Matching 2D Shapes using their Symmetry Sets*

Arjan Kuijper
RICAM, Linz
Austria

Ole Fogh Olsen, Philip Bille
IT-University of Copenhagen
Denmark

Peter Giblin
University of Liverpool
United Kingdom

Abstract

We introduce a shape descriptor that is based on the Symmetry Set. This set represents pairwise symmetric points and consists of several branches. The begin and end points of the branches relate to extrema of the curvature along the shape. Consequently, extrema of the curvature are pairwise connected via a Symmetry Set branch with a certain finite length. The novel shape descriptor is given by a string representing these extrema, together with the pair wise connections and a length measure. Next, an algorithm is given to match strings. This algorithm is based on a modified shortest path algorithm, taking into account the allowed changes of the Symmetry Set. Examples show the usability of the presented theory, applied to different types of shapes, including noise and occlusions.

1. Introduction

In shape analysis, much effort has been put into the research on the skeleton, or Medial Axis [1], as a simplified representation of the shape. As it was soon realised, the Medial Axis it itself didn't carry enough information and sophisticated extensions were built, like the Shock Graph method [6]. Basically, each points on the Medial Axis is endowed with some augments related to the distance to the shape itself or related to its neighbours. Next, the potential changes of the Medial Axis were investigated, yielding a set of possible transition [5]. In that way different shapes can be related to each other for shape indexing and retrieval [8, 9].

The results on transitions boiled down from the results on the possible transitions of the Symmetry Set. This set, containing the Medial Axis as subset, has been thoroughly studied in [2]. Recently, a data structure was presented for the Symmetry Set [7], using information of the evolute of the shape. The data structure can be visualized by a sequence of nodes that are pair wise joined. Its main advantage was claimed to be the low complexity of operations on it.

In this paper we use the idea of representing Symmetry Sets as a sequence. In major contrast to [7], we relate this sequence directly to the shape. As different shapes have different sequences, we propose to build a cost matrix based on these sequences. The similarity of shapes is then measured as a shortest path in this matrix.

2. Symmetry Sets

The Symmetry Set is defined as the closure of the loci of the circles tangent to a shape. Then for every pair of points p_i and p_j and their normal vectors N_1 and N_2, $i \neq j$, that contribute to the Symmetry Set point

$$(p_i - p_j).(N_i \pm N_j) = 0. \tag{1}$$

Next, the centre of the circle - the location of the Symmetry Set point - is given by solving r in $p_i - rN_i = p_j \pm rN_j$.

A branch of the Symmetry Set is given by a connected set of centres of circles. The end points of a branch are the closures of these sets, obtained when the two points p_i and p_j coincide. For the Medial Axis, such a point is an end point. In the Symmetry Set, these points come in pairs, as the Symmetry Set consists of distinct curves. At these points p_i the shape has a local extremum of the curvature κ. Following the shape, one can label the order of appearance of these points, yielding thus a sequence of end points. Connecting the end points pair wise and augmenting each connection with 'special points' that arise on the Symmetry Set, gives the string structure proposed in [7].

An example is given in Fig. 1. On the left, a fish shape is taken from a common data set [8, 9]. On the right, the string structure - without special points - is shown.

*This work is part of the DSSCV project supported by the IST Programme of the European Union (IST-2001-35443). A.K. acknowledges for funding the Radon Institute for Computational and Applied Mathematics, the Austrian Science Funding Agencies FWF, through the Wittgenstein Award (2000) of Peter Markowich, and FFG, through the Research & Development Project 'Analyse Digitaler Bilder mit Methoden der Differenzialgleichungen'.

The 18th International Conference on Pattern Recognition (ICPR'06)
0-7695-2521-0/06 $20.00 © 2006 IEEE

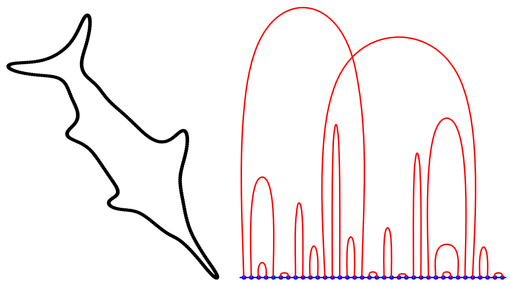

Figure 1. A fish shape and its corresponding sequential representation.

