Published in final edited form as:
Int J Comput Vis. 2004 December ; 60(3): 203–224. doi:10.1023/B:VISI.0000036835.28674.d0.

Area-Based Medial Axis of Planar Curves

Marc Niethammer,
Department of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0250, USA

Santiago Betelu,
Department of Mathematics, University of North Texas, Denton, TX 76203, USA

Guillermo Sapiro,
Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455, USA

Allen Tannenbaum, and
Department of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0250, USA

Peter J. Giblin
Department of Mathematics, University of Liverpool, Liverpool L69 3BX, UK

Guillermo Sapiro: guille@ece.umn.edu

Abstract

A new definition of affine invariant medial axis of planar closed curves is introduced. A point belongs to the affine medial axis if and only if it is equidistant from at least two points of the curve, with the distance being a minimum and given by the *areas* between the curve and its corresponding chords. The medial axis is robust, eliminating the need for curve denoising. In a dynamical interpretation of this affine medial axis, the medial axis points are the affine shock positions of the affine erosion of the curve. We propose a simple method to compute the medial axis and give examples. We also demonstrate how to use this method to detect affine skew symmetry in real images.

Keywords

medial axis; affine invariant; symmetry; area; shape; pattern recognition

1. Introduction

Object recognition is an essential task in image processing, and the *skeleton* or *medial axis* is a shape descriptor that is often used for this task.

Thus, the computation of skeletons and symmetry sets of planar shapes is a subject that received a great deal of attention from the mathematical (see Bruce et al., 1985; Bruce and Giblin, 1992 and references therein), computational geometry (Preparata and Shamos, 1990), biological vision (Kovács and Julesz, 1994; Lee et al., 1995; Leyton, 1992), and computer vision communities (see for example Ogniewicz, 1993; Serra, 1982 and references therein) since the original work by Blum (1967, 1973).

In the classical Euclidean case, *the symmetry set* of a planar curve (or of the boundary of a planar shape) is defined as the set of points equidistant from at least two different points on the given curve, providing the distances are local extrema (see Fig. 1). The fact that the distances are local extrema means that the symmetry set point lies on the intersection of the (Euclidean) normals at the corresponding curve points. This leads to an equivalent definition of the symmetry set, as the closure of centers of bitangent circles.

The importance of points which are the centers of bitangent circles was realized by Motzkin (1983a, 1983b), who used this concept to characterize the convexity and the non-convexity of a closed set S in the Euclidean plane. Particularly, he showed that if there is no point $p \in S^C$ (the complement of the set S), which is the center of a bitangent circle to S, S has to be convex. Klee (1949) and Phelps (1957) extended Motzkin's ideas to non Euclidean metrics.

Blum defined the symmetry set in a different, but equivalent way (the grassfire transform): If a fire is started at the boundary of the shape, and it propagates with uniform speed, the symmetry set is given by the set of points where two or more fire fronts collide (*shocks*) (Giblin and Kimia, 1999; Smoller, 1983). In other words, the symmetry set is composed of the shock points of the curve, which evolve according to Huygens' principle. The trace of this symmetry set is called the medial axis. Augmenting the medial axis by arrival time information leads to the medial axis function, which allows for perfect reconstruction of the boundary. Based on Blum's work, Calabi and Hartnett (1968) studied the extrinsic skeleton mathematically.

Medial axes describe the shapes of objects. To extract these shape descriptions automatically, computer algorithms have to be devised. This is not a trivial task. The transition from the continuous to the discrete (sampled) domain is problematic (Mott-Smith, 1970): The computations are sensitive to object orientation, and object topology is not necessarily preserved. Furthermore, most algorithms require the object's boundary as an input, and thus a segmentation step, which is in itself demanding. A segmentation does generally not represent the exact shape of the object to be segmented. Instead, the obtained boundary might be slightly off and/or noisy. Medial axis transforms can be very sensitive to boundary details (Katz and Pizer, 2002). Smoothing of the boundary can change the structure of a medial axis and can even create new structures (Shaked, 1996; August et al., 1999). Algorithms that compute a medial axis based on a segmentation of an object's boundary need thus to be robust with respect to boundary uncertainties. Many algorithms have been devised to tackle these problems. Following Kimia and Leymarie (2001) they can be classified into methods using thinning, the distance transform, boundary modeling, the Voronoi diagram, bisector computations and trimming, and surface evolution based approaches (see Kimia and Leymarie, 2001 for references for each class). The lines between these methods are not sharp, combinations exist.

For distance function based medial axis transforms, perfect reconstructability is usually given, i.e., given the medial axis function, the original shape of the object can be reconstructed. This is a crucial property for data compression applications. However, for certain purposes (e.g., in pattern recognition) it is enough to have a rough idea of the shape, and the reconstructability requirement can be relaxed. In case of noisy boundaries, relaxing this requirement yields the flexibility to remove spurious branches that are most likely noise induced. Rotational invariance is a desirable property for a medial axis transform. If the distance function is chosen to be Euclidean, this can be achieved. However, a good (and efficiently computable) approximation for the Euclidean distance function is then needed on the discrete grid. Meyer (1990) introduces digital Euclidean skeletons. The skeleton is given by connecting crest points (the set of points that only have neighboring points with lower distance transform values and so-called multiple points or pairs). Whereas in the continuous

case the crest points are identical to points which are the centers of maximal disks, this is not the case for the Euclidean distance transform on a discrete grid. Connectivity is not automatically guaranteed by translating continuous concepts to the discrete domain. An algorithm similar to Meyer's, that provably maintains connectivity, is suggested by Niblack et al. (1990).

Lam et al. (1992) and Lee et al. (1993) give comprehensive surveys on methods for thinning and skeletonization respectively. They discuss iterative and non-iterative thinning algorithms. Iterative methods can be subdivided into sequential and parallel methods. For sequential algorithms, the pixels of an image are visited in a prespecified order, e.g., by raster scanning or contour following. Pixels that are determined not to be part of the medial axis are deleted. This is not always natural and results can depend on the order the deletion is performed in. Most algorithms use arbitrary ordering. The resulting medial axis might not be truly medial. Ranwez and Soille (2002) thus proposes a thinning scheme that is order independent. In the sequential framework it is easy to ensure homotopy (i.e., ensuring that the connectedness of an objects is preserved). This is generally not the case for parallel algorithms, which on the other hand bear the promise of extremely fast implementations.

Many medial axis algorithms, e.g., algorithms based on ridge following, do not easily extend to higher dimensions (which would for example be crucial for countless medical imaging applications) or are relatively prone to boundary noise. One possible remedy is the use of anchor point based homotopic thinning. The idea is to perform iterative thinning, but not to remove a point whenever it is an anchor point or it would break the topology. A possible choice for the anchor points are the centers of maximal balls (or disks in 2D) obtained from a distance transform (here thresholding can be performed based on the level of desired detail). Pudney (1998) performs distance-ordered homotopic thinning using the centers of maximal balls. Svensson et al. (1999) propose a rotationally invariant algorithm based on similar ideas, which can be extended to any dimension, with any metric and connectivity. Ranwez and Soille (2002) propose an order independent scheme.

Anchor points have previously been used, for example by Leymarie and Levine (1992). They use an active contour based approach to determine the medial axis of an object. The active contour moves on a Euclidean distance function. Curvature extrema on the boundary are used as anchor points for the active contour. However, while the method guarantees the connectedness of the resulting medial axis, it relies on a well segmented boundary to compute the curvature maxima; for noisy boundaries pruning might be required.

Ogniewicz and Ilg (1992) compute the medial axis of an object by means of the the Voronoi diagram of its boundary elements. This approach guarantees the preservation of connectivity and uses a truly Euclidean metric, but any vertex introduces an additional branch in the medial axis. Pruning is thus necessary, but can nicely be incorporated into this framework by means of the "residual function", which measures how "deep" a medial axis part is inside the object. Pruning is then performed by thresholding based on the residual value of an edge of the medial axis. Hierarchical clustering leads to a multi-resolution representation of the medial axis and allows for the extraction of its main axis.

Most of the work on medial axes deals with binary images (i.e. objects with known boundaries). Furthermore, rotational invariance is desired (Lee et al., 1993). Extensions exist for gray valued/intensity images (see for example the survey by Serra (1982, 1988) or Ranwez and Soille (2002), Pizer et al. (1998),Wang et al. (1982), and Verwer et al. (1993)). More recently, affine invariant medial axis algorithms have been formulated (Giblin and Sapiro, 1998a, 1998b).

Inspired by Giblin and Sapiro (1998a, 1998b) and especially by the groundbreaking work of Moisan (1998), we define in this paper an analogous, affine invariant, symmetry set, the *affine area symmetry set* (AASS). Instead of using Euclidean distances to the curve, we define a new distance based on the areas enclosed between the curve and its chords. We define the symmetry set as the closure of the locus of points equidistant from at least two different points on the given curve, provided that the distances are local minima.

Further, we obtain a dynamical interpretation of the AASS, where, instead of using the classical Huygens' principle, we use a notion of *affine erosion* and show that the affine shocks of the corresponding evolution occur at the AASS.

As we will demonstrate, this definition based on areas makes the symmetry set remarkably noise-resistant, because the area between the curve and a chord "averages out" the noise. This property makes the method very useful to compute symmetry sets of real images without the need of denoising. In addition, as this is an area-based computation, the result is affine invariant. That implies, that if we process the image of a planar object, the skeleton will be independent of the angle between the camera that captures the image and the scene, provided the camera is far enough from the object. For these reasons, we believe that this symmetry set, in addition to having theoretical interest, has the strong potential of becoming a useful tool in invariant object recognition.

The proposed algorithm can be extended to the three dimensional case (Betelu et al., 2001a). Then, instead of areas enclosed by the curve and its chords, the volume enclosed between the surface of the object and its chords is used to define the distance measure.

We now summarize the contents of this paper. In Section 2, we define our key concepts of *affine distance, affine area symmetry set*, and the *affine skeleton*. In Section 3, we give our dynamical interpretation of the affine skeleton as the set of *affine shock points* derived via an affine notion of erosion. In Section 4, we give a robust algorithm for the computation of the affine skeleton. This algorithm is applied to some examples, both synthetic and real, in Section 5, where we also indicate its robustness in the presence of noise. Section 6 draws conclusions and proposes some directions for future research.

2. Affine Area Symmetry Set and Affine Invariant Skeletons

In this section, we formally introduce the key concepts of *affine distance, affine medial axis*, and *affine area symmetry set*. We begin with the following definition:

Definition 2.1. A *special affine* transformation in the plane ($\mathbb{R}^2$) is defined as

$$\tilde{X}=AX+B, \quad (1)$$

where $X \in \mathbf{R}^2$ is a vector, $A \in \mathrm{SL}_2(\mathbf{R}^2)$ (the group of invertible real 2×2 matrices with determinant equal to 1) is the affine matrix, and $B \in \mathbf{R}^2$ is a translation vector.

Transformations of the type (1) form a real algebraic group $\mathscr{A}$ called the *group of special affine motions* or the *group of area preserving motions*. The key property of the affine group is that the area enclosed by a curve is invariant to the group action.

In this work we deal with symmetry sets, which are affine invariant in the sense that if a curve C is affine transformed with Eq. (1), then its symmetry set is also transformed according to Eq. (1). Affine invariant symmetry sets are not new. In 1998 Giblin and Sapiro (1998a, 1998b) introduced them and proposed the definition of the *affine distance symmetry set* (ADSS) for a planar curve $C(s)$. In their definition, they used affine geometry and

defined distances in terms of the affine invariant tangent of $C(s)$, which involves second derivatives of the curve with respect to an arbitrary parameter. Later, they defined the ADSS analogously to the Euclidean case (see Fig. 1). In Betelu et al. (2001b) an efficient implementation for the computation of these sets was discussed. However, there is a fundamental technical problem with this definition: in curves extracted from real images, noise is always present; the higher order derivatives needed to compute the ADSS oscillate in a very wild fashion, unless considerable smoothing is performed. Giblin and Sapiro proposed a second definition, the *affine envelope symmetry set* (Giblin and Sapiro, 1998a, 1998b), which still requires derivatives and thus suffers from the same computational problem.[1]

As we will show with the new definition presented below, *we do not compute derivatives at all*. Consequently, the robustness of the computation is considerably higher, and smoothing is not required.

For simplicity, we shall always deal with simple closed curves $C(s) : [0, 1] \rightarrow \mathbf{R}^2$ with a countable number of discontinuities on the derivative $C'(s)$. Even though classical affine geometry is only defined for convex curves, our framework extends this to the non-convex case. We start by defining the building block of our *affine area symmetry set* (AASS), the *affine distance* (inspired by Moisan, 1998):

Definition 2.2. The *affine distance* between a generic point X and a point of the curve $C(s)$ is defined by the area between the curve and the chord that joins $C(s)$ and X:

$$d(X, s)=\frac{1}{2}\int_{C(s)}^{C(s')}(C - X) \times dC, \quad (2)$$

where s is the curve parameter, × yields the z component of the cross product of two vectors,[2] and the points $C(s)$ and $C(s')$ define the chord that contains X and that has exactly two contact points with the curve, as shown in Fig. 2.

This distance is invariant under the special affine transformation Eq. (1), and it is independent of the parameterization of the curve. For a simple convex curve, the function $d(X, s)$ is always defined for interior points, but for concave curves the function $d(X, s)$ may be undefined for some values of s in [0, 1], as sketched in Fig. 2. When the point is exterior to the curve (as the point Y in Fig. 1), the distance may not be defined. For example, if the curve is a circle, the distance from an exterior point X is not defined, because we cannot find a chord that intersects the curve exactly twice and contains the point X.

We can now define our affine symmetry set:

Definition 2.3. $X \in \mathbf{R}^2$ is a point in the *affine area symmetry set* of $C(s)$ (AASS) if and only if there exist two different points s_1, s_2, which define two different chords that contain X and have equal area,

$$d(X, s_1)=d(X, s_2), \quad (3)$$

[1]Even though many definitions for Euclidean skeletons are based on derivatives, it is not always necessary to use derivatives for the actual computations—sometimes algorithms can be devised that circumvent the use of derivatives, while only slightly changing the definition of the skeleton. If and how this is possible in the affine case is not clear at present.

