CHAPTER 1

COMPUTING GLOBAL SHAPE MEASURES

Paul L. Rosin
School of Computer Science, Cardiff University
P.O. Box 916, Cardiff, CF24 3XF, UK
E-mail: Paul.Rosin@cs.cf.ac.uk

Global shape measures are a convenient way to describe regions. They are generally simple and efficient to extract, and provide an easy means for high level tasks such as classification as well as helping direct low-level computer vision processes such as segmentation. In this chapter a large selection of global shape measures (some from the standard literature as well as other newer methods) are described and demonstrated.

1. Introduction

Formed by the eye and therefore, like the eye,
Full of strange shapes, of habits and of forms,
Varying in subjects as the eye doth roll
To every varied object in his glance

Love's Labour's Lost, Shakespeare

Shape is an important aspect of visual interpretation, and over the years many schemes have been proposed for representing the shape of objects (see reviews in Refs. 1, 2, 3, 4). Sometimes the goal is to recover three dimensional information from shape content present in the two dimensional scene (see examples in both computer vision[5,6] and human vision[7]), but in many instances full 3D recovery is not essential for interpretation. In fact, as Biederman and Ju showed[8], the human visual system performs very well even in the absence of many cues such as intensity, colour, texture, etc. In computer vision there are many uses for 2D shape analysis. These range from high level applications such as content based image retrieval; even the earliest systems such as QBIC[9] used circularity, eccentricity, moment invariants, etc., while more recent systems (particularly those aimed for trademarks and logos) use many more shape measures[10]. At the other end of the spectrum, shape analysis can be incorporated as a component in low level computer vision processes to direct image segmentation[11], line grouping[12], etc.

While some techniques focus on the interior components (e.g. pixels) of a region, and others prefer to analyse the boundary (see Pavlidis[1]), both provide essentially the same in-

formation if only binary regions are considered (i.e. interior colour and texture are ignored), and we will not distinguish between the two.

2. Shape Decomposition

One approach to analyse shape is to break up the region into simpler parts such as triangles, convex polygons, etc. An example of partitioning is shown in Fig. 1 where the global shape measure of convexity has provided the criterion for selecting appropriate dividing cuts. The overall convexity of the decomposition (into a specified number of parts) is maximised, although individual parts need not be perfectly convex.

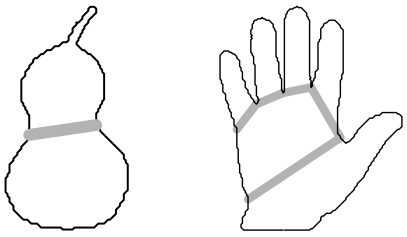

Fig. 1. The gray lines indicate the cuts used to partition the shapes into a pre-specified number of roughly convex parts.

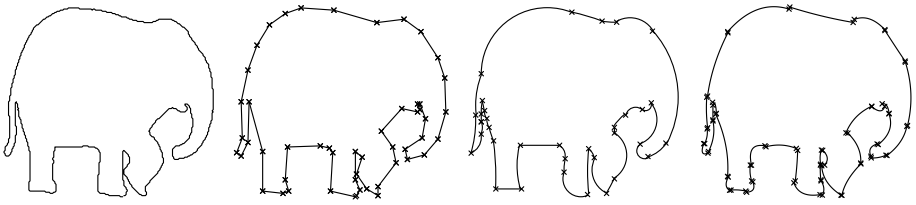

Fig. 2. The original curve is approximated in turn by straight line segments, circular arcs, and parabolic arcs. Breakpoints (i.e. the ends of primitives) are marked by crosses.

More often, segmentation is applied to the boundary which is broken up into a sequence of basic curve primitives. A popular scheme is polygonal approximation, which is straightforward and efficient, but other primitives are possible. Examples are shown in Fig. 2 of approximation by straight lines, circular arcs[13], and parabolic arcs[14]. All approximations are found by recursively splitting the curve at the point of maximum deviation from the fitted primitive (no parameters are required).

It can be useful to represent the boundary by multiple primitive types. For example, roughly straight sections are represented by straight line segments, and curved sections by arcs. This results in a more compact representation which may also be more effective for matching. Since the different representations generally yield different selected breakpoints

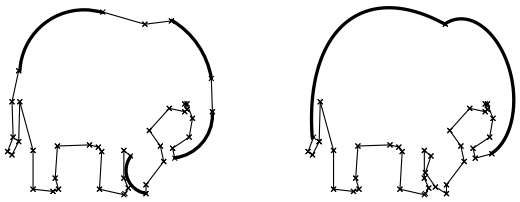

Fig. 3. The polygonal approximation in Fig. 2 is post-processed to replace appropriate sequences of lines by circular arcs, and by elliptic arcs.

the results in Fig. 2 cannot be easily merged. An approach that circumvents this difficulty is to sequentially find each representation, starting with approximating the pixel data by the lowest level primitive type, and then replacing primitive sequences by higher level representations. The results of this are shown in Fig. 3 where circular or elliptical arcs have been substituted for sequences of lines if they provide a better fit.

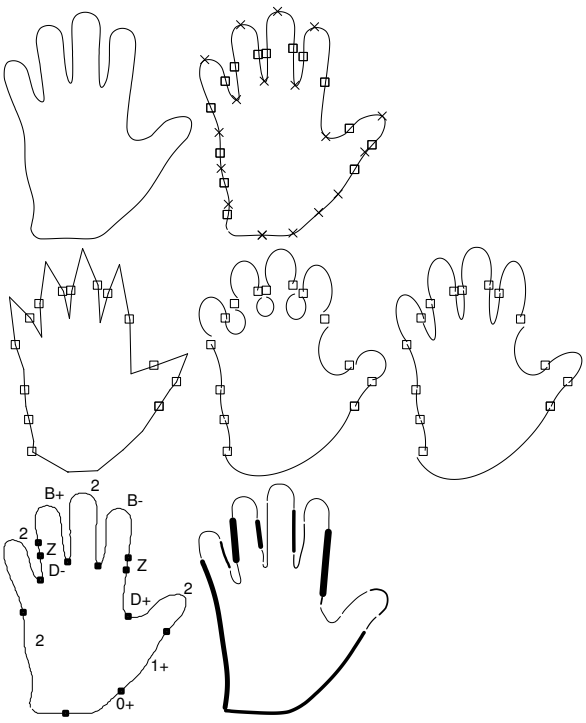

Fig. 4. The curvature extrema (marked by crosses) and point of inflection (boxes) are detected in the input data (top row). Each curve section is fitted by straight line segments, elliptical arcs, and superelliptical arcs (second row). The curvature extrema can also be used to determine codon labels (bottom row). The results of adaptive smoothing are also shown in the bottom row; line thickness is proportional to degree of smoothing.