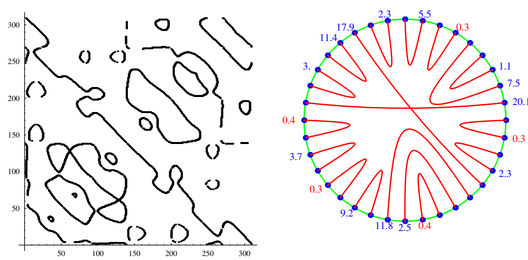

Figure 2. The pre-Symmetry Set and the circular representation of the fish shape of Fig. 1.

3. Closed form representation

The Symmetry Set can become complicated, especially for concave shapes. The centres of the circles can be located anywhere. By taking the locations of the tangency of the circle, instead of its centre, one obtains pairs of so-called 'pre-Symmetry Set' points. They are shown in Figure 2 on the left.

In this diagram, branches of the Symmetry Set are visible as curves. Note that the shape is closed, so the left part of the diagram is connected to the right part, and the bottom to the top. At end point of the Symmetry Set branches, $p_i = p_j$, which is the diagonal. This diagonal can also be regarded as an identity mapping of the shape on itself, and therefore as the shape.

Consequently, points on the shape (diagonal) are connected to points on the shape (diagonal) via the curves in the pre-Symmetry Set. As the shape is closed and not self-intersecting, it can be represented as a circle. The connections of points on the shape are visible as cords. An example is given in Figure 2 on the right.

Next, each cord can be assigned a weight. This weight is the number of points on a branch in the pre-Symmetry Set, divided by the sum of all branches in the pre-Symmetry Set that intersect the diagonal. So the weights sum up to 1. In Figure 2 this number is given as a percentage.

A straightforward manner to store the information given by the circle with cords, is by creating a vector with the same dimension as the number of end points. Each coordinate of the vector then get the value of the relative length of the cord that is related to it. Consequently, the coordinates sum up to 2.

When all cords have different length, the cords can easily be reproduced from this vector. However, the connectivity information is lost if two cords have the same length. Therefore, each coordinate of the vector contains besides the length also the coordinate to which it relates.

4. Matching strings

Given two shapes, comparison can be done visually by comparing their circle diagrams A and B. As the information of these diagrams consists of points and cord, the points are mapped such that the number of coinciding cords is highest. Obviously, the ordering of points may not change. As the parametrisation has an arbitrary begin point, also all rotated versions of A up to 2π must be taken into account. Furthermore, the number of cords of both circles may differ, as well as the way the cords are connected.

From the transitions of the Symmetry Set [5] it follows that a cord (a branch of the Symmetry Set) may (dis-) appear in a transition where two end points meet and a cord (dis-) appears. As the removal of a cord in one circle to optimise matching relates to introducing a cord in the other circle, it suffices to consider removing cords. Consequently, a cord connecting two neighbouring end points is allowed to vanish - in the mapping such a cord may be removed.

The matching of two circle diagrams A and B can be done as follows. Let $\{A_i\}_{i=1\ldots n}$ and $\{B_j\}_{j=1\ldots m}$ denote the vectors with the lengths of the branches. Then $M(i,j) = f(A_i, B_j)$ is the cost matrix, where f is some distance measure. In the remainder we shall use $f(x,y) = x.y/\|x\|\|y\|$, but other choices, like $f(x,y) = \|x-y\|$, can be applied as well.

If A=B and the starting positions are equal, tr M describes the inner product between two identical vectors and equals one. If the starting positions are different, say a, the trace of the a permuted version of M (the a^{th} column becomes the first one) equals one.

To maximise the matching, a path $P = \{M(i_k, j_l)\}$ is to be found in M, such that each row and column i_k and j_l are present only once - each point can be matched only once. For the two examples given above, this is simple. For different shapes, it must be taken into account that two neighbouring points and their connecting cord may be removed.