[2]Our distance is the integral of the distance used in Giblin and Sapiro (1998a, 1998b). This is *not* the standard definition of affine distance.

provided that $d(X, s_1)$ and $d(X, s_2)$ are defined and that they are local minima with respect to s (see Fig. 4). The AASS is then the closure of the locus of these points X.

This definition is analogous to the Euclidean case and the ADSS in Giblin and Sapiro (1998a, 1998b).

Commonly in shape analysis, a subset of the Euclidean symmetry set is used, and it is denoted as *skeleton* or *medial axis*. For example, the symmetry set of an ellipse contains segments on both symmetry axes of the ellipse, while the skeleton or medial axis corresponds only to the segment on the principal axis. In the Euclidean case, one possible way to simplify the symmetry set into a skeleton, inspired by original work of Giblin and colleagues (Bruce and Giblin, 1992; Wright et al., 1995), is to require that the distance to the curve is a global minimum. The affine definition is analogous:

Definition 2.4. The *affine skeleton* or *affine medial axis* is the subset of the AASS where $d(X, s_1)$ and $d(X, s_2)$ are global (not just local) minima with respect to s.

This means, that X belongs to the affine medial axis if there are two different chords which contain X and that have the same minimum area. The medial axis is invariant with respect to Eq. (1), because of the invariance of the areas. However, as the definition is based on the comparison of areas, it is also invariant with respect to more general transformations up to change of scale $\det(A) > 0$.

There are curves for which the medial axis may be computed exactly. For example, it is easy to verify that the medial axis of a circle is its center. We can make an affine transformation to the circle and transform it into an ellipse. By virtue of the affine invariance of our definitions, we conclude that the medial axis of an ellipse is its center (the midpoint of the line that connects the two foci of the ellipse), too.

Another important example is the triangle. In affine geometry, all triangles may be generated by affine-transforming an equilateral triangle. The medial axis of an equilateral triangle are the segments connecting the middle of the bases and the center of the figure (as can be verified by simple area computations). As a consequence, the medial axis of an arbitrary triangle are the segments connecting the middle of the sides with the center of gravity (where all medians cross).

An important property concerns curves with straight sections. If the curve contains two equal straight parallel non-collinear segments, the AASS will contain an identical segment parallel and equidistant to the former two. For instance, the AASS of a rhombus will contain its medians.

3. Dynamical Interpretation: Affine Erosion of a Planar Curve

We now define a notion of *affine erosion* of a curve and show its connection with the AASS. Our definition is analogous (but different) to that used by Moisan (1998). Let $C(s) : [0, 1] \rightarrow \mathbf{R}^2$ be a simple closed curve. Let $(C(s), C(s'))$ be a chord of $C(s)$, as shown in Fig. 3. As before, this chord intersects the curve exactly twice.

The difference of this definition, with respect to the definition given in Moisan (1998), is that in the Moisan set-up the chord is allowed to cross the curve outside the interval (s, s'). For our purposes, we opted for a definition which gives a unique chord for a given parameter s, when this chord exists, and which also separates the curve into two disjoint parts. The connected closed set enclosed by the chord and the curve is the *chord set*. Its area is denoted by A, so that

$$A=\frac{1}{2}\int_{C(s)}^{C(s')}(C-X)\times dC, \quad (4)$$

where X is any point of the chord. This area may have any sign, depending on the orientation of the curve. If the area is positive, we call the corresponding chord a *positive* chord and the corresponding chord set a *positive* chord set.

Definition 3.1. The *minimum distance* from a point X to the curve C is defined by

$$f(X)=\inf(d(X,s), s\in D), \quad (5)$$

where D is the domain of $d(X, s)$ for a fixed value of X.

For interior points X, the minimum distance $f(X)$ is always defined, because in a simple closed curve, we can always draw at least one chord that contains X. For exterior points it may be undefined, as for instance, in a circle.

Moisan defines the affine erosion of the curve C by A as the set of the points of the interior of C, which do not belong to any positive chord set with area less than A. Here, we define the affine erosion in terms of our affine distance:

Definition 3.2. The *affine erosion* $E(C, A)$ of the shape enclosed by a curve C, by the area A, is the set of the points X of the interior of C that satisfy

$$E(C,A)=\{X\in\mathbf{R}^2: f(X)\geq A\geq 0\}. \quad (6)$$

Roughly speaking, it is the area bounded by the envelope of all the possible chords of area $A \geq 0$ (see Fig. 3). We also define the *eroded curve* $C(A)$ as the "boundary" of the affine erosion

$$C(A)=\{X\in\mathbf{R}^2: f(X)=A\}. \quad (7)$$

If we consider the area to be a time parameter, $t = A$, the distance $f(X)$ represents the time that the eroded curve $C(A)$ takes to reach the point X, when traveling with constant affine velocity. Initially, when $A = 0$, we have the initial curve. At later times ($A > 0$), the curve $C(A)$ will be contained inside the original curve.

This statement is a particular case of a more general property analogous to the *maximum principle*, that is, if $0 < A_1 < A_2$, the affine erosion defined by A_2 will be included in the affine erosion defined by A_1. To prove that, we just have to see that if a point $X \in E(C, A_2)$ by the definition Eq. (6), $f(X) \geq A_2 > A_1$ and $X \in E(C, A_1)$. Then, $E(C, A_2) \subset E(C, A_1)$.

However, because of our "two-point contact" chord definition, if we take two different initial curves, $C_2(s)$ bounded by $C_1(s)$, we cannot guarantee that $C_2(A)$ will be bounded by $C_1(A)$ for $A > 0$.

Remark. Our chord definition has an interesting hydrostatic analogy: if we make a cylindrical container of length L, whose walls have the shape of a closed curve $C(s)$, we introduce a volume AL of an incompressible liquid inside it, and we put the container with its longitudinal axis at 90 degrees with respect to the gravity, then the liquid surface of the liquid will coincide with a chord of the curve $C(s)$. By setting the container in different positions and requiring that the liquid fills a connected region, the liquid surface will take the shape of different chords of the same area, in virtue of the conservation of the mass of

the liquid. The envelope of these liquid surfaces will be the *eroded curve*. For more details and examples we refer the reader to Betelu et al. (2001a).

There is a fundamental relationship between the affine erosion of a curve and the skeleton, namely, an affine shock point X is a skeleton point. This allows to define the medial axis as Blum did for the Euclidean case.

Definition 3.3. An *affine shock point* X is a point of the eroded curve $C(A)$ where two different chords $(C(s_1), C(s_1'))$ and $(C(s_2), C(s_2'))$ of equal area A intersect.

Clearly, the distance from X to these two points are equal, $d(X, s_1) = d(X, s_2)$. On the other hand, the distance is a global minimum at s_1, because as X belongs to $C(A)$ (see Eq. (7)), $f(X) = \inf(d(X, s), s \in D) = A$. Thus, an affine shock point X is a point of the medial axis (see Fig. 3).

3.1. Basic Properties of the Affine Area Distance and Symmetry Set

In this section, we shall present some theoretical results concerning the AASS, mostly without proof. Further results and details will appear elsewhere.[3] We shall show that the AASS has connections with our previously defined *affine envelope symmetry set* or AESS (Giblin and Sapiro, 1998b),[4] as well as with the *affine distance symmetry set* (ADSS) mentioned before. We will see that a point that is an element of the AASS will lie at the midpoint of its associated chords.

As mentioned above, it is not clear to us how far the area definition can be extended. That is, it is not yet clear whether we can define a *smooth* family of functions

$$d: C \times U \to \mathbf{R},$$

associating to (s, X) the 'area $d(s, X)$ of the sector of C determined by $C(s)$ and X', for all points X inside some reasonably large set U in the plane $\mathbf{R}^2$. We have seen that when C is convex, the function is well-defined and smooth for all X inside the curve C. We shall assume this below.

We have

$$2d(s, X) = \int_s^{t(s)} [C(s) - X, C'(s)] ds = \int_s^{t(s)} F(s, X) ds,$$

for any regular parameterization of C, where [,] means the determinant of the two vectors inside the square brackets, $'$ means $\frac{d}{ds}$ and $t(s)$ is the parameter value of the other point of intersection of the chord through X and $C(s)$ with the curve. Using standard formulae for differentiation of integrals,

$$2d_s := 2\frac{\partial d}{\partial s} = F(t(s), X)t'(s) - F(s).$$

[3]Some of the results below have also been obtained by Holtom (2000).

[4]This set is basically defined as the closure of the center of conics having three-point contact with at least two points on the curve.

We evaluate $t'(s)$ by using the fact that $C(s)$, X and $C(t(s))$ are collinear:

$$C(t(s)) - X = \lambda(C(s) - X) \text{for a scalar} \lambda, \text{i.e.} [C(s) - X, C(t(s)) - X] = 0.$$

Differentiating the last equation with respect to s we obtain

$$t'(s) = [C(t(s)) - X, C'(s)] / [C(s) - X, C'(t(s))],$$

and from this it follows quickly that, provided the chord is not tangent to C at $C(s)$ or $C(t(s))$, $d_s = 0$ if and only if $\lambda = \pm 1$. But $\lambda = 1$ means $C(s)$ and $C(t(s))$ coincide, so for us the interesting solution is $\lambda = -1$. This means that X is the mid-point of the segment (this result has also been obtained in Moisan (1998) and Brady and Asada (1984)): *The area function d has a stationary point at (s, X), i.e.,* $\partial d / \partial s = 0$, *if and only if X is the midpoint of the segment from C(s) to C(t(s))*.

One immediate consequence of this is: *The envelope of the chords cutting off a fixed area from C (the affine eroded set) is also the locus of the midpoints of these chords*.

Further calculations along the same lines show that: *The first two derivatives of d with respect to s vanish at (s, X) if and only if X is the midpoint of the chord, and also the tangents to C at the endpoints of the chord are parallel*.

In mathematical language this means that the 'bifurcation set' of the family d is the set of points X, which are the midpoints of chords of C at the ends of which the tangents to C are parallel. This set is also the envelope of lines parallel to such parallel tangent pairs and halfway between them. It has been called the *midpoint parallel tangent locus* (MPTL) by Holtom (2000). Various facts are known about the MPTL, for example, it has an odd number of cusps. These cusps coincide in position with certain cusps of the AESS (Giblin and Sapiro, 2000). The cusps in question occur precisely when the point X is the center of a conic having 3-point contact with C at two points where the tangents to C are parallel: *The first three derivatives of d with respect to s vanish at (s, X) if and only if X is the midpoint of the chord, the tangents to C at the endpoints are parallel, and there exists a conic with center X having 3-point contact with C at these points*.

The *full bifurcation set* of the family d consists of those points X for which (i) d has a degenerate stationary point ($d_s = d_{ss} = 0$) for some s, or else (ii) there are two distinct s_1, s_2, and d has an ordinary stationary point ($d_s = 0$) at each one and the same value there: $d(s_1, X) = d(s_2, X)$. The latter is precisely the AASS as defined above. Mathematically, the AASS and the MPTL 'go together' in the same way that the classical symmetry set and evolute, or the ADSS and the affine evolute go together: in each case the pair makes up a single mathematical entity called a *full bifurcation set*. A good deal is known about the structure of such sets, including the structure of full bifurcation sets arising from *families* of curves (see Bruce and Giblin, 1986). For instance, the symmetry set has endpoints in the cusps of the evolute, and in the same way the AASS has endpoints in cusps of the MPTL.[5]

When X lies on the AASS, there are two chords through X, with X the midpoint of each chord. The corresponding areas defined by d are equal. From the midpoint conditions alone, it follows that the four endpoints of the chords form a *parallelogram* (see Fig. 4). The tangent to the AASS at X is in fact parallel to two sides of this parallelogram.

[5]The ADSS as defined in Giblin and Sapiro (1998b) also has endpoints, and these are in the cusps of the *affine evolute*. The endpoints of the AASS are, by contrast, in the cusps of the *midpoint parallel tangent locus*.

Consider the AASS of a triangle, as in Section 2, where it was noted that the affine medial axis comes from the center of a side and stops at the centroid of the triangle. The 'full' AASS, allowing for non-absolute minima of d, stops at the point *halfway* up the median, as can be verified by an elementary calculation with areas. Presumably this is a highly degenerate version of a cusp on the AASS. With the triangle, the two branches of the cusp are overlaid on each other.

We mention finally one curious phenomenon connected with the AASS. Given a point $C(s_1)$, there will generally be an *area-bisecting chord* through this point. That is, the two areas on either side of the chord and within C are equal. In that case, let X be the midpoint of the chord and let $C(t(s_1))$ be the other end of the chord. Let $s_2 = t(s_1)$: then $t(s_2) = s_1$ by construction, and X satisfies the conditions to be a point of the AASS. That is, the midpoints of all area-bisecting chords automatically appear in the AASS (see Fig. 5). These points are in some sense anomalous: the 'genuine' AASS consists of the other points satisfying the defining condition.

4. Robust Numerical Implementation for Discrete Curves

Inspired by Betelu et al. (2001b), we propose the following algorithm to compute the affine medial axis.

1. *Discretization of the curve*. Discretize $C(s) = (x(s), y(s))$ with two vectors for the points $C_k = (x_k, y_k)$, with $1 \le k \le M$ (see Fig. 6).
2. *Discretization of the rectangular domain*. Discretize the domain that contains the curve, of dimensions $L_x \times L_y$, with a uniform grid of $N_x \times N_y$ points. Each point X_{ij} of the grid will have coordinates $X_{ij} = (i\Delta x, j\Delta y)$, where $\Delta x = L_x/N_x$, $\Delta y = L_y/N_y$ (see Fig. 6) and $0 \le i \le N_x$, $0 \le j \le N_y$.