An alternative approach is to determine the breakpoints by a fixed rule such as points of inflection as shown in Fig. 4. This enables all sections of curve to be easily represented by all the curve primitives (in this case straight line segments, elliptical arcs, and superelliptical arcs)[15].

The curvature extrema can be used to decompose the curve. One example is Asada and Brady's curvature primal sketch[16] which analyses the scale space tree of curvature extrema of a region's boundary to extract the following curve primitives: corner, smooth join, crank, end, and bump. Another well known scheme is the set of codons which were defined by Hoffman and Richards[17] as sections terminating in curvature minima, and divided into five types depending on the curvature extrema within the section. Their catalogue of codon types was extended by Rosin[18] to incorporate straight sections and endings, and an example of the application of this extended codon set is given in Fig. 4.

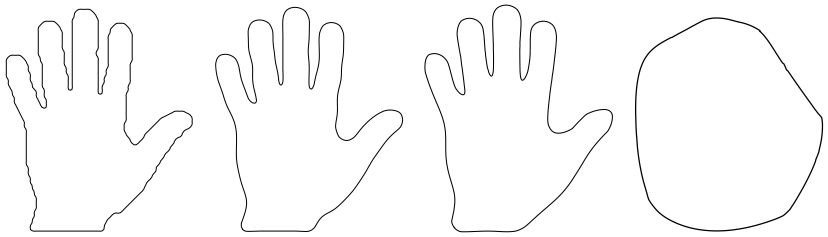

Fig. 5. The global natural scales of the original pixel data (left). The second natural scale was used for processing in Fig. 4.

Detecting curvature extrema requires some filtering to eliminate responses due to noise. One possibility is to determine adaptively an appropriate degree of smoothing for local sections of curve. This can be done in a similar manner to fitting curve primitives as described above[14], and results are shown in Fig. 4. Alternatively, rather than represent each part of the curve at a single scale multiple scales can be employed. The results of representing the hand data at its "natural scales"[19] are presented in Fig. 5. The curve has been smoothed over a range of scales, and those scales maximising a significance measure are considered to contain the most useful information, and are retained.

3. Global Shape Measures

Although the multiple component shape representations described above provide a detailed shape description they have the disadvantage that matching such descriptions is complicated. The same applies to other representations such as medial axes, shock graphs[20], etc. Algorithms such as dynamic programming, interpretation trees, multiple cliques in association graphs, etc. are often used to find the correspondences between subparts, based on their multiple properties (e.g. orientation, position, curvature) and the spatial inter-relationships between groups of subparts.

In contrast, describing regions by global shape measures enables much simpler match-

ing to be used (e.g. a Euclidean distance classifier), and the many standard property based classifiers in the literature can be directly used without modification (e.g. support vector machines, neural networks, etc.). Global shape measures (sometimes called global scalar transforms) just return a single value to describe a region, and therefore they obviously do not provide a one-to-one mapping between shapes and values; for example, there are infinitely many perfectly convex shapes, all of which should yield the same value for any given convexity measure. Nevertheless, if a region is described by a combination of shape measures this should be sufficient to provide discrimination between different categories of shapes. The main disadvantage of global shape measures is that they do not cope well with occlusion, since this can change the global nature of a region's shape to an arbitrary degree.

There are many standard global shape measures in the literature, e.g. moments, Fourier descriptors, auto-regressive models. These types of approaches have been proven to be effective, although a downside is that their values are often not easily understandable. In many applications it is preferable if the region descriptors can be analysed by the domain experts, as this aids validation and development of the computer vision system. In this chapter we shall concentrate on descriptors that have direct intuitive meanings, e.g. convexity.

A wide variety of global shape measures are described in the following sections, and a variety of methods have been used in their construction. They are generally based on the deviation (in some sense) of a region from the best matching perfect instance of the shape. However, while it is usually straightforward to determine whether a region satisfies a shape type or not (e.g. is a region convex, or is a region a rectangle?) there are no universally accepted standard methods for measuring the discrepancy between the region and the canonical form. Nevertheless, since real data tends to be noisy and distorted it is necessary to be able to cope with errors, and so continuous shape measures are required rather than binary shape tests.

Some of the approaches use a combination of primary measurements of the region such as perimeter, area, and diameter (for which there are accepted precise definitions although there may still be difficulty in estimating the values in the presence of noise). From these primary measurements the shape measures can be built, e.g. by defining them as algebraic combinations of the primary measurements. Other approaches explicitly fit a model to the data (using least squares, moments, Fourier transform, etc.) and measure the discrepancy of the fit using the geometric distance between the points and the model or between parameters (e.g. Fourier descriptors) of the fitted model and the region. In the extreme case, similarity can even be measured by determining how many pixels need to be moved, and how far, to change one shape into another[21].

One of the goals of the shape measures is to be able to capture a class of forms which may be fairly general, but also to be specific enough to discriminate against other shapes. A shape measure needs to be invariant under certain transformations, but the particular type of invariance will depend on the application. For instance, a convexity measure should ideally be invariant under any perspective projection. On the other hand, a shape measure to enable digits "6" to be distinguished from "9" in an OCR application should *not* be rotationally

invariant.

A further requirement is that the shape measures should be robust in the presence of noise. In other words, small changes to a region's form should have small effects on the values computed by the shape measures. As we shall see, several of the standard shape measures fail on this count.

