

Rectification of the Chordal Axis Transform and a New Criterion for Shape Decomposition

Lakshman Prasad

Space and Remote Sensing Sciences Group (ISR-2),
International, Space, and Response Division,
Los Alamos National Laboratory,
Los Alamos, NM 87545, USA.
prasad@lanl.gov

Abstract. In an earlier work we proposed the chordal axis transform (CAT) as a more useful alternative to the medial axis transform (MAT) for obtaining skeletons of discrete shapes. Since then, the CAT has benefited various applications in 2D and 3D shape analysis. In this paper, we revisit the CAT to address its deficiencies that are artifacts of the underlying constrained Delaunay triangulation (CDT). We introduce a valuation on the internal edges of a discrete shape's CDT based on a concept of approximate co-circularity. This valuation provides a basis for suppression of the role of certain edges in the construction of the CAT skeleton. The result is a rectified CAT skeleton that has smoother branches as well as branch points of varying degrees, unlike the original CAT skeleton whose branches exhibit oscillations in tapered sections of shapes and allows only degree-3 branch points. Additionally, the valuation leads to a new criterion for parsing shapes into visually salient parts that closely resemble the empirical decompositions of shapes by human subjects as recorded in experiments by M. Singh, G. Seyranian, and D. Hoffman.

Keywords: Shape, Delaunay triangulation, chordal axis transform, medial axis, skeleton, shape decomposition, morphology, co-circularity, shape graph, grouping, chord strength.

1 Introduction

The skeleton of a shape is an important descriptor that provides structural information about the shape. Skeletons are used to compare shapes, identify shape parts, and, in case of thin objects such as textual characters, even represent the shapes themselves. Blum [1] defined the skeleton of a two dimensional shape with a continuous closed contour as the locus of centers of maximal discs (i.e., discs touching the shape contour at two or more points) interior to the shape, with each center attributed the radius of the corresponding maximal disc. This definition of a shape's skeleton is known as the medial axis transform (MAT) of the shape. While the MAT is an elegant characterization of the skeleton of a shape with a continuous boundary, it has proved to be difficult to use as a practical tool to analyze shapes. Indeed, for example, minor oscillations in shape contours due to insignificant features or noise result in skeletal

branches that are not easy to isolate, a skeletal feature may be spatially far-removed from the contour feature it represents, and a skeleton part may greatly exaggerate or diminish the importance of the contour feature that gave rise to it (Figs. 1 & 2). The medial axis transform is not defined for shapes specified by discretely sampled contours, as typically encountered in digital imagery. Several extensions of the MAT to discrete shapes have been formulated using pixel morphology [3] and geometry [4, 5]. These methods, however, require uniform or well-sampled representations of the shape to yield satisfactory skeletons (Fig. 8d).




Fig. 1. Rectangle with boundary feature




Fig. 2. MAT skeleton exaggerating feature

2 Background

In earlier works [6, 7, 8] we proposed the chordal axis transform (CAT) as a more useable and stable definition of the skeleton of a shape that is robust to sparse and uneven sampling of shape boundary. Since then it has gained currency among researchers in the area of 2D and 3D shape analysis and modeling [12-15]. In this section, we will review the CAT, its strengths, and drawbacks to set the context for this paper.

Definition 1: A *maximal chord of tangency* (Fig. 3) connects two points of tangency of a maximal disc inscribed in a shape such that at least one of the two arcs of the maximal disc's bounding circle subtended by the chord is free of points of tangency with the shape's boundary.

Definition 2: The *Chordal Axis Transform* (CAT) of a planar shape is the set of all ordered pairs (p, δ) , where p and δ are either the midpoint and half the length, respectively, of a maximal chord of tangency, or the center and radius, respectively, of a maximal disc with three or more maximal chords of tangency.

Although the definition of the CAT appears to be a variation of that of the MAT, there are important differences between the two transforms. First, the CAT, as defined, yields a piecewise smooth disconnected protoskeleton (Fig. 5). By joining

the midpoints of the maximal chords of a maximal disc with three or more chords to the center of the maximal disc if the center lies within the polygon determined by the chords, or to the center of the longest chord otherwise, we obtain a connected skeleton of the shape (Fig. 6).