Now, for each point X_{ij} of the grid we perform the following steps:

a. *Compute the chord areas*. With Eq. (2), compute the areas between X_{ij} and each point in the curve C_k:

$$d(C_k, X_{ij}) = \frac{1}{2}\int_{C_k}^{C_k^*} (C - X_{ij}) \times dC \quad (8)$$

for $k = 1, \ldots, M$. The integral is computed by approximating the curve with a polygon that interpolates the points C_k and by computing the point C_k^* as the intersection between the curve and the line joining C_k and X_{ij}. As mentioned before, the distance is not defined if the chord does not touch exactly two points on the curve. The points have to be labeled with a logical vector E_k, indicating whether or not the distance is defined at C_k. Here, we detect these singular points just by scanning around the curve and counting the crossings between the line that contains C_k and X_{ij}. Then, we store the areas in a vector $d_k = d(C_k, X_{ij})$, with $k = 1, \ldots, M$.

b. *Search for local minima of the chord areas*, approximated by the local minima of the set d_k, with $k = 1, \ldots, M$. When d_{k-1}, d_k and d_{k+1} are defined, the local minimum condition is

$$d_{k-1} \ge d_k \le d_{k+1}. \quad (9)$$

When d_k and $d_{k\mp1}$ are defined, but $d_{k\pm1}$ is not defined, the condition is simply $d_k \le d_{k\mp1}$. Now, for each point of the grid X_{ij} we shall have after this step the set of l

local minima distances $(d_1^*, d_2^*, \ldots, d_l^*)$, corresponding to the points $(C_1^*, C_2^*, \ldots, C_l^*)$. All these quantities are functions of X_{ij}.

c. *Approximate AASS computation*. Compute the differences

$$\begin{aligned} D(X_{ij}) &= d_p^* - d_q^* \\ p &= 1, \ldots, l, \quad q = 1, \ldots, l, r \neq q, \end{aligned} \tag{10}$$

p and q represent the indices of the local minimum distances of *different* chords. If this difference is smaller in magnitude than a given tolerance ε, we can consider X_{ij} to be an approximate point of the area symmetry set (see Fig. 6). Then, as a first approximation, *add to the AASS all the points X_{ij} that satisfy*

$$|d_p^* - d_q^*| < \varepsilon. \tag{11}$$

The tolerance ε has to be of the order of the variation of the distance difference along a cell in the discretized domain, i.e.,

$$\varepsilon \approx \max\left(\left|\frac{\partial D}{\partial x}\right|\Delta x, \left|\frac{\partial D}{\partial y}\right|\Delta y\right). \tag{12}$$

The partial derivatives are taken with respect to the components of X_{ij}.

There is not a simple general formula for this expression. However, by using the distance definition Eq. (2) and by restricting ourselves to *regular* local minima, which satisfy $d_s(s^*) = 0$, we can demonstrate that

$$\nabla d_p^*(X_{ij}) = \frac{1}{2}(-\Delta y_p^*, \Delta x_p^*) \tag{13}$$

where $(\Delta x_p^*, \Delta y_p^*)$ are the components of the chords $C_p^{*'} - C_p^*$ corresponding to local minimum area. This formula is not general, since we may have non-regular minima, as for example, points where $d(X, s)$ has a discontinuity with respect to s. However, we still use this expression, because we only need an order of magnitude for the tolerance. Then, we define

$$\varepsilon = \max(|\Delta y_p^* - \Delta y_q^*|\Delta x, |\Delta x_p^* - \Delta x_q^*|\Delta y). \tag{14}$$

d. *Focusing of the AASS*. At this point, the medial axis is quite crude, and the branches have a spatial error of the order of the discretization of the domain $(\Delta x, \Delta y)$. We can compute, with negligible computational cost, the (small) correction vector $\Delta X_{ij} = (u, \upsilon)$ to the position X_{ij} that makes the difference of distances exactly equal to zero at $X_{ij} + \Delta X_{ij}$ (see Fig. 6):

$$d(X_{ij} + \Delta X_{ij}, C_p^*) = d(X_{ij} + \Delta X_{ij}, C_q^*)$$

At first order, we must solve

$$D(X_{ij}) + \Delta X_{ij} \cdot \nabla D(X_{ij}) = 0 \tag{15}$$

with ΔX_{ij} parallel to the gradient of D. We get

$$u = \frac{(\Delta y_p^* - \Delta y_q^*)D}{2G^2}, \quad (16)$$

$$v = \frac{-(\Delta x_p^* - \Delta x_q^*)D}{2G^2}, \quad (17)$$

$$4G^2 = (\Delta y_p^* - \Delta y_q^*)^2 + (\Delta x_p^* - \Delta x_q^*)^2. \quad (18)$$

After this step, the AASS for a discrete image is computed.

e. *Pruning*. If one of the distances d_p^* or d_q^* is not a positive global minimum: discard the corresponding point. In this way we obtain the affine medial axis. We also have to discard points which originated very close to each other on the curve, in such a way that they are effectively indistinguishable. Here, we discard points for which the areas of the triangles $A(C_p, X_{ij}, C_q)$ are smaller than the areas of the triangles defined by the discretization of the curve $A(C_p, X_{ij}, C_{p+1}) + A(C_p', X_{ij}, C_{p+1}')$ and the area defined by the discretization of the domain $\Delta x L_Y + \Delta y L_X$, as indicated in Fig. 7.

5. Examples

Let us start with the simplest nontrivial example: the triangle. In Fig. 8(a) we show a triangle and its medial axis. The exact medial axis should reach the border. The gap is due to the discretization errors of the numerical scheme. Better discretizations reduce this gap. In Fig. 8(b) we corrupted the shape by adding to each one of the 170 points of the curve a uniform random perturbation of absolute amplitude 0.01. The medial axis is barely perturbed by the noise (this has been obtained without denoising). This is not a self-evident property for medial axis algorithms. See for example Fig. 9, which shows a Euclidean medial axis obtained by topological thinning. The slight perturbation shown in Fig. 9(b) results in a completely new branch of the medial axis compared to the unperturbed case in Fig. 9(a). This is in direct contrast to the effect the slight boundary perturbation has on the affine medial axis. Figure 10 shows the medial axis computations of the unperturbed triangle in its left column and of the perturbed triangle in the right column. Specifically, Fig. 10(a) shows the gray level medial axis of the unperturbed triangle, Fig. 10(c) its thresholded version. Figure 10(b) and (d) show the corresponding figures for the perturbed triangle. The gray level medial axis shows all calculated medial axis points corresponding to global maxima. The gray value of a point corresponds to how close the point is to its discretization level. Specifically, the gray level is chosen based on the ratio of the dashed triangle in Fig. 7 to the area defined by the discretization of the curve. Darker values correspond to larger ratios. The thresholded results correspond to a thresholding of this ratio at 1 (this is exactly the described “pruning” strategy). Strikingly, the boundary perturbation of the triangle does not influence its affine medial axis. No additional branches are created. The medial axis is robust with respect to this perturbation.

Naturally the question arises, if this is always a desirable property. Is information grossly neglected as a trade-off for robustness? This is not the case as is exemplified by the example of a discretized circle, as shown in Fig. 12. Clearly, in the unperturbed and the perturbed cases, the center of the circle is the most significant point of the medial axis (as indicated by its darkness in Fig. 12(a) and (b)). Due to the boundary discretization, we cannot expect to get isolated points at the centers of the circles after thresholding (unless we adapt our

thresholding scheme). This is sensible, since the boundary discretization might as well represent the correct shape of the object, i.e., the algorithm does not try to introduce any additional assumptions regarding boundary smoothness. However, in the case of the discretized unperturbed circle, the detected medial axis (after thresholding) is confined to the center of the circle, whereas it shows the additional symmetry introduced by the perturbation in the perturbed case. We see, that the same perturbation that did not cause any changes to the medial axis for the triangle example might very well change the medial axis if the perturbation adds new symmetries. Noise will generally be random. For a very regular object, like a circle, this can pose a problem. Figure 12(e) and (f) show the gray level affine medial axis and the thresholded affine medial axis of a circle with noisy boundary respectively. Clearly, the center of the circle is identifiable as the main symmetry point in Fig. 12(e). However, since the circle is such a regular object, it is very likely that a perturbed boundary element will form a symmetry with another perturbed boundary element, thus introducing weak, but in this case spurious, symmetry points.

The medial axis may be used to detect symmetries, a topic that has been the subject of extensive research in the computer vision community (e.g., Brady and Asada, 1984; Brooks, 1981; Cham and Cipolla, 1995; Ponce, 1989; van Gool et al., 1985; van Gool et al., 1996). In particular, affine skeletons may be used to detect skew symmetries: numerical experiments show that if the skeleton contains a straight branch, a portion of the curve has skew symmetry with respect to this line.[6] This is illustrated now. In Fig. 13(a) we show the original figure with its corresponding skeleton, while in Fig. 13(b) we show the figure affine transformed by the matrix

$$A = \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix}.$$

In Fig. 13(c) we corrupted the shape by adding to each point of the discrete curve a random number of amplitude 0.025. The corresponding skeleton remains almost unchanged. In particular (as expected since our medial axis is affine invariant) the linear medial axis branch stays linear under this transformation. If we select a local minimum instead of a global minimum of the distance, we obtain the AASS, shown in Fig. 14. Figure 11 shows the skewed symmetry axes for the triangular example and a couple of skewed symmetry points. As in Fig. 13(a) we see that the skewed symmetry is represented by a straight line in the affine medial axis.

Figure 15 demonstrates the dynamic interpretation of the affine skeleton: the successive curves of the affine erosion $C(A)$ (thin lines) form affine shocks at skeleton points (thick line). Here we plotted the contour lines of $f(X)$ (see Eq. (5)).

Now, we show how to compute skeletons from real data: For an image showing two tools (Fig. 16), an image showing a tennis racket (Fig. 21), and an image showing a slice through a human heart (Fig. 19). First we need to obtain the points of the curve C which define the shape. If the points are to be extracted from a digital image with a good contrast, they may be extracted by thresholding the image in binary values and then getting the boundary with a boundary-following algorithm (Jain et al., 1995). This procedure was performed with the shape on the right in Fig. 16 and with the tennis racket of Fig. 21. When the border of the shape is more complex or fuzzier, as for the shapes on the left in Fig. 16 and in Fig. 19,

[6] Although the *Euclidean* skeleton of a symmetric shape contains a straight line, this is not necessarily true anymore after the shape is affine transformed, obtaining skew symmetry.

more sophisticated techniques can be used. Here, we extracted the contour with the active contour algorithm as formulated in Caselles et al. (1997), Kichenassamy et al. (1996).[7]

The resulting medial axes are shown in Figs. 17–21. Figure 18 shows the four strongest lines of the computed medial axis of the tools image, as determined by means of a Hough transform. The main parts of the medial axes are highlighted nicely (i.e. they are represented by dark pixels in the gray value Figs. 17, 20(a), and 22) and are connected. Figure 20(c) shows the Euclidean skeleton for the heart image 19. Figure 20(b) and (d) show the gray level affine medial axis and the Euclidean skeleton for the same heart image with a slightly perturbed boundary respectively. While both, the gray level medial axis and the Euclidean skeleton, change under the boundary perturbation, the gray level medial axis seems to change less than the Euclidean skeleton in this case.

The medial axes decompose into several large connected pieces for the sample images 16 and 21 after thresholding. However, some isolated pixels and some short connected pieces remain in all three examples. Since small segments or even individual pixels might very well indicate a significant symmetry of the image (just think about the circle example), it is not advisable to simply suppress these objects. Our current strategy is to do pruning based on thresholding a discretization measure (see the algorithm in Section 4). In this case the significance of every point on the gray level medial axis is evaluated individually. It is not clear at this point if this is the optimal pruning strategy—most likely it is not. Thus, further research should try to develop alternative criteria for point significance testing to facilitate pruning, possibly including more global information.

6. Conclusions

In this paper, we introduced a new definition of affine medial axis and a robust method to compute it. The definition based on areas makes the medial axis remarkably insensitive to noise, and thus, useful for processing real images.

There are still theoretical problems that have to be solved in order to build a consistent theory to support our definition:

- **a.** The medial axis in the Euclidean case may be found by detecting the shocks in the solutions of the Hamilton-Jacobi equation (this is equivalent to the Huygens principle in Blum's method). We do not know at this point, whether there is a differential equation which would allow us to compute our affine medial axis in an analogous way.
- **b.** The properties of the area-distances $d(X, s)$ in relation with geometric characteristic of the curves, such as normals, curvatures, etc., are yet unknown.
- **c.** The extension of the definition to multiply connected curves would be of paramount importance in practical applications.
- **d.** The assessment of the importance of individual parts of the medial axis to describe the object shape would be useful to further refine the obtained medial axis.

These are some of the open questions that occupy us and others in the area of symmetry sets. Nevertheless, at this point, we believe that we have presented a reliable method for affine

[7]Specifically, we minimize the cost functional $L=\int_0^1 g(\|\nabla I\|)\|C_p\|dp$, where C describes the curve, $p \in [0, 1]$ is the curve's parameterization, and g is some potential function on the image I. Compared to the classical "snake" approach (Kass et al., 1988), this formulation results in a curve evolution equation based on intrinsic geometric properties of the curve being evolved.

invariant medial axis computations, which hopefully will lead to some interesting practical results in studying affine invariant symmetries of objects in computer vision.

Acknowledgments

This work was supported in part by grants from the National Science Foundation ECS-0137412, the Air Force Office of Scientific Research AF/F49620-94-1-00S8DEF, AF/F49620-94-1-0461, AF49620-01-1-0017, the Army Research Office DAAD19.02-1-0378, the Office of Naval Research ONR-N00014-97-1-0509 and Young Investigator Award to GS, the Presidential Early Career Awards for Scientists and Engineers (PECASE) to GS, the National Science Foundation CAREER Award to GS, by MURI, and NIH (through NAC, Brigham and Women's Hospital).

References

August, J.; Tannenbaum, A.; Zucker, SW. Proceedings of the International Conference on Computer Vision. Vol. vol. 1. IEEE; 1999. On the evolution of the skeleton; p. 315-322.

Betelu, S.; Sapiro, G.; Tannenbaum, A. Proceedings of the International Conference on Computer Vision. Vol. vol. 2. IEEE; 2001a. Affine invariant erosion of 3D shapes; p. 174-180.

Betelu S, Sapiro G, Tannenbaum A, Giblin P. On the computation of affine skeletons of planar curves and the detection of skew symmetry. Pattern Recognition. 2001b; 34:943–952.

Blum, H. Models for the Perception of Speech and Visual Form. MIT Press; 1967. p. 362-380.chapter A Transformation for Extracting New Descriptors of Shape

Blum H. Biological shape and visual science. Journal of Theoretical Biology. 1973; 38:205–287. [PubMed: 4689997]

Brady JM, Asada H. Smoothed local symmetries and their implementation. International Journal of Robotics Research. 1984; 3:36–61.