3.1. *Eccentricity and Elongatedness*

There are various formulations which measure some notion of the aspect ratio of a region, and they go under several names such as eccentricity, elongatedness, etc. A common and convenient measure uses the central moments μ_{pq}, and is computed as the ratio of the lengths of the axes of the image ellipse $\frac{\mu_{20}+\mu_{02}+\sqrt{(\mu_{20}-\mu_{02})^2+4\mu_{11}^2}}{\mu_{20}+\mu_{02}-\sqrt{(\mu_{20}-\mu_{02})^2+4\mu_{11}^2}}$, which can be simplified and reformulated as $\sqrt{(\mu_{20}-\mu_{02})^2+4\mu_{11}^2}/\left(\mu_{20}+\mu_{02}\right)$ to provide a measure in the range $[0,1]$.

Alternatively, an example of the boundary based techniques is to compute the ratio of the region's width and diameter. This requires determining the diameter (i.e. the longest chord between pairs of boundary points), which can be done in linear time[22]. The width which is the longest chord perpendicular to the diameter can then be easily found. Alternatives such as the ratio of the shortest Feret diameter to the longest chord perpendicular to it are also used.

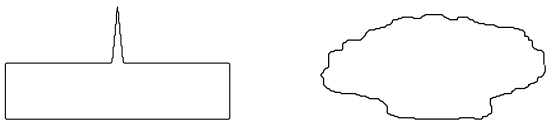

Fig. 6. The spike in the rectangular region on the left makes it appear less elongated than the shell on the right according to the aspect ratio measure, but not according to the moments based measure.

Because the aspect ratio is based on maximum distances it is more sensitive to narrow area spikes than the moments based measure, as is demonstrated in Fig 6.

3.2. *Circularity*

The most common approach to measuring circularity (also called compactness, shape factor, etc.) is to compute $\frac{P^2}{A}$ where P and A are the perimeter and area of the region. Low values correspond to circular shapes (4π in the limit).

Although this is a convenient formulation, there are problems due to the artifacts caused by digitisation[23,24], and moreover, since perimeter estimates are extremely sensitive to noise (prone to overestimation), then in such cases the value of the measure becomes inflated, incorrectly assigning non-circular values to noisy circles.

An alternative that avoids this problem was suggested by Haralick[25]. Given the distance R between the centre and any point on the perimeter his measure is $\frac{\mu_R}{\sigma_R}$, where high values

correspond to circular shapes. Note that the centre should be calculated as the centroid of the polygon and not the average boundary coordinate values, as otherwise it will be biased in the event of noise being unevenly distributed along the boundary.

Figure 7 demonstrates how the $\frac{P^2}{A}$ measure is confused by the circle's boundary texture, and incorrectly estimates the pear as the more circular of the two regions, Although the undulation superimposed on the circle causes its value of $\frac{\mu_R}{\sigma_R}$ to decrease compared to a perfect circle it is still higher than the pear's value, which is therefore correctly estimated as the less circular of the regions.

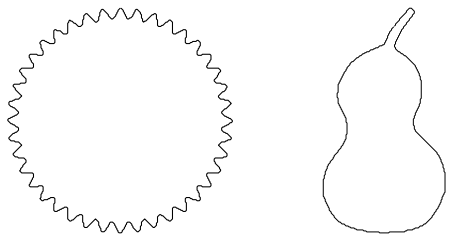

Fig. 7. The undulations on the circle cause it to be measured as less circular than the pear according to $\frac{P^2}{A}$, but not according to $\frac{\mu_R}{\sigma_R}$.

3.3. *Squareness*

Squareness is not such a common shape measure in computer vision, but it is straightforward to define such a measure in various ways. One approach is to use complex Fourier descriptors[26]. If the boundary of the region is given by the points (x_j, y_j) then the coordinates are represented by complex numbers $z_j = x_j + iy_j$ and applying the discrete Fourier transform leads to the complex coefficients comprising the descriptors $c_k = a_k + ib_k = \frac{1}{N}\sum_{m=1}^{N} z_m \exp(-i2\pi mk/N)$. Just the magnitude will be used $r_k = \sqrt{a_k^2 + b_k^2}$, and it can be shown that r_0 represents the region's centroid, and r_1 indicates the size of the region. To make the descriptors scale invariant they are normalised as $w_k = r_k/r_1$. Bowman *et al.*[27] use w_{-3} to measure squareness. Here we modify this approach slightly to include the harmonics and also divide to take into account the remaining harmonics which do not contribute to squareness: $(w_{-3} + w_{-7} + w_{-11} + \ldots)/\sum_{\forall i \notin \{-1,0,1\}} w_i$. Computing the FFT takes $O(N \log N)$ time.

3.4. *Ellipticity*

Voss and Süße describe a method for fitting geometric primitives by the method of moments[28]. The data is normalised into a canonical frame (which for an ellipse they take as the unit circle), by applying an affine transformation. Applying the inverse transformation to the primitive (i.e. the circle) produces the fitted ellipse. An ellipticity measure can be calculated as the RMS of the differences between the normalised moments of the data

(m'_{ij}) and the moments of the canonical primitive (m_{ij}) where only the moments not used to determine the normalisation are included: $1/\left(1+\sum_{i+j\leq 4}(m'_{ij}-m_{ij})^2\right)$.

Alternative approaches to fitting exist, and in Ref. 29 a Least Median of Squares (LMedS) approach is taken. Minimal subsets (i.e. five points) of the data are selected from the boundary at random and used to generate the ellipses through them. The ellipse with the lowest median error over the full data set is retained as the best fit. The boundary point errors are approximated using the orthogonal conic distance approximation method[30] since the true Euclidean distance from each point to the ellipse requires solving a quartic equation. The LMedS is robust to outliers, and enables the ellipse to be fitted more reliably. It is then refitted to the inliers by Fitzgibbon *et al.*'s ellipse-specific least squares algorithm[31] to increase the accuracy.