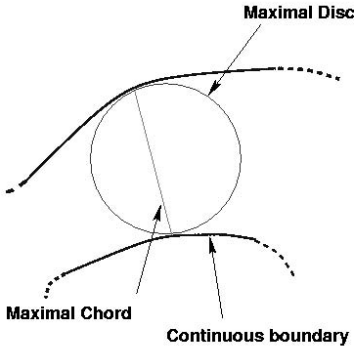


Fig. 3. A maximal disc and the associated maximal chord of tangency of a shape with continuous boundary

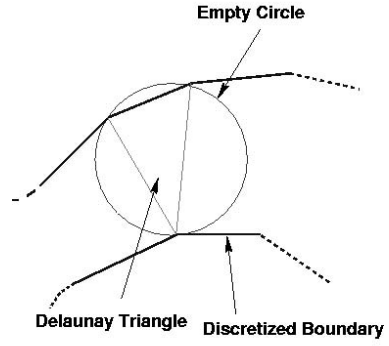


Fig. 4. An empty circle and the associated Delaunay triangle of a shape with discrete boundary




Fig. 5. CAT proto-skeleton of shape in Fig. 1




Fig. 6. Connected CAT Skeleton

Second, the CAT can be stably defined for a shape whose boundary is discrete (i.e., specified as sequences of points separated in space) (Fig. 8). This is done by replacing maximal discs by empty circles that pass through three or more points of the shape's discrete boundary such that no circle contains a boundary point in its interior that is visible to two boundary points lying on the circle (Two vertices u and v of a simple polygon are *visible* to each other, if the line segment joining u and v does not intersect the exterior of the polygon). Each such empty circle identifies a triangle whose edges lying in the shape's interior replace maximal chords of tangency in the discrete

version of the CAT (Fig. 4). The triangles so formed are indeed the Delaunay triangles of a constrained Delaunay triangulation (CDT) [2] of the shape's interior. It is worth noting here that this extension of the CAT to discrete shapes is natural from the point of view of constructing skeletons. This is because the constrained Delaunay triangulation is the geometric dual of the generalized Voronoi axis [2] of the contour point set. Indeed, the MAT is essentially the Voronoi skeleton of a shape. In using the dual of the Voronoi axis, we can define a more robust and manipulable skeleton than the MAT that applies to discrete shapes whose boundaries are sparsely and unevenly sampled. We can also ensure strong invertibility of the skeleton to recover the shape [8]. The CDT of a shape's interior gives rise to three kinds of triangles, namely *Junction* triangles (J) that have all their edges inside the shape and signify shape bifurcations, *Sleeve* triangles (S) that have one edge in common with the shape boundary and signify shape prolongations, and *Terminal* triangles (T) that have two edges in common with the shape boundary and signify shape terminations. The connected CAT skeleton for a discretized shape is obtained from its CDT by i) joining the midpoints of the internal edges of each S-triangle by a line segment, ii) joining the midpoints of the internal edges of each J-triangle to its circumcenter if the triangle is acute, or to the midpoint of its longest side if it is obtuse (Figs. 7, 8). These localized constructions ensure that the CAT skeleton does not cross the boundary of the shape irrespective of the sparsity of sampling. Henceforth we will restrict ourselves to the structure of the CAT skeleton in the rest of the paper and direct the interested reader to [8, 9] for other details and implications of the CAT.




Fig. 7. Construction of skeletal segments in the various types of triangles in the CDT of a shape: from left to right, skeleton segments in a sleeve triangle, an obtuse junction triangle, an acute junction triangle, and a terminal triangle




Fig. 8. Construction of the CAT skeleton of a sparsely and unevenly sampled shape and comparison with its discrete MAT skeleton: (a) A discrete shape, (b) CDT and construction of sleeve skeletal segments, (c) Connected skeleton after construction of skeletal segment in the junction triangle, (d) Discrete MAT skeleton obtained by connecting adjacent Voronoi vertices (circumcenters of triangles) of the contour points. The comparison shows the stability of the CAT skeleton over that of the MAT of a shape with sparsely sampled contour

The CAT skeleton of a discrete shape is robust in the face of sparse and irregular presentations of the shape boundary. It also allows easy excision of insignificant and noisy features via a simple pruning criterion [6, 7, 8]. Finally, it enables parts-based decomposition of shapes into structurally meaningful components [8, 9].