Brooks RA. Symbolic reasoning among 3D models and 2D images. Artifical Intelligence. 1981; 17:285–348.

Bruce JW, Giblin PJ. Growth, motion and one-parameter families of symmetry sets. Proceedings of the Royal Society of Edinburgh. 1986; vol. 104A:179–204.

Bruce, JW.; Giblin, PJ. Curves and Singularities. 2nd edition. Cambridge University Press; 1992.

Bruce JW, Giblin PJ, Gibson CG. Symmetry sets. Proceedings of the Royal Society of Edinburgh. 1985; vol. 101A:163–186.

Calabi L, Hartnett WE. Shape recognition, prarie fires, convex deficiencies and skeletons. The American Mathematical Monthly. 1968; 75(4):335–342.

Caselles V, Kimmel R, Sapiro G. On geodesic active contours. International Journal of Computer Vision. 1997; 22(1):61–79.

Cham TJ, Cipolla R. Symmetry detection trough local skewed symmetries. Image and Vision Computing. 1995; 13:439–450.

Giblin, PJ.; Kimia, BB. Proceedings of the International Conference on Computer Vision. IEEE; 1999. On the local form and transitions of symmetry sets, and medial axes, and shocks in 2D; p. 385-391.

Giblin, P.; Sapiro, G. Affine invariant symmetry sets and skew symmetry; Proceedings of the International Conference on Computer Vision; 1998a.

Giblin PJ, Sapiro G. Affine invariant distances, envelopes and symmetry sets. Geometriae Dedicata. 1998b; 71:237–261.

Giblin, PJ.; Sapiro, G. Real and Complex Singularities, vol. 412 of Research Notes in Mathematics, chapter Affine Versions of the Symmetry Set. Chapman and Hall/CRC; 2000. p. 173-187.

Van Gool, L.; Moons, T.; Proesmans, M. Mirror and point symmetry under perspective skewing; Proceedings of the Conference on Computer Vision and Pattern Recognition; 1996. p. 285-292.

Van Gool L, Moons T, Ungureanu D, Oosterlinck A. The characterization and detection of skewed symmetry. Computer Vision and Image Understanding. 1985; 61:138–150.

Holtom, P. PhD thesis. Liverpool University; 2000. Affine-invariant symmetry sets.

Jain, R.; Kasturi, RR.; Schunck, B. Machine Vision. McGraw Hill; 1995.

Klee VL Jr. A characterisation of convex sets. The American Mathematical Monthly. 1949; 56(4): 247–249.

Kass M, Witkin A, Terzopoulos D. Snakes: Active contour models. International Journal of Computer Vision. 1988:321–331.

Katz RA, Pizer SM. Untangling the Blum medial axis transform. International Journal of Computer Vision. 2002 (submitted).

Kichenassamy S, Kumar A, Olver P, Tannenbaum A, Yezzi A. Conformal curvature flows: from phase transitions to active vision. Archive for Rational Mechanics and Analysis. 1996; 134:275–301.

Kimia, BB.; Leymarie, FF. Proceedings of the International Conference on Image Processing. IEEE; 2001. Symmetry-based representations of 3D data; p. 581-584.

Kovács I, Julesz B. Perceptual sensitivity maps within globally defined visual shapes. Nature. 1994; 370:644–646. [PubMed: 8065449]

Lam L, Lee S-W, Suen CY. Thinning methodologies—a comprehensive survey. IEEE Transactions on Pattern Analysis and Machine Intelligence. 1992; 14(9):869–885.

Lee S-W, Lam L, Suen CY. A systematic evaluation of skeletonization algorithms. International Journal of Pattern Recognition and Artificial Intelligence. 1993; 7(5):1203–1225.

Lee TS, Mumford D, Zipser K, Schiller P. Neural correlates of boundary and medial axis representation in primate striate cortex. Investigative Opthalmology and Visual Science Supplement. 1995 (ARVO Abstract).

Leymarie F, Levine MD. Simulating the grassfire transform using an active contour model. IEEE Transactions on Pattern Analysis and Machine Intelligence. 1992; 14(1):56–75.

Leyton, M. Symmetry, Causality, Mind. MIT Press; 1992.

Meyer, F. Visual Communications and Image Processing ’90. Vol. vol. SPIE-1360, Society of Photo-Instrumentation Engineers; 1990. Digital Euclidean skeletons; p. 251-262.

Moisan L. Affine plane curve evolution: A fully consistent scheme. IEEE Transactions on Image Processing. 1998; 7:411–420. [PubMed: 18276261]

Mott-Smith, JC. Picture Processing and Psychopictorics, chapter Medial Axis Transformations. Academic Press; 1970. p. 267-278.

Motzkin, TS. Sur quelques propriétés charactéristiques des ensembles bornés non convexes. In: Motzkin, Theodore S., editor. Selected Papers. Birkhäuser; 1983a. p. 138-143.

Motzkin, TS. Sur quelques propriétés charactéristiques des ensembles convexes. In: Motzkin, Theodore S., editor. Selected Papers. Birkhäuser; 1983b. p. 138-143.

Niblack, CW.; Capson, DW.; Gibbons, PB. Proceedings of the International Conference on Pattern Recognition. Vol. vol. 1. IEEE; 1990. Generating skeletons and centerlines from the medial axis transform; p. 881-885.

Ogniewicz, R.; Ilg, M. Voronoi skeletons: Theory and applications; Proceedings of the Conference on Computer Vision and Pattern Recognition; 1992. p. 63-69.

Ogniewicz, RL. Discrete Voronoi Skeletons. Hartung-Gorre Verlag; 1993.

Phelps RR. Convex sets and nearest points. Proceedings of the American Mathematical Society. 1957; vol. 8:790–797.

Pizer SM, Eberly D, Fritsch DS. Zoom-invariant vision of figural shape: The mathematics of cores. Computer Vision and Image Understanding. 1998; 69(1):55–71.

Ponce, J. On characterizing ribbons and finding skewed symmetries; Proceedings of the International Conference on Robotics and Automation; 1989. p. 49-54.

Preparata, FP.; Shamos, MI. Computational Geometry. Texts and Monographs in Computer Science. Springer Verlag; 1990.

Pudney C. Distance-ordered homotopic thinning: A skeletonization algorithm for 3D digital images. Computer Vision and Image Understanding. 1998; 72(3):404–413.

Ranwez V, Soille P. Order independent homotopic thinning for binary and grey tone anchored skeletons. Pattern Recognition Letters. 2002; 23:687–702.

Serra, J. Image Analysis and Mathematical Morphology. Vol. vol. 1. Academic Press; 1982.

Serra, J. Image Analysis and Mathematical Morphology. Vol. vol. 2. Academic Press; 1988.

Shaked, D. PhD thesis. Technion; 1996. Symmetry, invariance, and evolution in planar shape analysis.

Smoller, J. Shock Waves and Reaction-Diffusion Equations. Springer Verlag; 1983.

Svensson S, Borgefors G, Nyström I. On reversible skeletonization using anchor-points from distance transforms. Journal of Visual Communication and Image Representation. 1999; 10:379–397.

Verwer BJH, van Vliet LJ, Verbeek PW. Binary and grey-value skeletons: Metrics and algorithms. International Journal of Pattern Recognition and Artifical Intelligence. 1993; 7(5):1287–1308.

Wang S, Rosenfeld A, Wu AY. A medial axis transformation for grayscale pictures. IEEE Transcations of Pattern Analysis and Machine Intelligence. 1982; 4(4):419–421.

Wright MW, Cipolla R, Giblin PJ. Skeletonization using an extended Euclidean distance transform. Image and Vision Computing. 1995; 13(5):367–375.

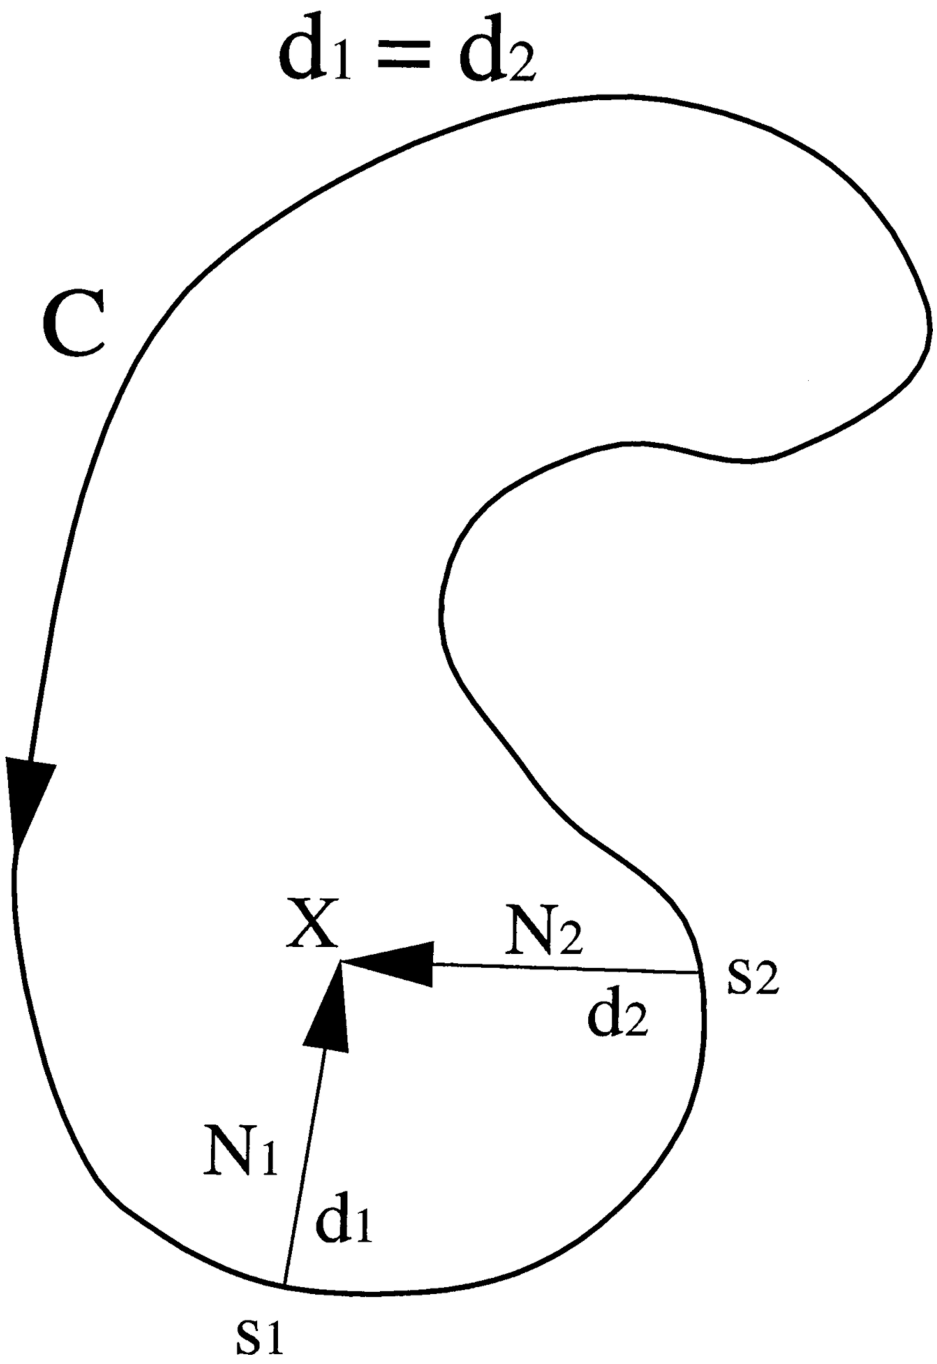

Figure 1.
Sketch of a Euclidean symmetry set. The distance from X to the curve is equal at points s_1 and s_2, and a point of the symmetry set X is contained in the direction of the normals.

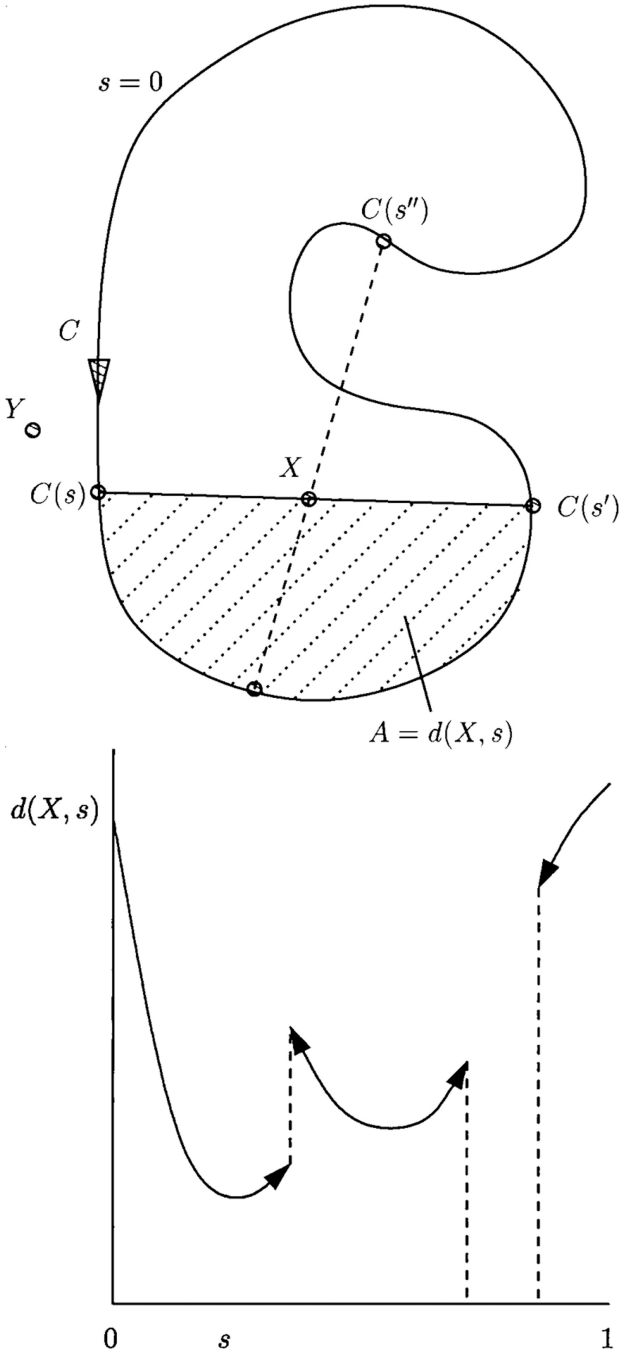

Figure 2.