The 18th International Conference on Pattern Recognition (ICPR'06)
0-7695-2521-0/06 $20.00 © 2006 IEEE
IEEE Computer Society

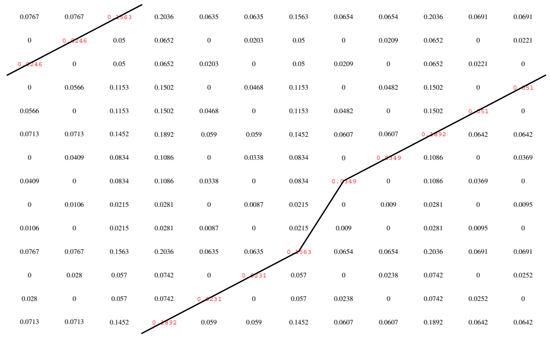

Figure 3. Cost matrix and optimal path for two different shape circles.

This relates to the matrix in removing two subsequent rows or columns.

Next, when two points are matched, automatically the two points to which they are connected, must be matched. For simplicity, one can state that when two cords are given by (i_k, i_{k+1}) and (j_l, j_{l+1}), i_k and j_l can only be matched, if i_{k+1} and j_{l+1} are matched, and that the matchings $M(i_k, j_{l+1})$ and $M(i_{k+1}, j_l)$ (connecting a begin point in circle A to an end point point in B, vice versa) are forbidden.

An example of a matrix M is given in Figure 3. The origin is bottom left. The line through the matrix denotes the optimal match. As one can see, the matrix contains zeros, denoting the forbidden entries. When two subsequent values along the line are equal, the off-diagonal neighbouring points are zero, as described above. As the vectors have different length, the line makes a jump. The jump skips two rows. In general, jumps skip an even number of rows or columns, since a jump resembles the removal of a number of cords, each with two points.

The derivation of the Symmetry Set given a shape is described in [2, 7]. It basically boils down in computing all zero crossings in Eq. 1 for all point pairs (p_i, p_j). These points pairs form the pre-Symmetry Set as shown in Fig. 2, left. Then the distinct Symmetry Set branches that intersect the diagonal are derived, with the locations at the diagonal and their lengths. This gives a set with elements $A_i = (e_1, e_2, L)_i$, with e_1 and e_2 the e_1^{th} and e_2^{th} position on the diagonal, and L the relative length of the branch.

Next, on each cord that is allowed to vanish, the two points are marked as 'begin' or 'end' point. Note that if two cords are nested, both are allowed to vanish. If the cross each other, they cannot be removed.

Let $L_i \in A$ and $L_j \in B$, then the cost matrix is built up as $M(i,j) = 0$ if A_i and B_j are a combination of a begin and an end point, and $M(i,j) = L_i L_j$, elsewhere.

The path with maximal value is found by using a shortest path algorithm [4] on $-M$. A graph structure on M can be defined by taking as vertices the entries of M. The directed edges for $M(i,j)$ are potential connections between the three neighbouring entries $M(i+1,j)$, $M(i,j+1)$, and $M(i+1,j+1)$. The edge weight is set as follows:

- If $M(i+1,j+1) = M(i,j)$ and $M(i+1,j) = M(i,j+1) = 0$ two begin points of a cord are matched and the only allowed edge is $M(i+1,j+1) \rightarrow M(i,j)$ with cost $M(i+1,j+1)$.
- If $M(i+1,j+1) = 0$, this position is not allowed and the only allowed edges, denoting a possible skip, are $M(i+1,j+1) \rightarrow M(i+1,j)$ and $M(i+1,j+1) \rightarrow M(i,j+1)$, both with cost 0.
- Else three edges are possible: $M(i+1,j+1) \rightarrow M(i,j)$ with cost $M(i+1,j+1)$, and $M(i+1,j+1) \rightarrow M(i+1,j)$ and $M(i+1,j+1) \rightarrow M(i,j+1)$, both with cost 0.

Obviously, to compute the complete path from a point to itself, one should handle the boundaries of M properly. To find the shortest path solution, it suffices to take the shortest paths through the entries of one column or row and take the minimum of them. To find the optimal solution taking into account all possible starting positions, this procedure is ran for all entries $M(i,1)$ and $M(1,j)$.

5. Results

In the remaining we used shapes from an existing data base [8, 9]. These shapes are the boundary of 128×128 pixel sized black and white images. These images are derived from an automatic segmentation that included errors due to shadows and highlights [9].