3 Drawbacks of the Discrete CAT Skeleton

However, the CAT skeleton, as defined above, has certain structural deficiencies (which are also present in discrete realizations of the MAT skeleton) that warrant rectification. The CAT skeleton exhibits oscillations through shape regions that are tapered (Fig. 9). The CAT skeleton allows only branches of degree three to represent shape ramifications even when higher degree branches are more natural to represent them (Fig. 10). These deficiencies are artifacts of giving equally important roles to all chords (i.e., internal triangle edges of the shape's CDT) in constructing the CAT skeleton. Indeed, in a previous work [8], we considered special cases where more than three points on a shape boundary are co-circular with respect to an empty circle. We noted that the triangle edges that form the chords of the polygon determined by the co-circular points are not uniquely defined (i.e., any triangulation of the interior of the co-circular polygon will be consistent with the Delaunay criterion of triangulation.) We proposed that the restriction of the shape's skeleton to such a polygon be constructed by joining all the midpoints of the polygon's edges that are internal to the shape to the circumcenter of the polygon or, if the latter falls outside the polygon, to the midpoint of longest edge of the polygon that does not lie on the shape boundary. In effect, we discarded internal shape edges that are common to two co-circular triangles in the CDT of a shape. We will generalize this notion of co-circularity to define a valuation on the chords of a shape that will help filter chords which are common edges of nearly co-circular triangles. The motivation for this is to prevent common edges of nearly co-circular triangles from participating in the construction of the skeleton. This will greatly reduce skeleton oscillations in tapered regions of shapes. Indeed, consider two adjacent, nearly co-circular, sleeve triangles. Unless the external edges of the triangles are parallel, the midpoints of the internal edges of the two triangles will not lie on a straight line, thus producing an oscillation in the skeleton. If the common internal edge of this triangle pair is discounted, then the skeleton of the polygon determined by the triangle pair is given by the line segment joining the midpoints of the remaining two internal edges, thus locally rectifying the CAT skeleton of the shape (Fig. 12).

4 A Measure of Chord Strength

We introduce a valuation on the chords of a discrete shape's CAT. The chords of the CAT are edges of the CDT of the shape that are shared by two triangles. Let the angles opposite a chord c in its two flanking triangles be θ and ϕ . We then define the strength of c by

$$S(c) = 1 - (\theta + \varphi) / \pi. \tag{1}$$

Thus S is a valuation on the set of all chords of a shape, with values in the half-open interval $[0, 1)$. This is because the empty circle condition of the CDT ensures that the sum of the angles across from a chord does not exceed π radians. We will refer to this valuation as the chord strength. S takes the value 0 on chords that are flanked by co-circular triangles. This observation is based upon an elementary fact of Euclidean geometry that the opposite angles of a cyclic quadrilateral add up to π radians. Hence, the smaller the strength of a chord, the closer its flanking triangles are to being co-circular, and vice versa. We are now ready to suppress the chords of low strength in the construction of the CAT skeleton. In what follows, we will address shapes without holes to keep the discussion simple. The techniques described can easily be extended to shapes with holes as well.




Fig. 9. CAT skeleton showing oscillations in tapered regions of a rectangular shape




Fig. 10. CAT skeleton showing degree three branching for a star-like shape




Fig. 11. Corner detail of Fig 9. showing CDT and skeleton oscillations




Fig. 12. Rectification of skeleton (dotted line) by suppressing weak chord AB




Fig. 13. Skeleton in Fig. 11 rectified by suppressing weak chords




Fig. 14. Rectified CAT skeleton of shape in Fig. 10, with a degree 8 branch point




Fig. 15. Schematic of CDT in the neighborhood of chords at junction of part with shape (left) and shape necks (center and right)

5 Chord Strength Profile Extrema and Shape Decomposition

The directed chords of a shape can be ordered by traversing the boundary of the shape and recording in sequence the chords encountered. For instance, at boundary point \mathbf{p} , a directed chord \mathbf{pq} joining \mathbf{p} to another boundary point \mathbf{q} will be recorded, as will the chord \mathbf{qp} on arrival at point \mathbf{q} . Thus each chord will be visited twice in traversing the shape contour. We call the plot of chord strengths versus chord numbers the chord strength profile (CSP). The CSP is a circular function in that the first and last chords in this enumeration are neighbors. Using the CSP, we select chords whose strengths are strictly greater than that of at least one neighbor and greater than or equal to the strengths of both neighbors. More precisely, the chord c_i is selected if and only if

$$[(S(c_i) > S(c_{i-1})) \ \& \ (S(c_i) \geq S(c_{i+1}))] \ \vee \ [(S(c_i) \geq S(c_{i-1})) \ \& \ (S(c_i) > S(c_{i+1}))] \quad (2)$$

i.e., the selected chords' strengths are at least *one-sided* local maxima of the CSP. Only these selected chords will be allowed to play a role in determining the skeleton of the shape.