The summed errors of the ellipse fit $E=\sum_{i=1}^{N} d_i$, where d_i are the conic distance approximations of the N boundary points to the ellipse to the ellipse, need to be made scale invariant. This is done by weighting by the square root of the region's area A, giving the ellipticity measure[32]: $1/\left(1+\frac{1}{N\sqrt{A}}E\right)$.

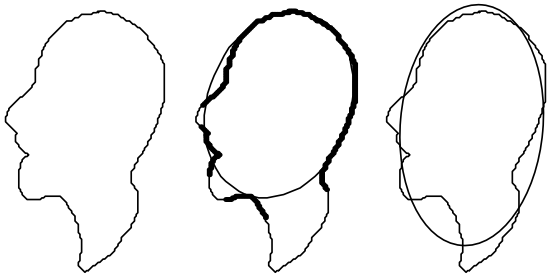

Fig. 8. The region is shown with the fitted ellipse using the LMedS method (inliers plotted bold) and Voss and Süße's moments method.

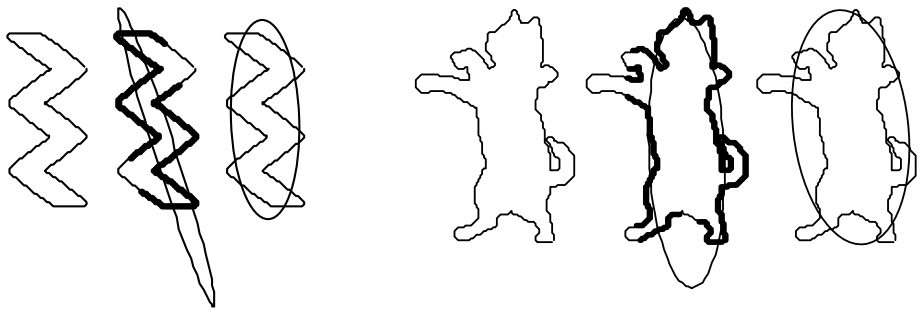

Fig. 9. The two regions show discrepancies between the two ellipticity measures. The LMedS method rates the cat as more elliptical while Voss and Süße's method considers the zigzag region to be more elliptical.

As shown in Figs. 8–9 the two approaches fit quite different ellipses. In the first example (Fig. 8) this does not substantially affect the measured ellipticity values – the overall errors of fit are similar. Figure 9 shows more difference where the LMedS fits most of the cat fairly closely, and therefore assigns it a higher ellipticity value than the zigzag region, whereas Voss and Süße's method ranks the zigzag region as more elliptical than the cat.

3.5. *Triangularity*

Triangles remain triangular under affine transformations, and this naturally leads to affine moment invariants being employed to measure triangularity. We will describe the methods for computing moments based on the region's interior pixels[33], but similar methods are available that operate directly on the region boundary (line moments).[34]

To avoid using higher order moments (which are less reliable) we use the simplest affine moment invariant[35] of the triangle to characterise it: $I_1 = \frac{\mu_{20}\mu_{02}-\mu_{11}^2}{\mu_{00}^4}$, where μ_{pq} are the central moments. For all triangles it can be shown that $I_1 = \frac{1}{108}$. The triangularity measure defined in Ref. 32 is normalised to $[0, 1]$ as $108I_1$ if $I_1 \leq \frac{1}{108}$ and $\frac{1}{108I_1}$ otherwise.

A second approach fits a geometric triangle model to the data, and measures the error between the model and data. Fitting is performed by finding a three line polygonal approximation of the boundary. The summed error E of points along the boundary is computed, and the final triangularity measure is defined in the same manner as the previous ellipticity measure: $1/\left(1 + \frac{1}{N\sqrt{A}}E\right)$. The computation complexity of this measure depends on which polygonal approximation algorithm is used; complexities range from linear to $O(N^3)$.

Fig. 10. The moment based triangularity measure incorrectly ranks the bird between the two triangles, but the polygonal approximation algorithm does not.

A weakness of the moment based approach is that other shapes can yield the same moment invariant (since an infinite number of moments are required to encode all the information in a region). This can leads to errors in the shape measure as shown in Fig. 10.

3.6. *Rectangularity*

The standard approach to measuring rectangularity is to use the ratio of the area of the region X to the area of its minimum bounding rectangle (MBR), i.e. $\frac{area(X)}{area(\mathrm{MBR})}$. The MBR of a convex polygon can be calculated in linear time by Toussaint's[36] "rotating caliper" method. Since the convex hull of a simple polygon can be found in linear time the overall algorithm remains linear.

A weakness of the MBR is its sensitivity to narrow protrusions (while it is relatively insensitive to intrusions). If a modified MBR is found that need contain only most of the region then it should be more robust in the presence of small area deviations in the boundary[32,37]. This robustified MBR is fitted by iteratively minimising the functional $\frac{area(\mathrm{MBR}\,/X)+area(X/\,\mathrm{MBR})}{area(X\cap\mathrm{MBR})}$. The set operations can be done in $O(N\log N)$ time[38]. The functional provides a trade-off between forcing the rectangle to contain most of the data while keeping the rectangle as small as possible. The resulting functional is subtracted from one and provides a rectangularity measure in $[0,1]$.

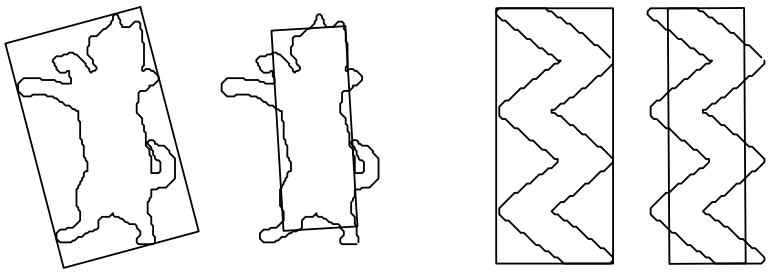

Fig. 11. Two regions with their bounding boxes, both the standard MBR and the robustified MBR.