We define the affine distance from X to the curve as the area between the curve C and the chord $(C(s), C(s'))$. The chord touches the curve exactly twice. For example, there is no "legal" chord which contains the pair of points $(C(s''), X)$. As a consequence, the function $d(X, s)$ may have discontinuities as shown in the last sketch.

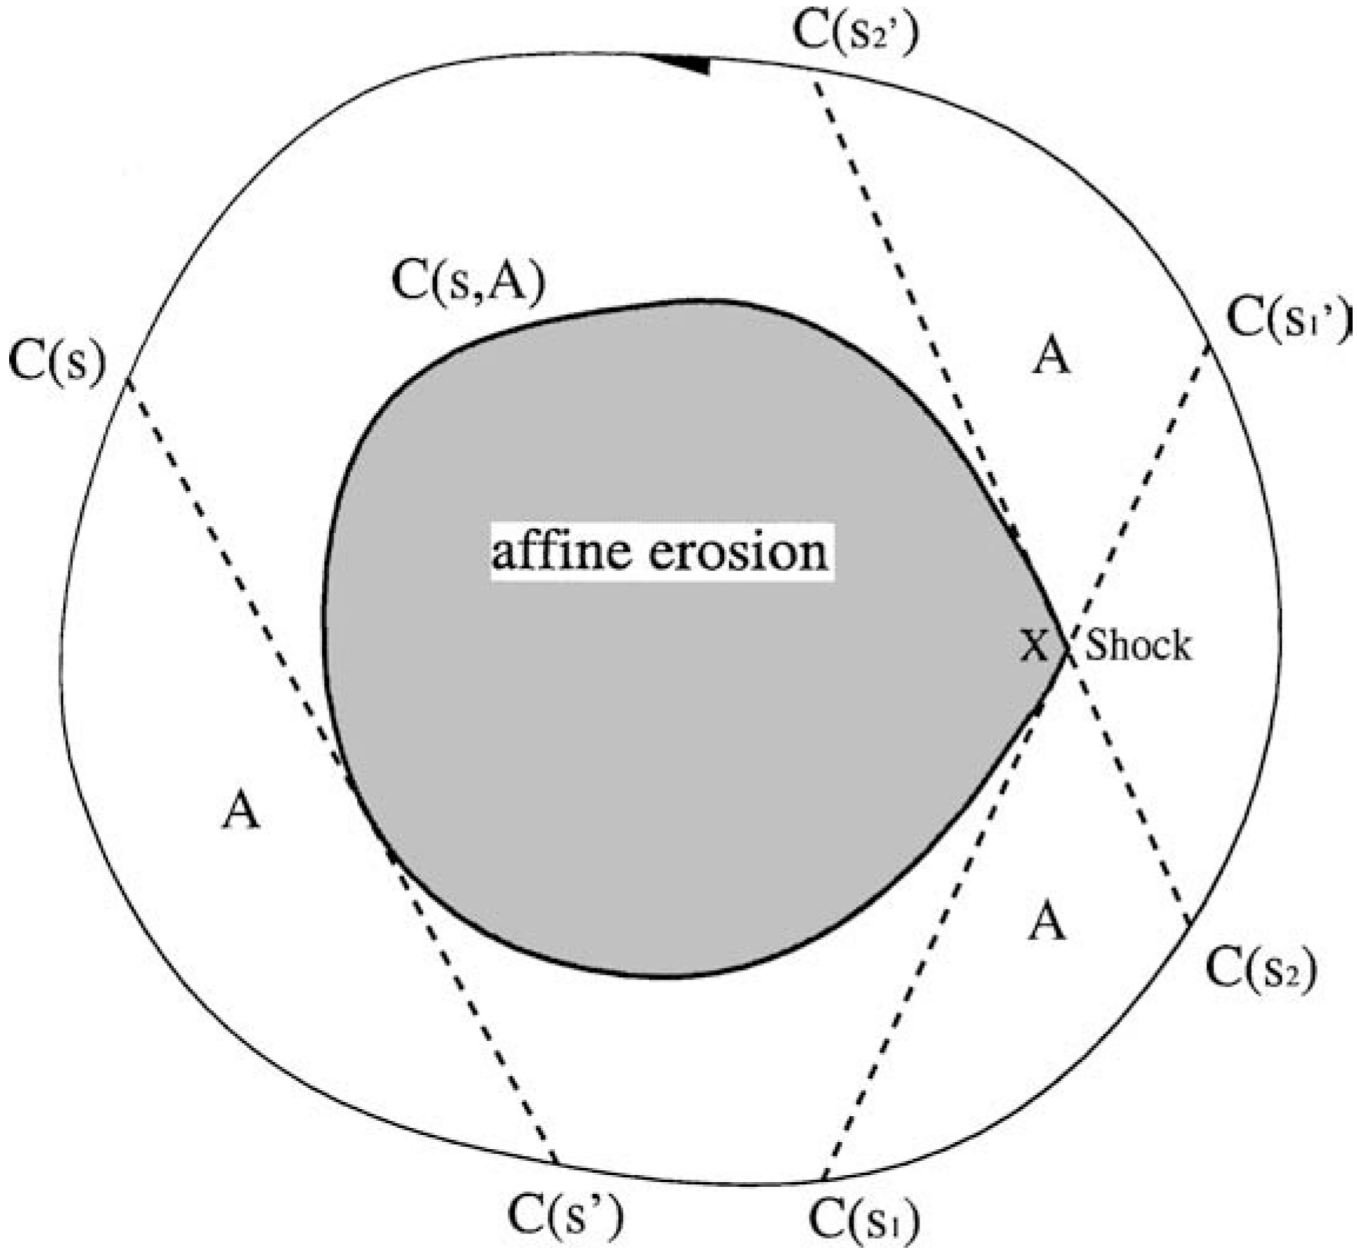

Figure 3.
Affine area based erosion of a curve. After a time $t = A$, the eroded curve is the envelope of all the chords that have the same area.

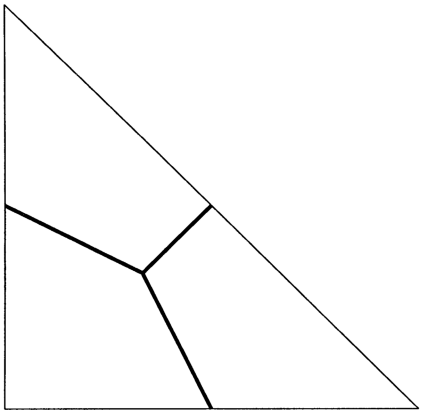

(a) Medial axis of an isosceles triangle.

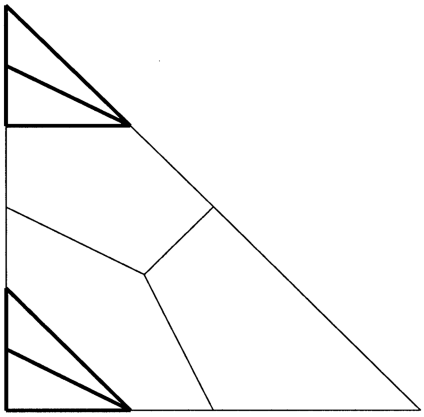

(b) The median lines of two small similar triangles are parallel to the median line of the large triangle.

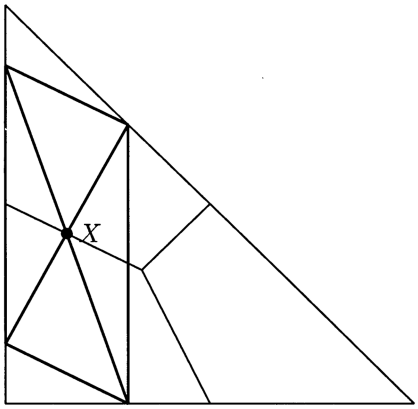

(c) The intersections of the medial lines of the small triangles define a parallelogram whose diagonals intersect on the medial axis of the large triangle and define two chords of which X is the midpoint.

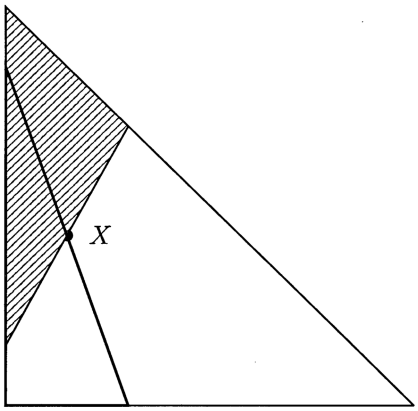

(d) The medial point X is given by the intersection of two chords which enclose an equal area.

Figure 4.

For X being a point on the AASS, X is the midpoint of its defining chords. These chords enclose equal areas and form a parallelogram. The tangent to the AASS at X is parallel to two sides of this parallelogram.

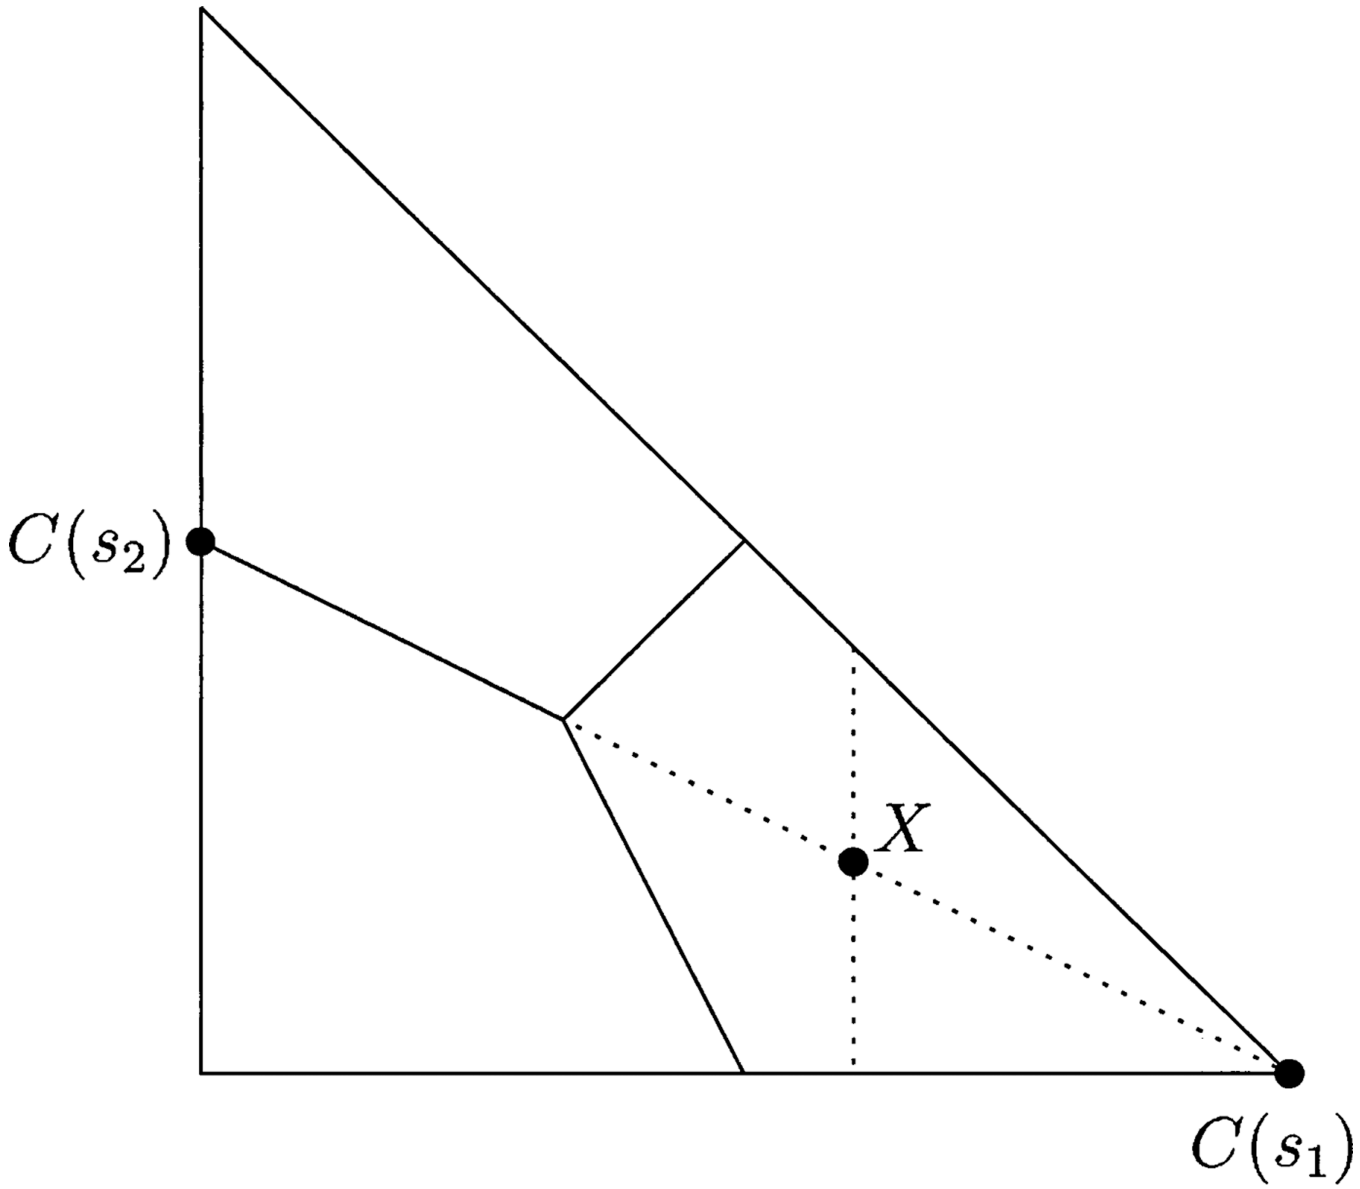

Figure 5.
X is an anomalous point of the AASS.