5.1 Polygonal Decomposition of Shapes

Next, we construct a triangle grouping graph whose vertices are the triangles of the CDT of the shape, and with an edge between two vertices corresponding to adjacent triangles if and only if their common edge is a suppressed chord. A connected component analysis via a depth-first-search traversal of this grouping graph yields a polygonal decomposition of the shape, with each polygon comprised of pairwise adjacent, approximately co-circular triangles belonging to the same connected component. We refer to such polygons decomposing a shape as Delaunay polygons. As in the case of the CDT of a shape, these polygons can be classified into terminal, sleeve, and junction Delaunay polygons depending on whether they have one, two, or

more chords, respectively, among their bounding edges. Again, as in the case of triangles, skeletal segments are constructed in each Delaunay polygon to obtain a skeleton of the shape. The midpoints of chords of a sleeve Delaunay polygon are joined together to yield a sleeve skeletal segment. In the case of a junction Delaunay polygon, we define its *barycenter* as the weighted average of the midpoints of its chords, where the weight of each chord's midpoint is the (normalized) length of the chord. The midpoints of the chords are then joined to this barycenter to yield a skeletal segment of the junction Delaunay polygon. The collection of all the skeletal segments, with their adjacencies inherited from the adjacencies of the parent Delaunay polygons, form a connected rectified CAT skeleton of the shape. The suppression of weak chords remedies not only the skeletal oscillations (Figs. 9, 11, 12, 13), but also the purely degree-3 branch points (Figs. 10, 14) forced by the CDT in the original CAT skeleton prior to rectification.

5.2 Skeleton Pruning

A pruning criterion for excising skeleton segments corresponding to insignificant shape features is easily specified. For each chord of a junction Delaunay polygon, the length of the shape boundary arc subtended by it (and not including the polygon,) as a fraction of total shape boundary length is computed. In our experiments we have found that this fraction provides a satisfactory measure of the saliency of the part bounded by the chord and the boundary arc that is also efficient to compute. If this fraction falls below a predetermined threshold, the chord is an external boundary segment of a pruned shape. Accordingly, a new barycenter of the junction polygon is computed with the remaining chords taken into consideration if the remaining chords number greater than two. Otherwise, the junction Delaunay polygon is demoted to a sleeve or terminal Delaunay polygon and appropriate skeleton segments are constructed anew. Other pruning criteria such as excising parts based on the ratio of distance of furthest point of part from subtending chord to chord length, or alternately, to the distance of chord from barycenter, may also be implemented at additional computational cost.

5.3 Visually Salient Shape Decomposition

The selection criterion for a shape's chords, specified in condition (2), may be applied repeatedly to the CSPs of successive generations of selected chords, yielding chords whose strengths are higher order maxima in the original CSP (Fig. 18). These strong chords correspond to cuts of the shape into visually salient parts. They typically occur at the intersection of limbs with the shapes and necks of the shape where there is a narrowing of the shape girth. The reason for this is intuitively captured in Fig. 15 where the structure of the shape around the chord at part junctions and necks forces the sum of the angles opposite the chord in the flanking triangles of the shape's CDT to be smaller, and hence the chord to have greater strength than in other places of the shape. The strength of a chord weakens with increase in its length for the same boundary geometry in the vicinity of its endpoints. Interestingly, this behavior is consistent with the *shortcut rule* of parsing shapes using shortest cuts, proposed in

[11] based on experiments in human vision. However, in order to ensure the saliency of the parts resulting from cuts induced by the maxima of the CSP, we bias the CSP by multiplying the strength of each chord with the ratio of the (smaller) arclength of the shape subtended at the chord to the length of the chord. This has the effect of enhancing the strength of chords that subtend visually salient parts of the shape, while diminishing the strengths of chords that subtend minor or noisy protuberances of the shape. It is important to note that we do not process the boundaries of shapes by removing noise or smoothing.

Thus, the CSP provides not only a means of rectifying the CAT skeleton (Fig. 16), but also a means of decomposing shapes into visually meaningful parts (Fig. 17). A well known work in this area is that of K. Siddiqi et al [10]. Their approach to shape decomposition is also motivated by considerations of visual saliency and yields good results. In contrast, our approach addresses obtaining good shape skeletons as well as good shape decompositions in a unified manner by proposing a single criterion for solving both problems. The CSP maps the two-dimensional problem of shape analysis to the analysis of a one-dimensional function's extrema. This opens up a host of well known techniques for analyzing 1-D signals such as wavelet transforms to obtain hierarchical decompositions of shapes. We note that the shape decompositions obtained by our method closely resembles the outcomes of experiments in shape decomposition by M. Singh, G. Seyranian, and D. Hoffman [11] using human subjects. We believe that the property of strong chords of CDTs to yield visually meaningful decomposition of shapes has the potential to be developed into a useful and elegant tool in investigating and understanding shapes.