Figure 11 demonstrates how the MBR fitted by the robust method is a more appropriate description of the gross shape of the regions, and consequently leads to more intuitively correct results. The robustified MBR measure ranks the cat as more rectangular while the standard MBR measure ranks the zigzag region are more rectangular since the cat's standard MBR needs to be very large (compared to the size of the cat) to include the protruding limbs.

3.7. *Rectilinearity*

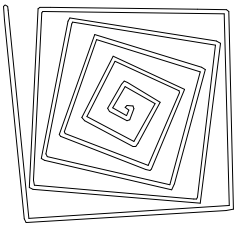

Fig. 12. Although all the polygon's angles are very close to $\pm 90°$ the polygon should not be considered as rectilinear.

A relatively new shape measure is rectilinearity – which is the degree to which a polygon's angles conform to $\pm 90°$. However, directly measuring angles can lead to counter-

intuitive results (see Fig. 12), and so Žunić and Rosin[39] defined the following measure. Let us denote the perimeter of X under the L_1 and L_2 norms as $\mathcal{P}_1(X)$ and $\mathcal{P}_2(X)$. That is, for each straight line segment $e = [(x_1, y_1), (x_2, y_2)]$ making up the boundary of X its length is calculated as $l_1(e) = |x_1 - x_2| + |y_1 - y_2|$ for the L_1 norm and $l_2(e) = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$ for the L_2 norm.

Now let $\mathcal{P}_1(X, \alpha)$ be the L_1 perimeter obtained by first rotating X by the angle α with the origin as the centre of rotation. It can be shown that X is rectilinear if and only if $\mathcal{P}_1(X, \alpha) = \sqrt{2}\mathcal{P}_2(X)$ for some α. This leads to a measure of rectilinearity $\max_{\alpha \in [0, 2\pi]} \frac{\mathcal{P}_1(X,\alpha)}{\sqrt{2}\mathcal{P}_2(X)}$ which is normalised to obtain a value in the range $(0, 1]$: $\frac{\pi}{\pi - 2\sqrt{2}} \left(\max_{\alpha \in [0, 2\pi]} \frac{\mathcal{P}_1(X,\alpha)}{\sqrt{2}\mathcal{P}_2(X)} - \frac{2\sqrt{2}}{\pi} \right)$. While it is possible to solve for α symbolically it is more straightforward to numerically maximise the measure over $[0, 2\pi]$ which can be performed efficiently by multi-resolution sampling of α.

3.8. *Sigmoidality*

A sigmoidality measure determines how much a region is S-shaped. The methods described here preprocess the region to extract a single centreline curve which is the focus of all further analysis. This is easily performed by smoothing the region until the skeleton (obtained by any thinning algorithm) is non-branching.

A modified version[40] of Fischer and Bunke's[41] technique is to fit the cubic $y = ax^3 + bx + c$ after the centreline is first rotated so that its principal axis lies along the X axis. The correlation coefficient ρ is used to measure the quality of fit. Inverse correlation is not expected, and so the value is truncated at zero.

We now describe an alternative method for computing sigmoidality based on a less restrictive characteristic of the sigmoid, namely its single, central point of inflection. For perfect data its presence would be easy to test, but in practise, with real data the sensitivity of curvature estimation to noise makes this difficult. Therefore, instead of identifying zero crossings of curvature the overall distribution of curvature values along the curve is analysed. The curvature at each point is estimated as κ_i using kernels of the Gaussian function and its first two derivatives. Separating the positive and negative curvature values as

$$\kappa_i^+ = \begin{cases} 0 & \text{if } \kappa_i < 0 \\ \kappa_i & \text{otherwise} \end{cases} \qquad \kappa_i^- = \begin{cases} 0 & \text{if } \kappa_i > 0 \\ -\kappa_i & \text{otherwise} \end{cases}$$

then the positive and negative curvature values are summed over the curve to the left and right respectively of the midpoint $\frac{N}{2}$. In addition the total curvature is computed for normalisation purposes:

$$A^+ = \sum_{i=1}^{\frac{N}{2}} \kappa_i^+ \qquad A^- = \sum_{i=\frac{N}{2}}^{N} \kappa_i^- \qquad S = \sum_{i=1}^{N} |\kappa_i| \,.$$

The sections of positive and negative curvature should be restricted to either side of the inflection point and so the quantity $A^+ + A^-$ should be large. Also, the amount of positive

curvature on the left should equal the (absolute) value of the negative curvature on the right, and the discrepancy is measured by $|A^+ - A^-|$. These values are scaled and combined to obtain the following measure $[2(A^+ + A^-)/S - 1][1 - |A^+ - A^-|/S]$. To cope with the curve bending in either direction it is analysed both for κ_i and $-\kappa_i$, and the larger of the two values returned.

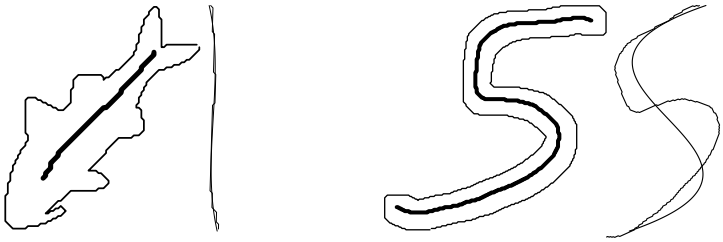

Fig. 13. The regions are shown with their extracted centrelines. Next to each region is its centreline (aligned with the Y axis) and the fitted cubic. The cubic fit measure rates the fish more sigmoidal than the "5".

The weakness of the cubic fit is demonstrated in Fig. 13 in which it rates the fish as more sigmoidal than the digit "5". This is because the fish's centreline is well fit by a cubic, even though it is almost straight, while the shape of the digit does not conform to the smooth, symmetric cubic. The curvature method correctly assigns the digit a higher sigmoidality value since its model is more general.

3.9. *Convexity*

In a similar approach to the MBR method for rectangularity, convexity can be measured with respect to its convex hull. If we denote the convex hull of region X by $\mathrm{CH}(X)$ then the standard convexity measure is defined as $\frac{area(X)}{area(\mathrm{CH}(X))}$. The convex hull of a simple polygon can be found in linear time[42] and so the overall computational complexity of the measure is linear.

