Shape Partitioning by Convexity

Paul L. Rosin

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

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, architecture, biological visual perception, psycholinguistics, qualitative reasoning, and computer vision. 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 [2]. 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 [3].

Shape can be analysed by considering either a region's interior (i.e. the enclosed area) or exterior (i.e. its boundary) [10]. 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 [13] determined part boundaries by a process of segmentation, pruning, and merging. A variation of this approach that makes the combination of interior and exterior information more explicit is presented by Abe *et al.* [1] who segment the axes based on dominant points detected along the boundary.

Based on psychophysical and ecological considerations Siddiqi and Kimia [15] 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.* [16] 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 [6] 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

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]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.

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

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 [9, 11], 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 [4] definition of axes, but this appears to be an arbitrary choice since no justification is given.

   (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.

The second major weakness is that the sole determinant of salience is the length of the cut. Obviously there is more to salience than cut length. In fact, in another paper Hoffman and Singh [7] 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.


(a)   (b)   (c)

Figure 2: 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?

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 2a 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 2b), but the longer cuts in figure 2c look better yet. As Singh *et al.*

point out in their concluding remarks their short-cut rule and minima rule do not incorporate all the important perceptual factors involved in shape perception (e.g. the Gestalt principles of symmetry and good continuation).

(a)   (b)

Figure 3: 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 3 the single cut joining the cusps looks preferable to the shortest cut.

(a)   (b)

Figure 4: Problems with the short-cut rule; (a) selecting closest boundary points, (b) selecting closest cusps

In figure 4a 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 4b. Here the cusps closest to the ends are closest to the opposite central cusp leading to oversegmentation. 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 [10] 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 [14] 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 many parameters, namely thresholds for cluster overlap, compactness, association, and size.

Held and Abe [5] 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.* [1]. 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 [8] 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 terminate at 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

$$C_P = \frac{1}{A_R} \sum_{i=1}^{n} \frac{A_i^2}{H_i}$$

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 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.


Figure 5: 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 5 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

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

(a)   (b)   (c)   (d)   (e)   (f)

Figure 6: Saliency measures calculated over variations of the shape in figure 5; (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 6 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 7: Example shapes with best single cut

4 Examples

Figure 8: Example shapes with best pair of cuts

The application of the convexity rule to some examples shapes drawn from previous papers in the field and other sources is shown in figures 7 and 8. 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 7). This includes shapes which are problematic for the algorithms of Latecki and Lakämper (figure 7a and 7k), Siddiqi and Kimia (figure 7b and 7f), and Singh *et al.* (figure 7h). In cases where a single cut is inappropriate the result providing maximal convexity is sometimes appropriate (figure 7i) while at other times less appropriate (figure 7j).

Further examples showing the decomposition resulting from pairs of cuts are shown in figure 8. 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 8b have qualitatively different characteristics: elongated, tapering, and circular. In figures 8e–8f 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. It's 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. However, there remain some limitations with the proposed approach; these are listed with some possible solutions.

- *Efficiency.* The results in section 4 were obtained using an exhaustive search at two scales. First a version subsampled by a factor between 5 and 10 was pro-

Figure 9: Example shapes with five cuts

cessed. This was subsequently refined on the full resolution version. However, for larger numbers of cuts this becomes computationally excessive. We have experimented using 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 n cuts many sets of n randomly selected cuts are tested, refined by shifting their endpoints, and the best (i.e. producing the most convex partitioning) is retained. Figure 9 shows some results partitioned with five cuts. Although good results are achievable this simple scheme still requires fairly 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.

• *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. One approach is to look for a discontinuity in the convexity versus number of parts graph. 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 [12].

• *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.


Figure 10: 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 10 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] K. Abe, C. Arcelli, T. Hisajima, and T. Ibaraki. Parts of planar shapes. *Pattern Recognition*, 29(10):1703–1711, 1996.

[2] I. Biederman. Recognition by components: A theory of human image understanding. *Psychological Review*, 94(2):115–147, 1987.

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

[4] M. Brady and H. Asada. Smoothed local symmetries and their implementation. *International J. Robotics Research*, 3(3):36–61, 1984.

[5] A. Held and K. Abe. On the decomposition of binary shapes into meaningful parts. *Pattern Recognition*, 27(5):637–647, 1994.

[6] D. Hoffman and W. Richards. Parts of recognition. *Cognition*, 18:65–96, 1984.

[7] D. Hoffman and M. Singh. Salience of visual parts. *Cognition*, 63:29–78, 1997.

[8] L.J. Latecki and R Lakämper. Convexity rule for shape decomposition based on discrete contour evolution. *Computer Vision and Image Understanding*, 73(3):441–454, 1999.

[9] M. Leyton. Symmetry-curvature duality. *Computer Vision, Graphics and Image Processing*, 38(3):327–341, 1987.

[10] T. Pavlidis. A review of algorithms for shape analysis. *Computer Graphics and Image Processing*, 7(2):243–258, 1978.

[11] A. Rosenfeld. Axial representations of shape. *Computer Vision, Graphics and Image Processing*, 33:156–173, 1986.

[12] P.L. Rosin. Unimodal thresholding. In *Proc. Scand. Conf. Image Anal.*, pages 585–592, 1999.

[13] G. Sanniti di Baja and E. Thiel. (3,4)-weighted skeleton decomposition for pattern representation and description. *Pattern Recognition*, 27(8):1039–1049, 1994.

[14] L.G. Shapiro and R.M. Haralick. Decomposition of two-dimensional shapes by graph-theoretic clustering. *IEEE Trans. PAMI*, 1(1):10–20, 1979.

[15] K. Siddiqi and B.B. Kimia. Parts of visual form: Computational aspects. *IEEE Trans. PAMI*, 17(3):239–251, 1995.

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