyranian, and D.D. Hoffman, "Parsing silhouettes: the short-cut rule", *Perception and Psychophysics*.

[14] D. Hoffman and W. Richards, "Parts of recognition", *Cognition*, vol. 18, pp. 65–96, 1984.

[15] M. Leyton, "Symmetry-curvature duality", *Computer Vision, Graphics and Image Processing*, vol. 38, no. 3, pp. 327–341, 1987.

[16] A. Rosenfeld, "Axial representations of shape", *Computer Vision, Graphics and Image Processing*, vol. 33, pp. 156–173, 1986.

[17] M. Brady and H. Asada, "Smoothed local symmetries and their implementation", *International J. Robotics Research*, vol. 3, no. 3, pp. 36–61, 1984.

[18] D. Hoffman and M. Singh, "Salience of visual parts", *Cognition*, vol. 63, pp. 29–78, 1997.

[19] T. Pavlidis, "Analysis of set patterns", *Pattern Recognition*, vol. 1, pp. 165–178, 1968.

[20] L.G. Shapiro and R.M. Haralick, "Decomposition of two-dimensional shapes by graph-theoretic clustering", *IEEE Trans. PAMI*, vol. 1, no. 1, pp. 10–20, 1979.

[21] A. Held and K. Abe, "On the decomposition of binary shapes into meaningful parts", *Pattern Recognition*, vol. 27, no. 5, pp. 637–647, 1994.

[22] L.J. Latecki and R Lakämper, "Convexity rule for shape decomposition based on discrete contour evolution", *Computer Vision and Image Understanding*, vol. 73, no. 3, pp. 441–454, 1999.

[23] D. McCallum and D. Avis, "A linear algorithm for finding the convex hull of a simple polygon", *Inform. Process. Lett.*, vol. 9, pp. 201–206, 1979.

[24] P.L. Rosin, "Unimodal thresholding", in *Proc. Scand. Conf. Image Anal.*, 1999, pp. 585–592.