Again, like the MBR, the convex hull is sensitive to protrusions but not indentations, and so would estimate the left region in Fig. 14 as more convex then the right region even though the former has greater deviations from the underlying rectangle.

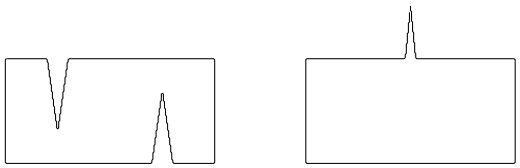

Fig. 14. The convex hull based convexity measure considers the left shape to be the more convex, while Žunić and Rosin's measure more correctly determines the right shape to be the more convex.

Žunić and Rosin[43] proposed a scheme based on perimeter rather than area measurements that responds more equally to intrusions and protrusions. Instead of being based on the convex hull it utilises the minimal rectangle with edges parallel to the coordinate axes which includes X, denoted by $\mathbf{R}(X)$. Also, let $\mathbf{R}(X,\alpha)$ be the minimal rectangle with edges parallel to the coordinate axes which includes X rotated by an angle α around the origin. X is convex if and only for any choice of the coordinate system the L_1 perimeter of X equals the perimeter of the minimal rectangle whose edges are parallel to the coordinate axes and which includes X, i.e., $\mathcal{P}_1(P,\alpha) = \mathcal{P}_2(\mathbf{R}(P,\alpha))$ for any $\alpha \in [0,2\pi]$. This leads to the following convexity measure: $\min_{\alpha\in[0,2\pi]} \frac{\mathcal{P}_2(\mathbf{R}(X,\alpha))}{\mathcal{P}_1(X,\alpha)}$ whose values are in the range $[0,1]$. There is a $O(N^2)$ procedure for the analytic computation of the measure, but like the rectilinearity measure it is convenient to estimate it numerically.

3.10. *Symmetry*

There are many schemes for detecting and quantifying symmetry[44,45]. A simple scheme for measuring n-fold rotation symmetric is to measure the amount of overlap between the region and rotated versions of itself. This can be averaged over $n-1$ rotations by increments of $\frac{2\pi}{n}$ radians. If the result of rotating X by α is $\mathbf{T}(X,\alpha)$ then the symmetry measure is $\frac{1}{n\cdot area(X)} \sum_{i=1}^{n-1} area(X \cap \mathbf{T}(X, \frac{2\pi i}{n}))$. The easiest approach is to use the centroid of X as centre of rotation, but of course for imperfectly symmetric regions this may not be the best centre. An alternative is to optimise the symmetry measure as a function of the centre of rotation, although this is considerably more computationally expensive. Even without the optimisation stage this approach to symmetry detection is more expensive than the previous measures encountered so far since the polygon intersection stage takes $O(N^2)$ time[38].

The two measures of symmetry (with and without optimisation of the centre of rotation) are demonstrated in Fig. 15. Both successfully detect the leftmost region as perfectly symmetric. However, the occlusion in the middle region causes the centroid (shown as the white circle) to be a poor estimate of the centre of rotation and therefore rates this region as less symmetric than the rightmost region. The iterative fitting method correctly determines the centre of rotation (shown as the black circle) and consequently rates this region as more symmetric than the rightmost region.

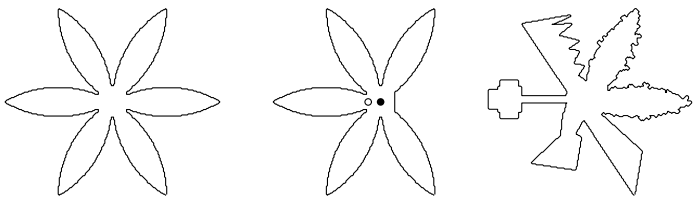

Fig. 15. Both measures of symmetry assign the leftmost region maximal values. For the middle region the centroid (white circle) is a poor estimate of the centre of rotation (black circle) and therefore its use leads to this region being rated as less symmetric than the rightmost region. The iterative fitting method correctly determines the centre of rotation and determines the middle region to be more symmetric than the rightmost region.

3.11. *Chirality*

Chirality is the degree of left or right handedness of a region, and is a useful measure in many disciplines[46] apart from computer vision[47]. A chiral region cannot be superimposed on its mirror image, and this directly leads to a measure of chirality similar to the symmetry measure above. Since the axis of mirror symmetry is not known the overlap is maximised over all possible orientations through the centre, which as above can also be optimised. If the mirror image of X with the line of reflection through (x_c, y_c) at angle α is $\mathbf{M}(X, \alpha, x_c, y_c)$ then the chirality measure is $\max_{\alpha,x_c,y_c}\{1 - area(X \cap \mathbf{M}(X, \alpha, x_c, y_c))\}$ since chiral shapes have minimal overlap. Like symmetry detection this measure is computationally expensive.

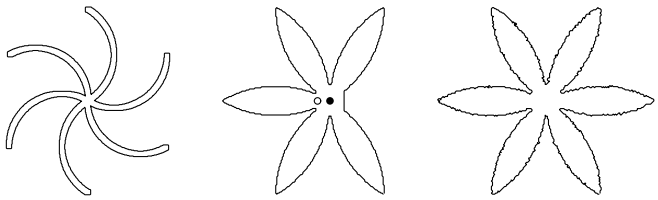

Fig. 16. The left hand region is chiral while the other two regions are not.

Figure 16 shows three regions whose chirality was measured with and without optimisation of the centre of line of reflection. The leftmost region gets a high score from both since there is no mirror reflection that produces any significant overlap with the region. Although the centroid of the middle figure is displaced due to occlusion both measures give zero chirality values. Reflecting about the horizontal line through the centroid is sufficient to get perfect overlap. Finally, the noisy boundaries in the rightmost region cause a fairly small chirality value (0.027) to be initially estimated using the centroid. The chirality value is subsequently refined to zero (to six decimal places) by optimising the centre of line of reflection.