Fig. 16. Shapes and their rectified CAT skeletons based on 1st order maxima of their CSPs

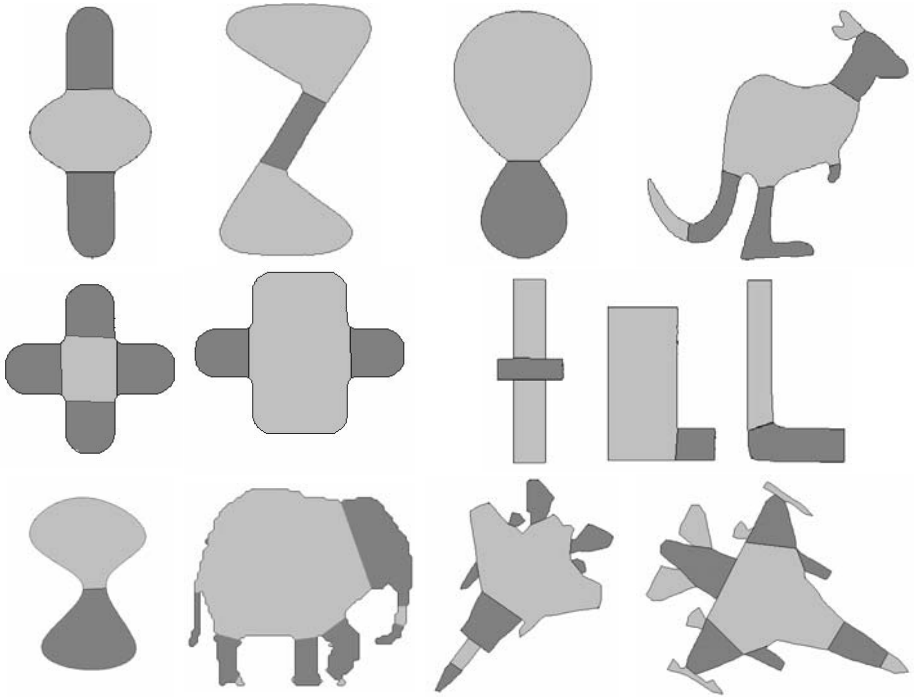


Fig. 17. Shapes and their decomposition based on higher order (order 4) maxima of their CSP. Adjacent parts are shown in alternating shades




Fig. 18. The normalized chord strength profile vs. chord visit number of the top left shape in Fig. 17. indicating 4th order maxima (circled). The four tallest peaks correspond to the two horizontal cuts (two peaks for each cut, as each chord is visited twice) decomposing the shape into three salient parts. The shorter circled peaks correspond to cuts across small noisy protuberances

6 Shapes with Holes

In the case of shapes with g holes, one can define $g+1$ CSPs, one for each hole contour, and one for the outer shape contour. The CSPs may then be analyzed individually for n^{th} order maxima just as in the case of a single CSP of a shape without holes. In the case of a CSP of a shape with holes, a chord will be encountered only once in the CSP if it joins two different contours and twice if it joins two points on the same contour. A chord is retained as a strong chord of the shape if and only if it is a strong chord (n^{th} order maximum) of at least one CSP. Based on the strong chords obtained thus, one can construct rectified CAT skeletons and obtain visually salient decompositions of shapes with holes as well.

Alternatively, one can formulate the problem of finding strong chords of a discrete shape on the dual graph of its CDT, where the vertices of the graph represent triangles, and with an edge between each pair of vertices corresponding to adjacent triangles in the CDT. Each graph edge can be weighted with the strength of the chord that separates the triangles whose representative vertices the edge connects. Edge weight maxima can be computed along minimal cycles corresponding to hole contours and the cycle corresponding to the outer contour. One can then group vertices that are path connected by edges whose weights correspond to non-maximal chord strengths to obtain a new graph whose vertices correspond to Delaunay polygons described before. Iterations of this process yield higher order maxima and hierarchical shape decompositions.