The derivatives of a Gaussian filter are applied to this sequence to find a reasonable estimation of the derivatives of the shape parametrisation. The normal vector is obtained at a scale of 4.5 pixels. We note that using a small scale resembles applying a (small) mean curvature motion to the shape [3]. This blurring of shapes has the property that it removes small details. This may be regarded as a disadvantage, but on the other hand no removal of spurious details, or whatever adjustments to the data need to be carried out. The corresponding string, pre-Symmetry Set and circle diagram are shown in Figures 1-2.

Next, 33 different fish and tool shapes are compared. The results are shown in Figure 4. The images show the shapes while the numbers are the score of the match. The first column shows the best match, the second column the second-best match and so on. As the matching of any shape with itself equals 1, the first column also represents the shape to be matched.

Figure 4. Matching of fishes and tools

The left part shows the fish-like part. The matching has a preference for matching fins. This is due to the fact that fins are represented by prominent extrema of curvature and therefore cords in the Gauss diagram. The presence of the three tools is due to the fact that these tools at coarse scale can be identified as a fish with only two tail fins and no other fins. The two artificial shapes in rows 3 and 4 are matched with the fishes

The tools part on the right in Figure 4 yields matchings with only tools. Note that differences in size of orientation do not influence the matching.

Although partial occlusion is already present in the data base, as result of the segmentation procedure [9], we manually occluded the fish in row 13 of 4 as well. In row 14 the top tail fin of this fish is removed, while in row 15 its bottom tail fin is occluded. These two shapes match to each other and the original fish. A match with the fishes in row 3 and 4 is also found. This may be due to the fact that these shapes are both artificially drawn.

Since the small scaling to obtain the normal vectors reduces the influence of noise, we removed and added manually black and white pixels to the binary image in row 13.This avoids the introduction of relatively "harmless" Gaussian-like noise. The noisy fish is shown in row 16. Still, the first match is the original fish with the highest matching score of all non-trivial matchings.

Both examples shows the stability of the method under moderate noise and occlusion.

6. Summary and Conclusions

We introduced a new way to represent and compare shapes based on the Symmetry Set, a generalisation of the Medial Axis. This string representation uses the end point of the Symmetry Set branches and the relative length of the branch in the pre-Symmetry Set diagram. The end points represent the extrema of curvature of the shape. Therefore, the representation links these extrema pair wise.

The representation allows the matching of shapes by comparing strings, where we choose to take the inner product of appropriate sub sets of these strings. These sub sets are defined by applying allowed changes of the Symmetry Set. The maximal matching is found by an adapted shortest pad algorithm that finds the optimal sub sets.

Examples show the usability of the proposed method, also under influence of occlusion and noise. Future work will focus on tests with larger data bases, improvement of the shortest path based algorithm, and on the influence of alternative difference measures besides the inner product.

References

[1] H. Blum. Biological shape and visual science (part i). *Journal of Theoretical Biology*, 38:205–287, 1973.

[2] J. W. Bruce, P. J. Giblin, and C. Gibson. Symmetry sets. *Proceedings of the Royal Society of Edinburgh*, 101(A):163–186, 1985.

[3] F. Cao. *Geometric Curve Evolution and Image Processing*, volume 1805 of *Lecture Notes in Mathematics*. Springer Verlag, 2003.

[4] T. Cormen, C. Leiserson, and R. Rivest. *Introduction to Algorithms*. MIT Press, 1993.

[5] P. J. Giblin and B. B. Kimia. On the local form and transitions of symmetry sets, medial axes, and shocks. *International Journal of Computer Vision*, 54(1/2):143–156, 2003.

[6] B. Kimia. On the role of medial geometry in human vision. *Journal of Physiology - Paris*, 97(2-3):155–190, 2003.

[7] A. Kuijper, O. Olsen, P. Giblin, P. Bille, and M. Nielsen. From a 2D shape to a string structure using the symmetry set. In *Proceedings of the 8th European Conference on Computer Vision - ECCV 2004, Part II (Prague, Czech Republic, May 11-14, 2004)*, volume II, pages 313–326, 2004. LNCS 3022.

[8] M. Pelillo, K. Siddiqi, and S. Zucker. Matching hierarchical structures using association graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(11):1105–1120, 1999.

[9] T. Sebastian, P. Klein, and B. B. Kimia. Recognition of shapes by editing shock graphs. In *Proceedings of the 8th International Conference on Computer Vision (2001)*, pages 755–762, 2001.