4. Conclusions

Many shape measures have been described in this chapter, but this is not an exhaustive list. Many alternative measures for these shapes exist, as well as measures for other, specialised shapes. The majority of the measures have linear computational complexity in the number of (boundary or interior) points. Some are sensitive to noise or small area protrusions; while more robust methods are available (and were described) they are often more computationally expensive. As can be seen, not all measures are specific, but can also produce their peak response for shapes other than their target shape (e.g. compare the two triangularity measures). Deciding on the most appropriate measures depends on their suitability for particular applications.

To conclude, Figs. 17–27 demonstrate applying all the shape measures described in

this chapter to a variety of regions. The regions are presented ranked into descending order according to each measure, except for width/diameter and $\frac{P^2}{A}$ which are put into ascending order for easier interpretation. These results help provide a qualitative understanding of the performances of the measures; more quantitative comparisons are given in the context of using the shape measures for classification in Refs 10, 32, 37, 39, 40, 43, 48.

$\sqrt{(\mu_{20}-\mu_{02})^2+4\mu_{11}^2}/\,(\mu_{20}+\mu_{02})$

width/diameter

Fig. 17. Eccentricity and Elongatedness

P^2/A

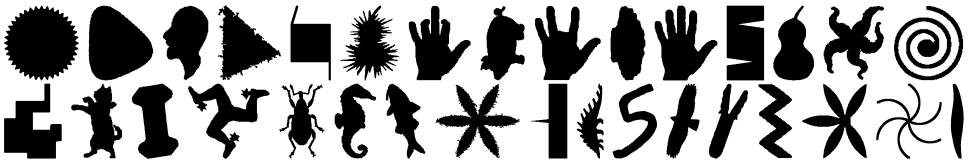

Haralick's μ_R/σ_R measure

Fig. 18. Circularity

References

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

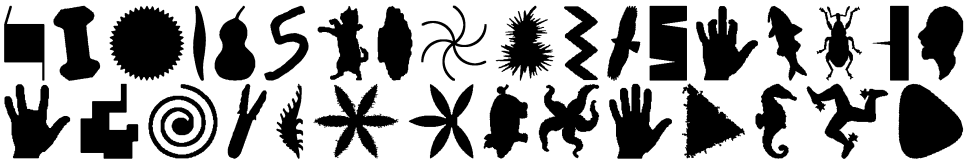

normalised w_{-3} and harmonics

Fig. 19. Squareness

Voss and Süße's $1/\left(1+\sum_{i+j\leq 4}(m'_{ij}-m_{ij})^2\right)$ measure

Rosin's LMedS fit measure

Fig. 20. Ellipticity

measure based on first affine invariant

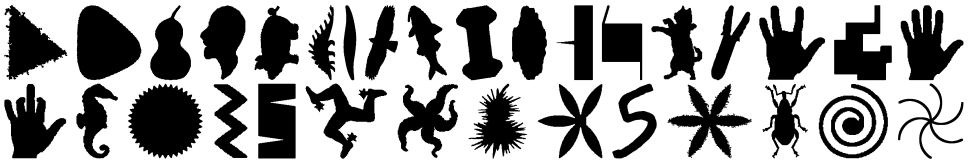

measure based on error of three line polygonal fit

Fig. 21. Triangularity

2. S. Marshall. Review of shape coding techniques. *Image and Vision Computing*, 7(4):281–294., 1989.
3. S. Loncaric. A survey of shape analysis techniques. *Pattern Recognition*, 31(8):983–1001, 1998.
4. D. Zhang and G. Lu. Review of shape representation and description techniques. *Pattern Recog-*

$area(X)/area(\mathrm{MBR})$

$1 - \frac{area(\mathrm{MBR}/X)+area(X/\mathrm{MBR})}{area(X \cap \mathrm{MBR})}$

Fig. 22. Rectangularity

$\frac{\pi}{\pi-2\sqrt{2}}\left(\max_{\alpha\in[0,2\pi]}\frac{\mathcal{P}_1(X,\alpha)}{\sqrt{2}\mathcal{P}_2(X)}-\frac{2\sqrt{2}}{\pi}\right)$

Fig. 23. Rectilinearity

Cubic fit

Curvature distribution

Fig. 24. Sigmoidality

nition, 37(1):1–19, 2004.

5. H.G. Barrow and J.M. Tenenbaum. Retrospective on interpreting line drawings as three-dimensional surfaces. *Artificial Intelligence*, 59(1–2):71–80, 1993.
6. D.P. Mukherjee, A. Zisserman, and J.M. Brady. Shape from symmetry – detecting and exploiting

$area(X)/area(\mathrm{CH}(X))$

$$\min_{\alpha\in[0,2\pi]} \mathcal{P}_2(\mathbf{R}(X,\alpha))/\mathcal{P}_1(X,\alpha)$$

Fig. 25. Convexity

2-fold rotational symmetry – not optimising over centre of rotation

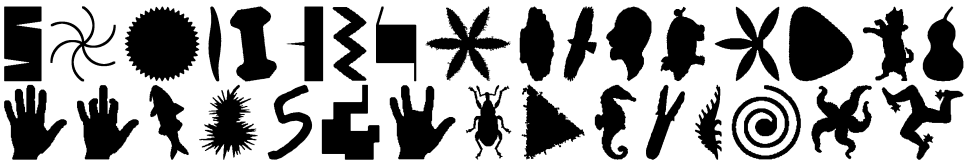

2-fold rotational symmetry – optimising over centre of rotation

3-fold rotational symmetry – optimising over centre of rotation

Fig. 26. Rotational Symmetry

symmetry in affine images. *Phil. Trans. Royal Soc. London, series A*, 351(1695):77–106, 1995.