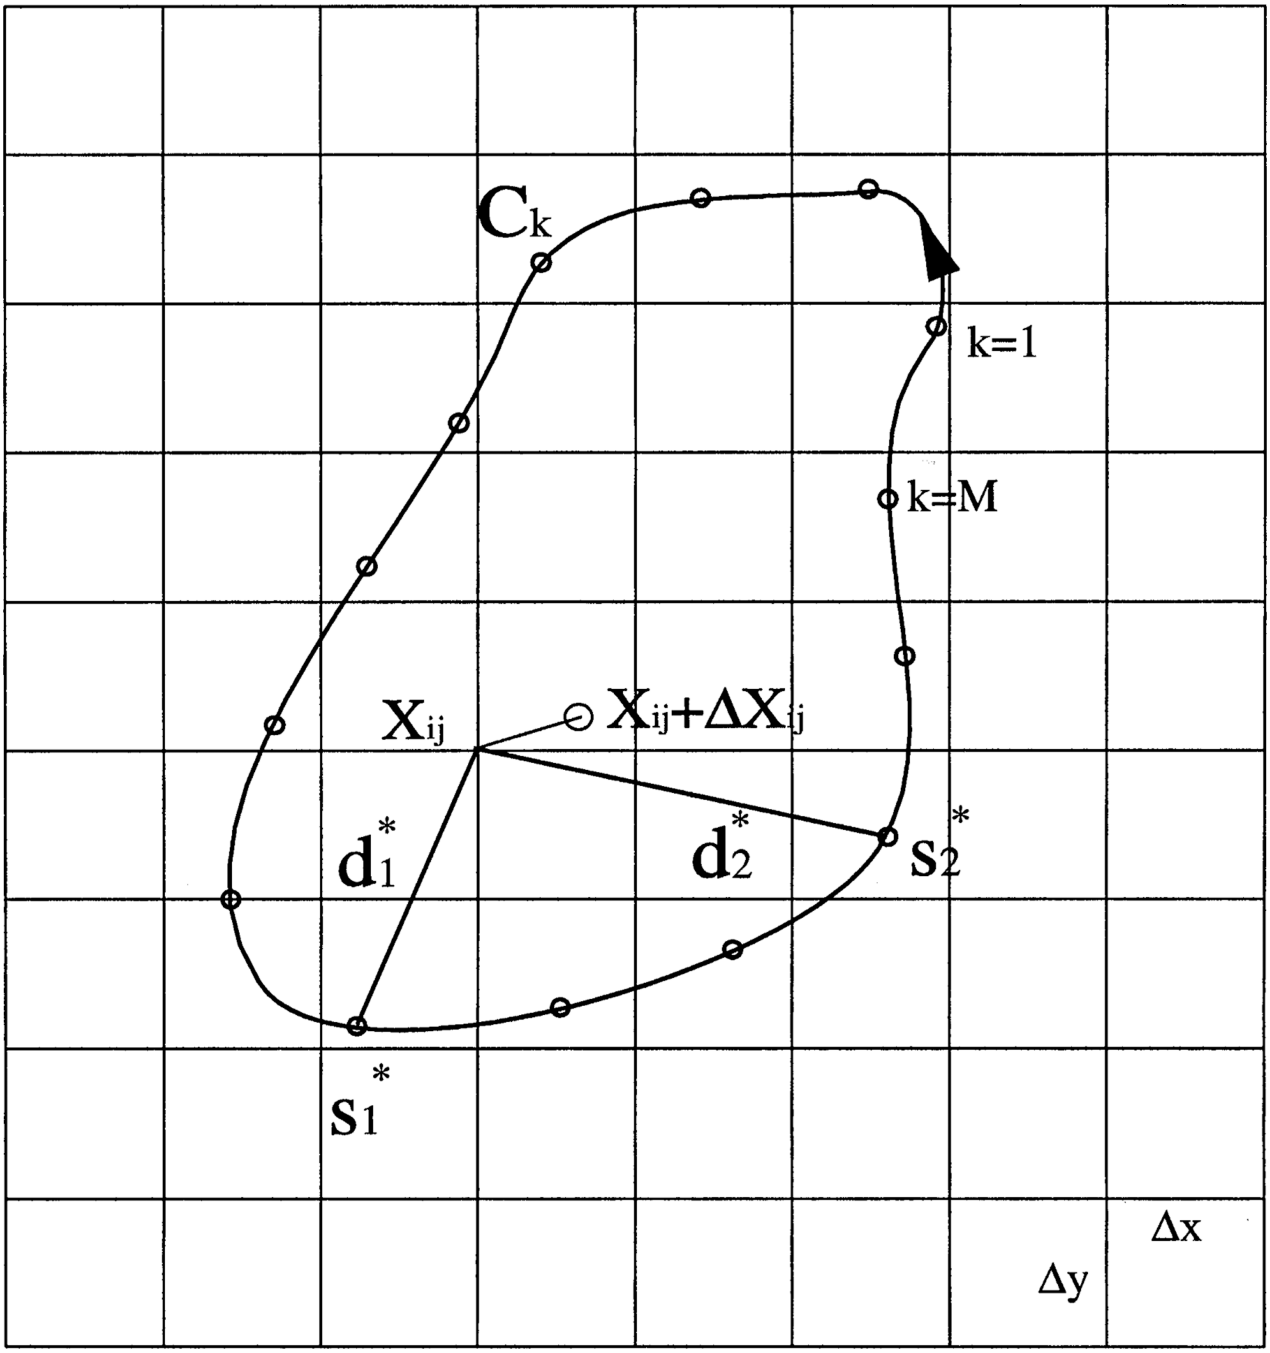

Figure 6.
Discretization of the curve and the domain.

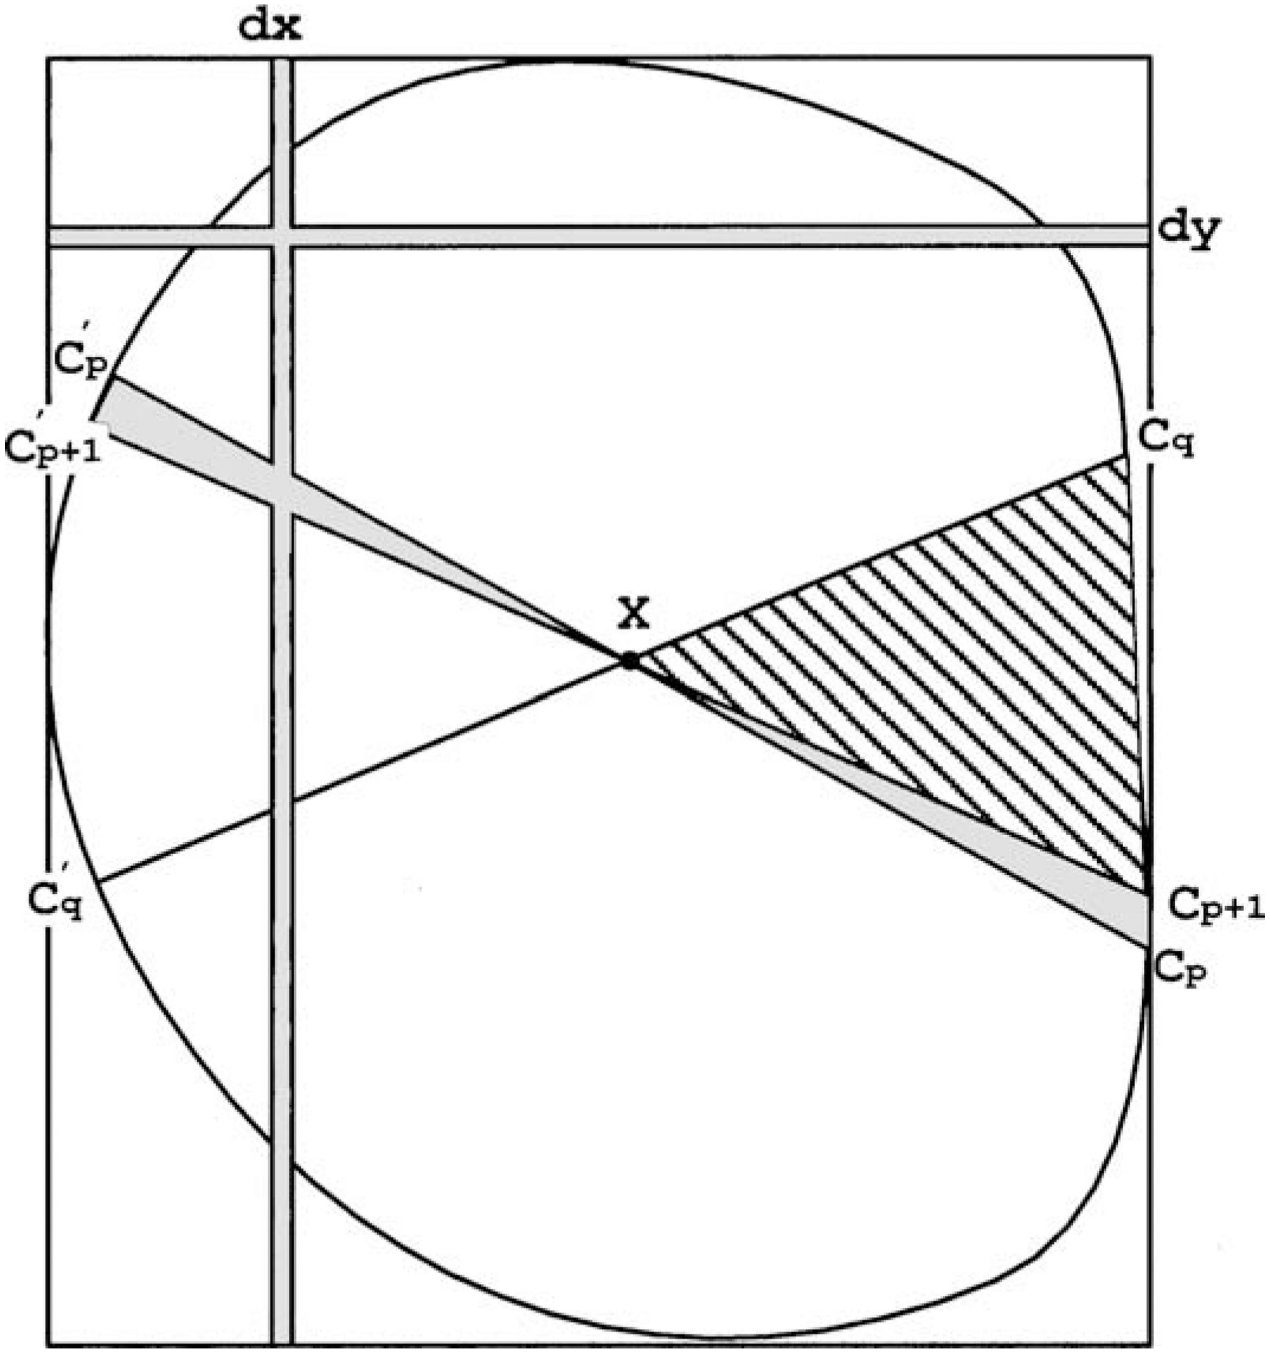

Figure 7.

We consider two points *p, q* "indistinguishable" if the area of the triangle $A(C_p, X, C_q)$ (dashed) is smaller than the area defined by the discretization of the curve $A(C_p, X, C_{p+1}) + A(C'_p, X, C'_{p+1})$, and the discretization of the domain (shaded regions).

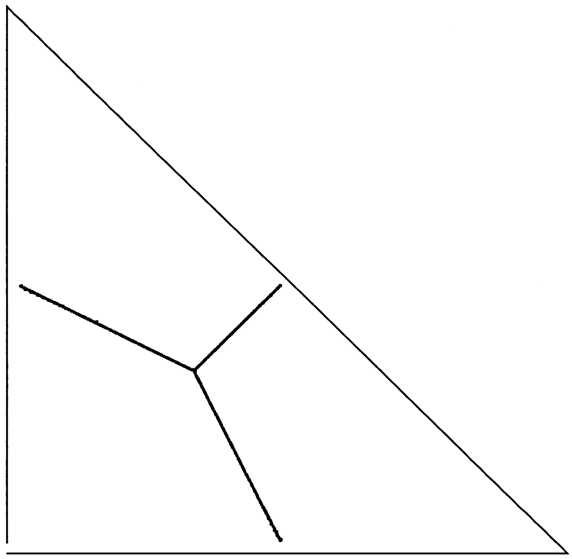

(a) Medial axis of a triangle with noise-free boundary.

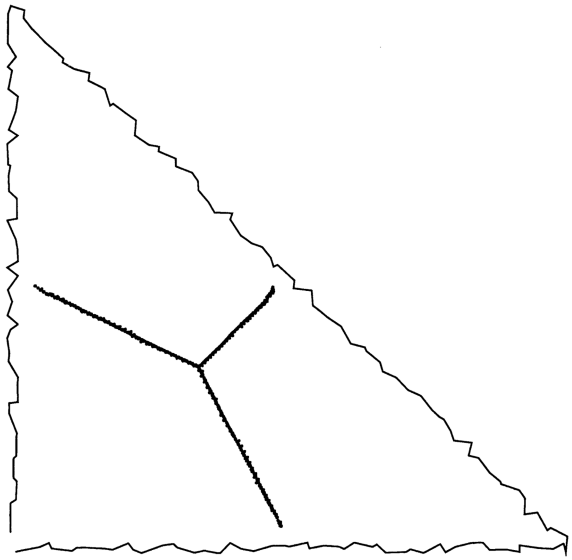

(b) Medial axis of the same triangle with added boundary noise.

Figure 8.
Medial axis of a triangle. When noise corrupts the curve, the main branches of the skeleton remain almost unaltered.