7 Computational Issues

The constrained Delaunay triangulation of a planar shape specified by n points runs in $o(n \log n)$ time. The construction of the CSP and detection of maxima is linear in the number of edges which are $2n-3+3g$ in number, where $g < n/3$ is the number of holes. The connected component analysis on the grouping graph is linear in the number of vertices (representing the triangles of the shape's CDT which are $n-2+2g$ in number). Shape pruning is, again, linear in the number of edges. Finally, skeletonization is also linear in the number of edges. Thus the overall efficiency of our shape decomposition and skeletonization scheme is high and amenable to real-time applications.

8 Conclusion

In this paper we have demonstrated a property of the chords of constrained Delaunay triangulations of 2-D shapes that induces a hierarchy of visually salient decompositions by defining a valuation on the chords. This valuation, which we call chord strength, along with the ordering induced by the shape boundary on the chords, maps the two-dimensional problem of shape decomposition into a one-dimensional problem of analyzing a function's extrema. We have briefly demonstrated how one can obtain rectified shape skeletons as well as visually meaningful shape decompositions using successive selection of strong chords. The

search for chords that best decompose a shape using the extrema of the chord strength profile function introduced here can be improved upon significantly and is part of our ongoing work in shape analysis. The unified approach provided by our method to both skeletonization and decomposition of shapes is the key contribution of this paper.

In three dimensions, one can define a valuation characterizing approximate co-sphericity of two face-adjacent Delaunay tetrahedra by defining the strength of their interface as the ratio of the distance between their circumcenters to the mean of their circumsphere diameters. However, the 3-D case is more complex. For instance, two adjacent co-spherical tetrahedra need not form a convex polytope, and therefore their interface may be an important structural partition. Thus, additional conditions will be required in order to extend our ideas presented here to three-dimensional shapes.

Acknowledgement

This work has been fully supported by the U. S. DOE under contract No. W-7405-ENG-36 through an LDRD ER (#20030162) research grant. We would like to thank the reviewers for their helpful comments.

Reference

1. H. Blum.: A Transformation for Extracting New Descriptors of Shape. Symp. Models for Speech and Visual Form Weiant Whaten-Dunn (Ed) MIT Press (1967)
2. J. E. Goodman, J. O'Rourke. (Eds.): Handbook of Discrete and Computational Geometry, CRC Press (1997)
3. J. R. Parker: Algorithms for Image Processing and Computer Vision, John Wiley & Sons (1997)
4. R.L. Ogniewicz: Skeleton-space: A multiscale shape description combining region and boundary information," in Proc. IEEE CVPR, Seattle, WA, June 1994
5. D. Attali, A. Montanvert.: Computing and Simplifying 2D and 3D Continuous Skeletons, Computer Vision and Image Understanding, Vol. 67, Sept. 1997
6. L. Prasad: Morphological Analysis of Shapes, CNLS Newsletter, No. 139, July 1997, LALP-97-010-139, Center for Nonlinear Studies, Los Alamos National Laboratory
7. L. Prasad, R. L. Rao & G. Zweig: Skeletonization of shapes using Delaunay triangulations, Fifth SIAM Conference on Geometric Design, Nashville, TN, Nov. 1997
8. L. Prasad, R. L. Rao.: A Geometric Transform for Shape Feature Extraction, Proc. SPIE, vol. 4117, Vision Geometry IX (2000)
9. L. Prasad, A. Skourikhine, B. Schlei.: Feature-based Syntactic and Metric Shape Recognition, Proc. SPIE, vol. 4117, Vision Geometry IX (2000)
10. K. Siddiqi, B. B. Kimia: Parts of Visual Form: Computational Aspects, IEEE Trans. on PAMI, vol. 17, No. 3, March 1995
11. M. Singh, G. Seyranian, D. Hoffman: Parsing Silhouettes: The Short-cut Rule, Perception and Psychophysics, 61, 636-660 (1999)
12. T. Igarashi, S. Matsuoka, H. Tanaka.: Teddy: A Sketching Interface for 3D Freeform Design, ACM SIGGRAPH'99, 409-416, Los Angeles, CA (1999)

13. P. Felzenszwalb.: Representation and detection of deformable shapes. Proc. CVPR, vol.1, pp. 102—108 (2003)
14. J. Arvo, K. Novins : Smart Text: A Synthesis of Recognition and Morphing *AAAI Spring Symposium on Smart Graphics*, Stanford, California, pp 140-147, March 2000
15. S. Yamakawa, K. Shimada.: Quad-Layer: Layered Quadrilateral Meshing of Narrow Two-Dimensional Domains by Bubble Packing and Chordal Axis Transform, ASME/DETC/DAC (2001)