7. Norman J.F. and Raines S.R. The perception and discrimination of local 3-d surface structure from deforming and disparate boundary contours. *Perception & Psychophysics*, 64(7):1145–1159, 2002.
8. I. Biederman and G. Ju. Surface versus edge-based determinants of visual recognition. *Cognitive Psychology*, 20:38–64, 1988.

optimising over rotation and centre of rotation

Fig. 27. Chirality

9. M. Flickner *et al.* Image and video content: The QBIC system. *IEEE Computer*, 28(9):23–32, 1995.
10. J.P. Eakins, J.D. Edwards, J. Riley, and P.L. Rosin. A comparison of the effectiveness of alternative feature sets in shape retrieval of multi-component images. In *SPIE Conf. Storage and Retrieval for Media Databases*, pages 196–207, 2001.
11. M.A. Shahin and S.J. Symons. Lentil type identification using machine vision. *Canadian Biosystems Engineering*, 45:3.5–3.11, 2003.
12. D.W. Jacobs. Robust and efficient detection of salient convex groups. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 18(1):23–37, 1996.
13. P.L. Rosin and G.A.W. West. Non-parametric segmentation of curves into various representations. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 17:1140–1153, 1995.
14. P. L. Rosin. Non-parametric multi-scale curve smoothing. *Int. J. Pattern Recognition and AI*, 8:1381–1406, 1994.
15. P.L. Rosin. A hierarchy of representations for curves. In Arcelli et al., editor, *Aspects of Visual Form Processing*, pages 465–474, 1994.
16. H. Asada and M. Brady. The curvature primal sketch. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 8(1):2–14, 1986.
17. D. Hoffman and W. Richards. Parts of recognition. *Cognition*, 18:65–96, 1984.
18. P. L. Rosin. Multi-scale representation and matching of curves using codons. *CVGIP: Graphical Models and Image Processing*, 55:286–310, 1993.
19. P. L. Rosin. Representing curves at their natural scales. *Pattern Recognition*, 25:1315–1325, 1992.
20. P.J. Giblin and B.B. Kimia. On the local form and transitions of symmetry sets, medial axes, and shocks. *Int. Journal of Computer Vision*, 54(1):143–156, 2003.
21. H. Sanchez-Cruz and E. Bribiesca. A method of optimum transformation of 3d objects used as a measure of shape dissimilarity. *Image and Vision Computing*, 21(12):1027–1036, 2003.
22. F.P. Preparata and M.I. Shamos. *Computational Geometry*. Springer-Verlag, 1985.
23. A. Rosenfeld. Compact figures in digital pictures. *IEEE Trans. on Systems, Man and Cybernetics*, 4:221–223, 1974.
24. Bogaert J., Rousseau R., Van Hecke P., and Impens I. Alternative area-perimeter ratios for measurement of 2D-shape compactness of habitats. *Applied Mathematics and Computation*, 111:71–85, 2000.
25. R.M. Haralick. A measure for circularity of digital figures. *IEEE Trans. on Systems, Man and Cybernetics*, 4:394–396, 1974.
26. G.H. Granlund. Fourier preprocessing for hand print character recognition. *IEEE Trans. on Computers*, 21:195–201, 1972.
27. E.T. Bowman, K. Soga, and T. Drummond. Particle shape characterization using Fourier analysis. *Geotechnique*, 51(6):545–554, 2001.
28. K. Voss and H. Süße. Invariant fitting of planar objects by primitives. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 19(1):80–84, 1997.

29. P.L. Rosin. Further five-point fit ellipse fitting. *CVGIP: Graphical Models and Image Processing*, 61(5):245–259, 1999.
30. P.L. Rosin. Ellipse fitting using orthogonal hyperbolae and Stirling's oval. *CVGIP: Graphical Models and Image Processing*, 60(3):209–213, 1998.
31. A.W. Fitzgibbon, M. Pilu, and R.B. Fisher. Direct least square fitting of ellipses. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 21(5):476–480, 1999.
32. P.L. Rosin. Measuring shape: Ellipticity, rectangularity, and triangularity. *Machine Vision and Applic.*, 14(3):172–184, 2003.
33. S. Maitra. Moment invariants. *Proc. IEEE*, 67:697–699, 1979.
34. M.H. Singer. A general approach to moment calculation for polygons and line segments. *Pattern Recognition*, 26(7):1019–1028, 1993.
35. J. Flusser and T. Suk. Pattern recognition by affine moment invariants. *Pattern Recognition*, 26:167–174, 1993.
36. G.T. Toussaint. Solving geometric problems with the rotating calipers. In *Proc. IEEE MELECON '83*, pages A10.02/1–4, 1983.
37. P.L. Rosin. Measuring rectangularity. *Machine Vision and Applic.*, 11:191–196, 1999.
38. M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 2nd edition, 2000.
39. J. Žunić and P.L. Rosin. A rectilinearity measurement for polygons. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 25(9):1193–1200, 2003.
40. P.L. Rosin. Measuring sigmoidality. In *Computer Analysis of Images and Patterns*, volume 2756 of *LNCS*, pages 410–417, 2003.
41. S. Fischer and H. Bunke. Identification using classical and new features in combination with decision tree ensembles. In J.M.H. du Buf and M.M. Bayer, editors, *Automatic Diatom Identification*, pages 109–140. World Scientific, 2002.
42. D. McCallum and D. Avis. A linear algorithm for finding the convex hull of a simple polygon. *Inform. Process. Lett.*, 9:201–206, 1979.
43. J. Žunić and P.L. Rosin. A new convexity measurement for polygons. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 26(7):923–934, 2004.
44. J.J. Leou and W.H. Tsai. Automatic rotational symmetry determination for shape analysis. *Pattern Recognition*, 20(6):571–582, 1987.
45. G. Marola. On the detection of the axes of symmetry of symmetric and almost symmetric planar images. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 11(1):104–108, 1989.
46. M. Petitjean. Chirality and symmetry measures: A transdisciplinary review. *Entropy*, 5:271–312, 2003.
47. Y. Hel-Or, S. Peleg, and D. Avnir. Characterization of right handed and left handed shapes. *Computer Vision, Graphics and Image Processing*, 53(2):297–302, 1991.
48. M.L. Hentschel and N.W. Page. Selection of descriptors for particle shape characterization. *Part. Part. Syst. Charact.*, 20:25–38, 2003.