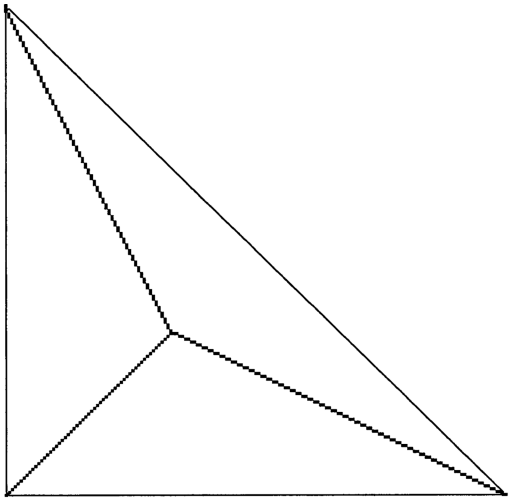

(a) Euclidean topology preserving thinning of a triangle.

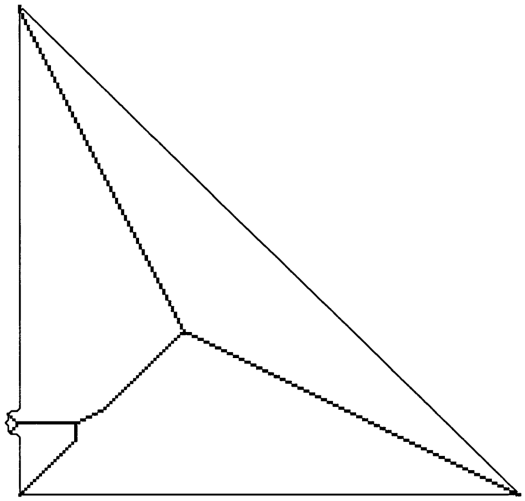

(b) Euclidean topology preserving thinning of a slightly perturbed triangle.

Figure 9.
Small boundary perturbations can have large effects.

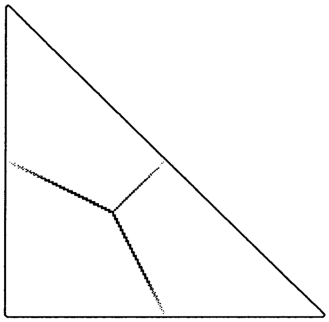

(a) Gray level **affine** medial axis of a triangle.

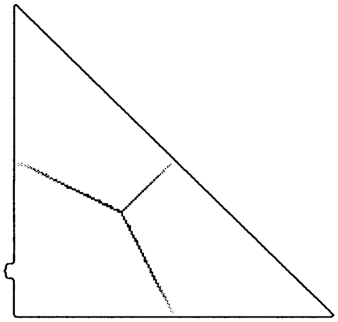

(b) Gray level **affine** medial axis of a slightly perturbed triangle.

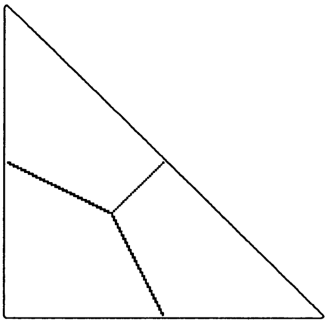

(c) Thresholded **affine** medial axis of a triangle.

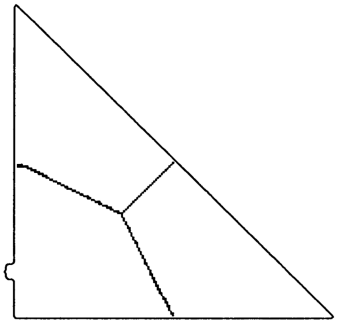

(d) Thresholded **affine** medial axis of a slightly perturbed triangle.

Figure 10.
The perturbed boundary of the triangle does not affect the affine medial axis.

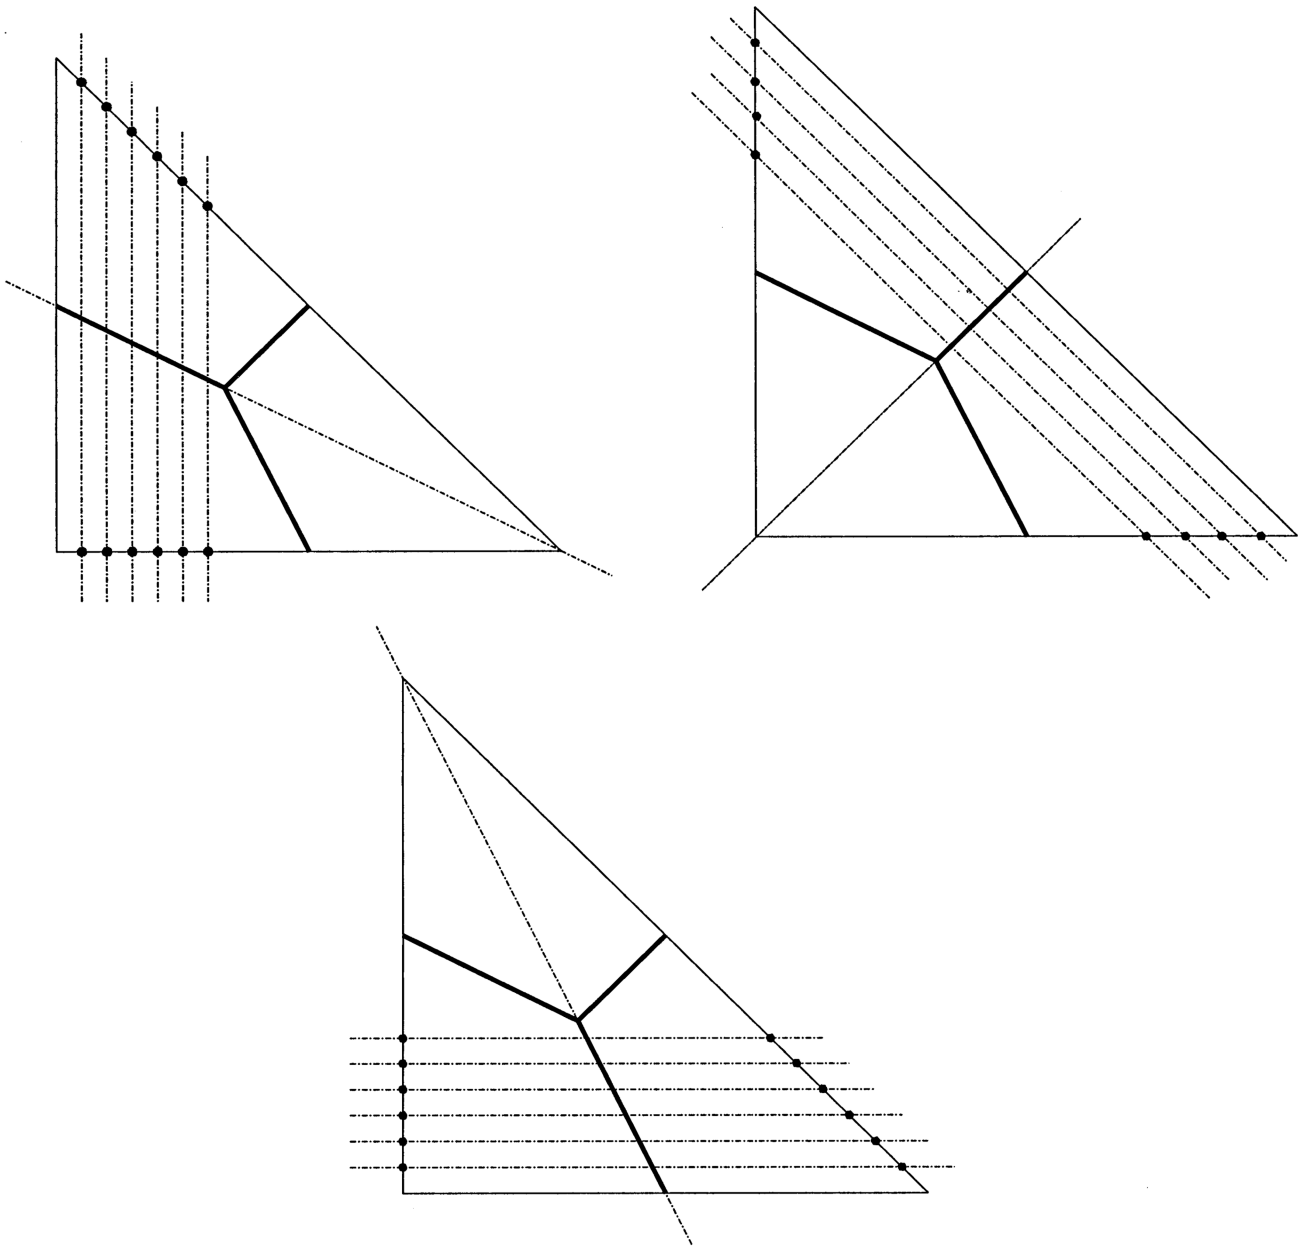

Figure 11.
Exemplary skewed symmetries for the triangle.

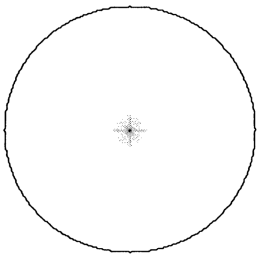

(a) Gray level affine medial axis of a circle.

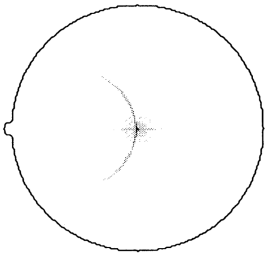

(b) Gray level affine medial axis of a slightly perturbed circle.

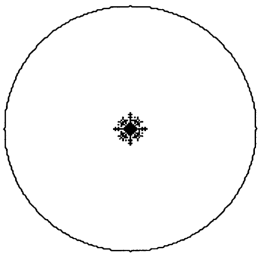

(c) Thresholded affine medial axis of a circle.

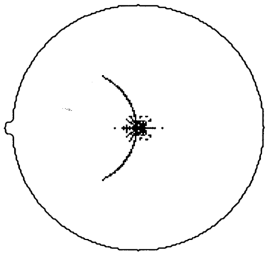

(d) Thresholded affine medial axis of a slightly perturbed circle.

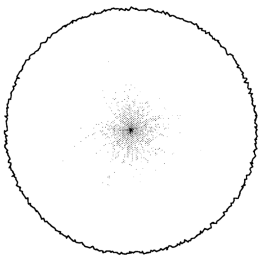

(e) Gray level affine medial axis of a noisy circle.

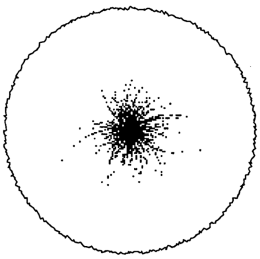

(f) Thresholded affine medial axis of a noisy circle.

Figure 12.
If the perturbation contributes to the symmetry of the object, it will contribute to the medial axis.

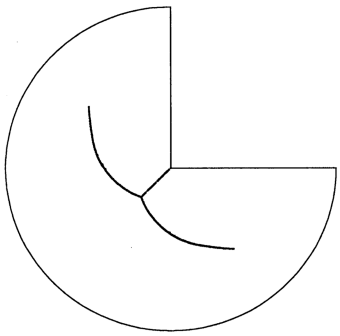

(a) A concave curve and its skeleton.

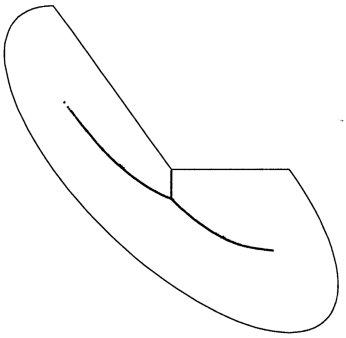

(b) After an affine transformation, the skeleton keeps the information about the original symmetry.

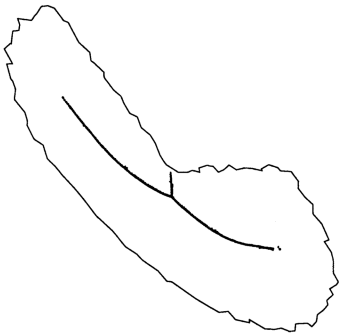

(c) When noise corrupts the curve, the main branches of the skeleton may still be recognized.

Figure 13.
Medial axis of a concave curve.

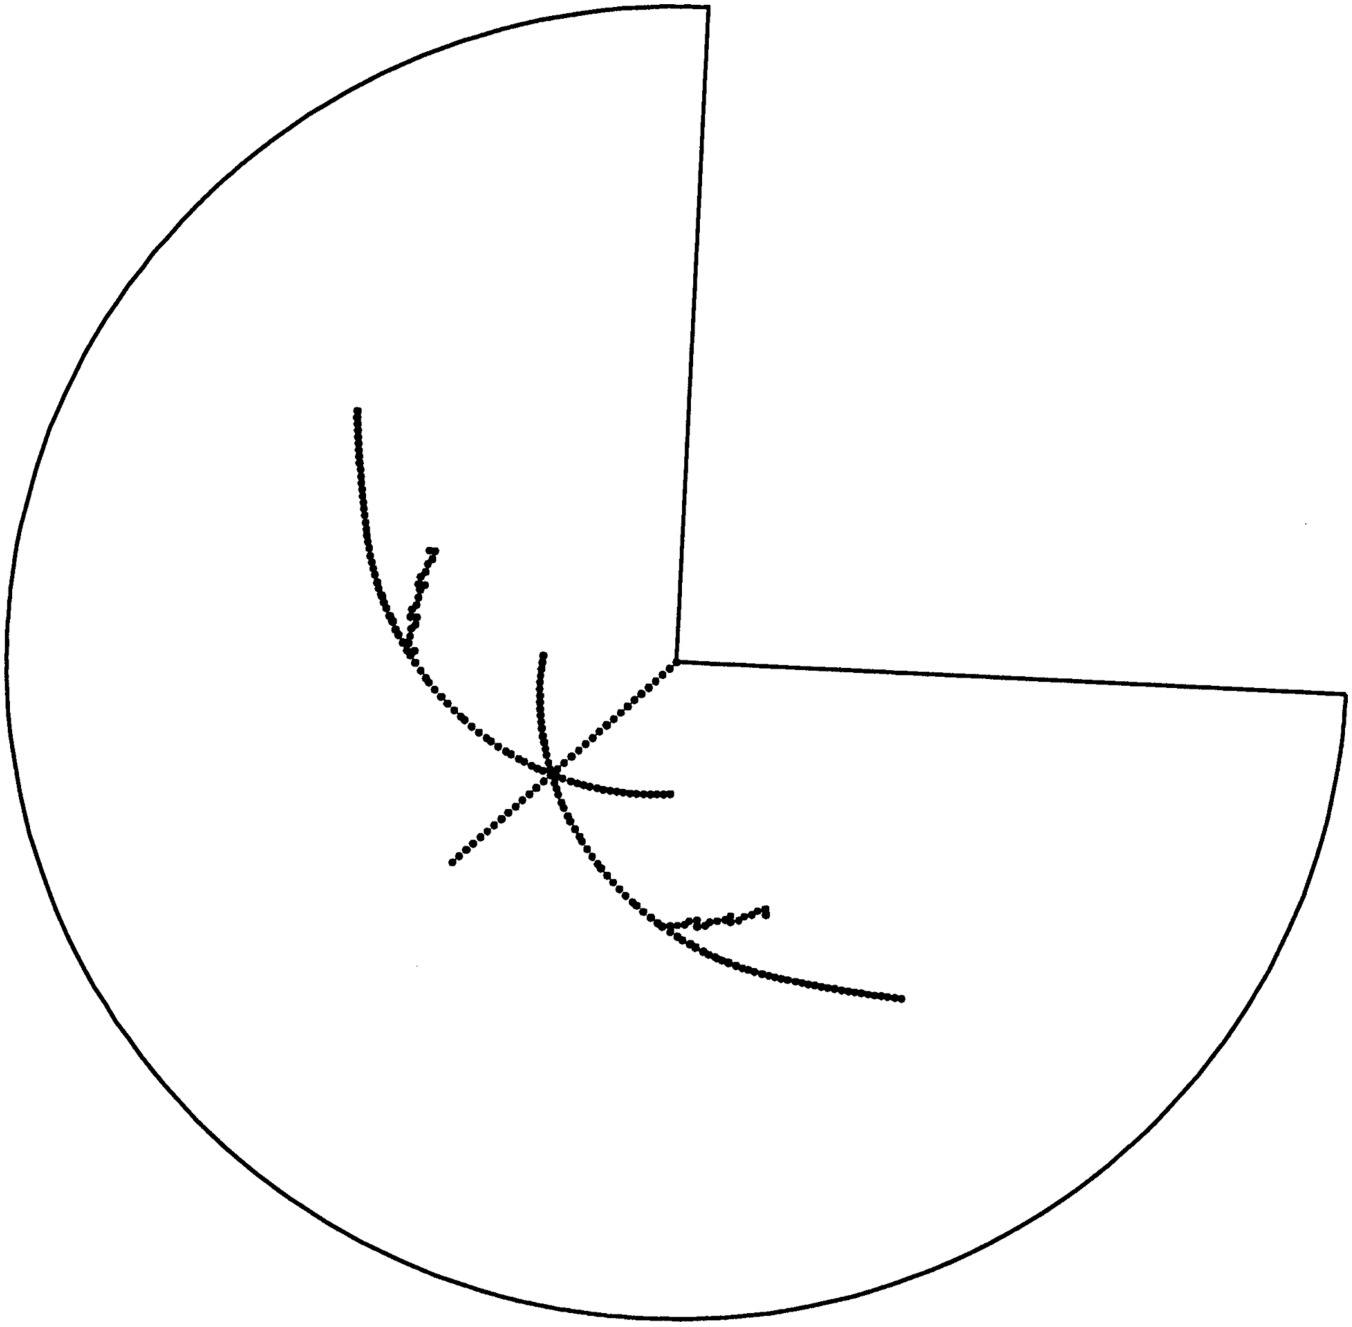

Figure 14.
Without requiring the distances to be global positive minima, we obtain the AASS.

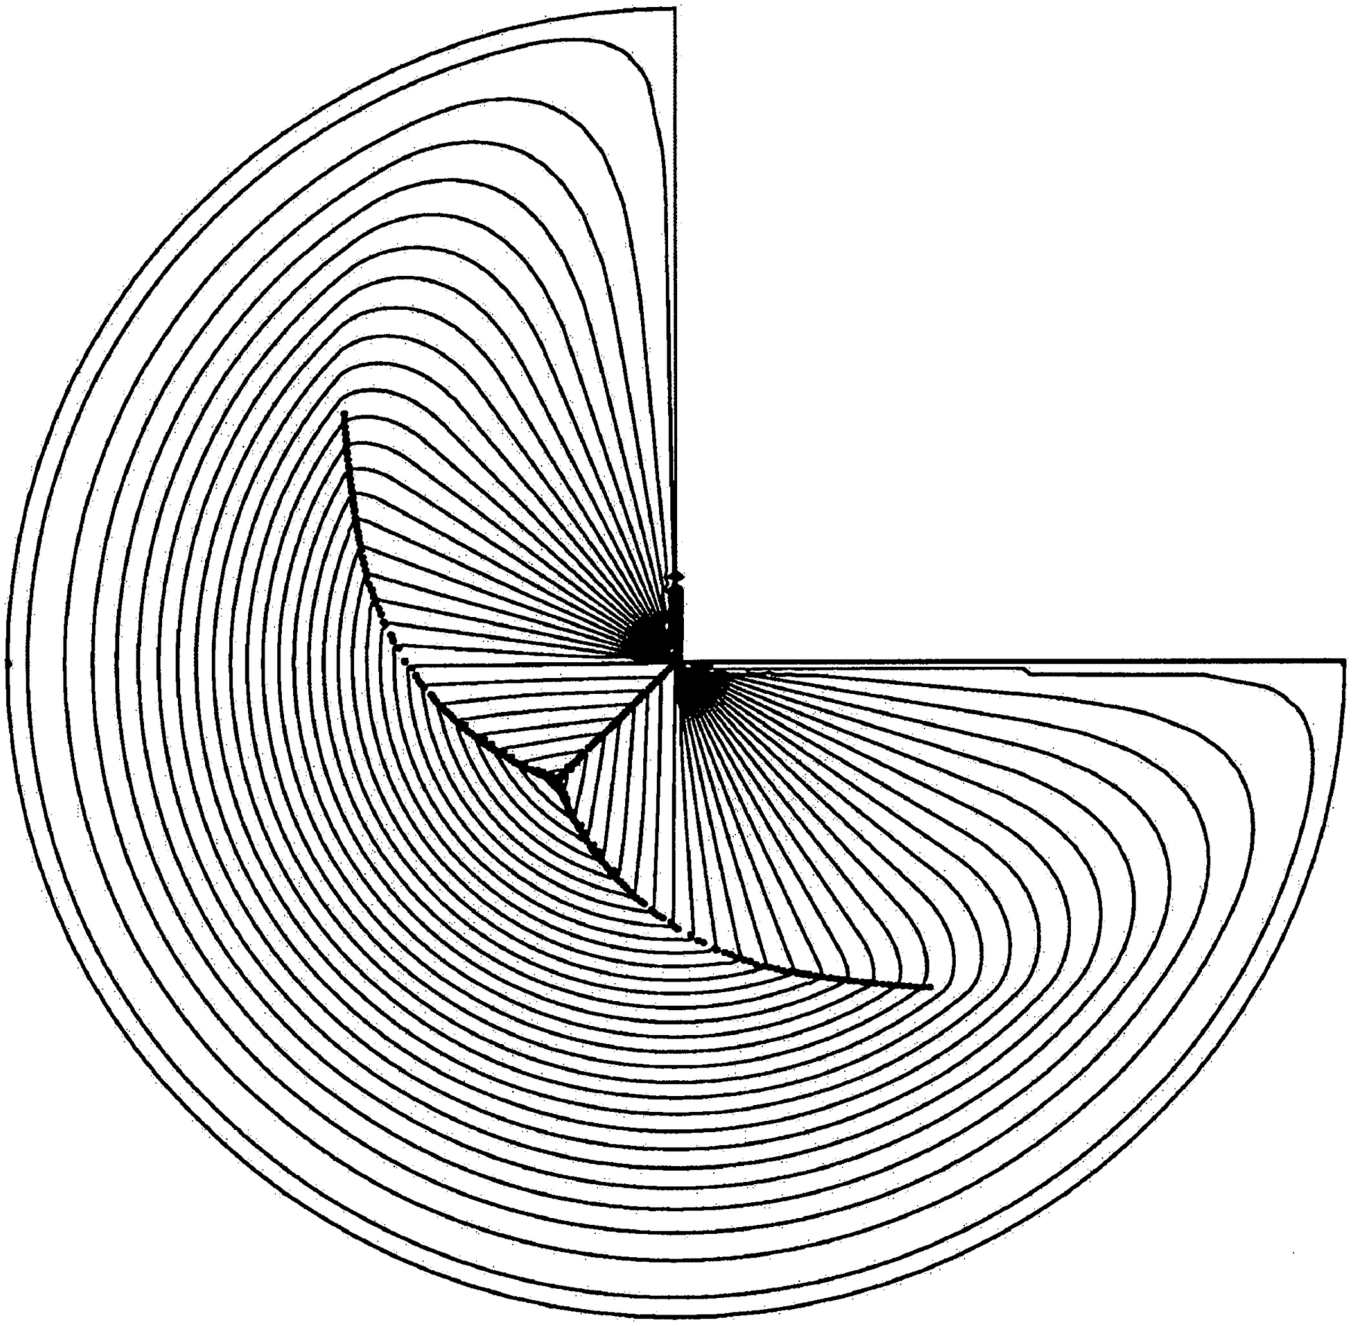

Figure 15.
Affine erosion of a truncated circle. Each line represents a curve of constant chord area. The affine shocks on the evolution are the skeleton points (dots).

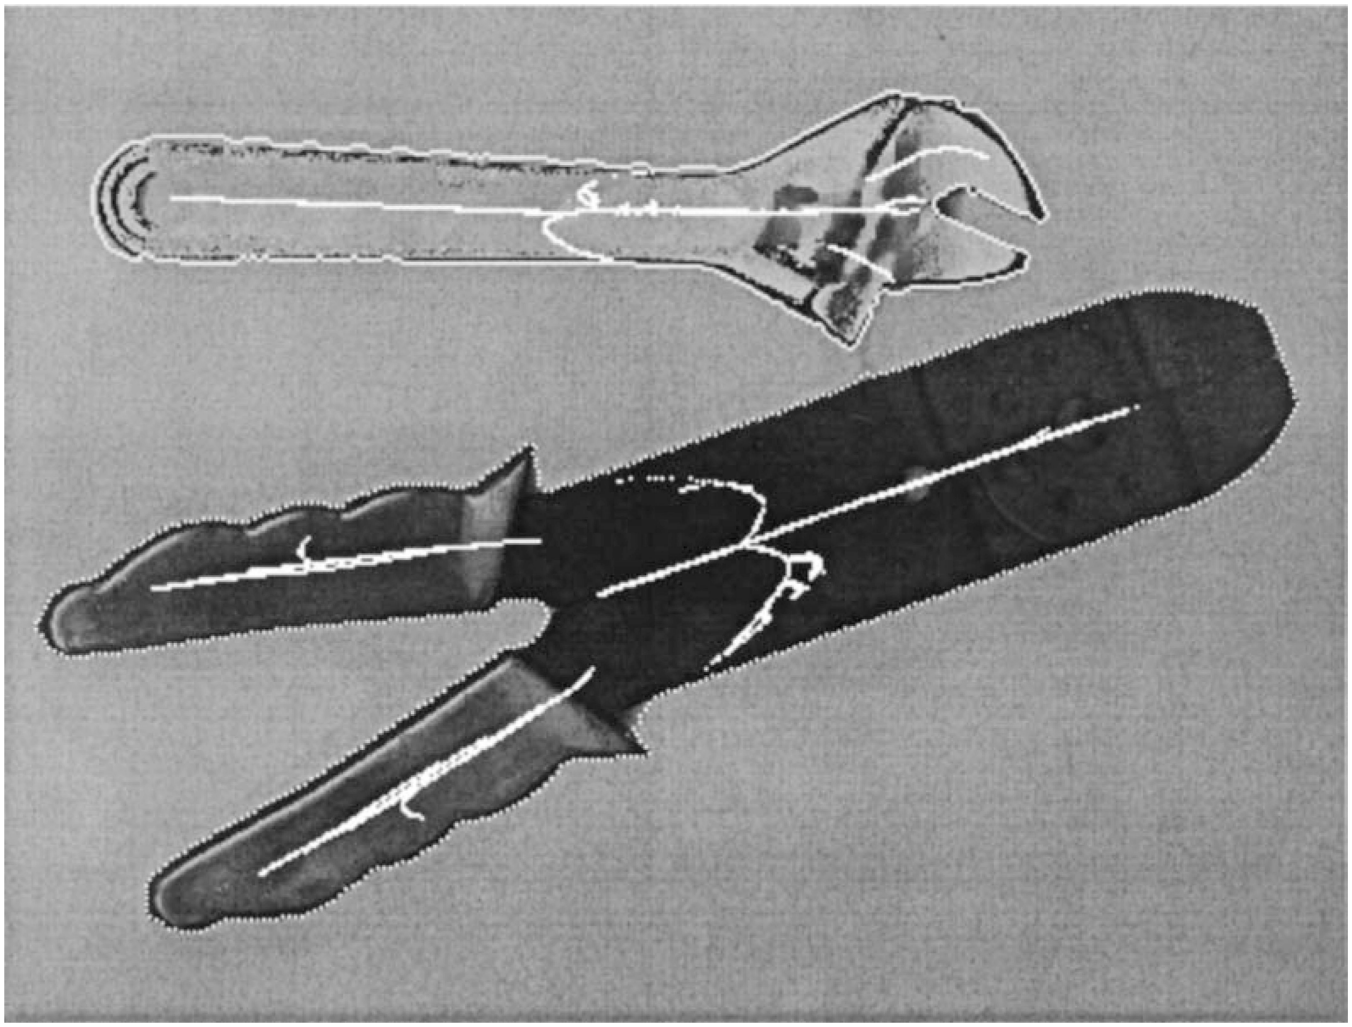

Figure 16.
Thresholded affine medial axis of the tool shapes.

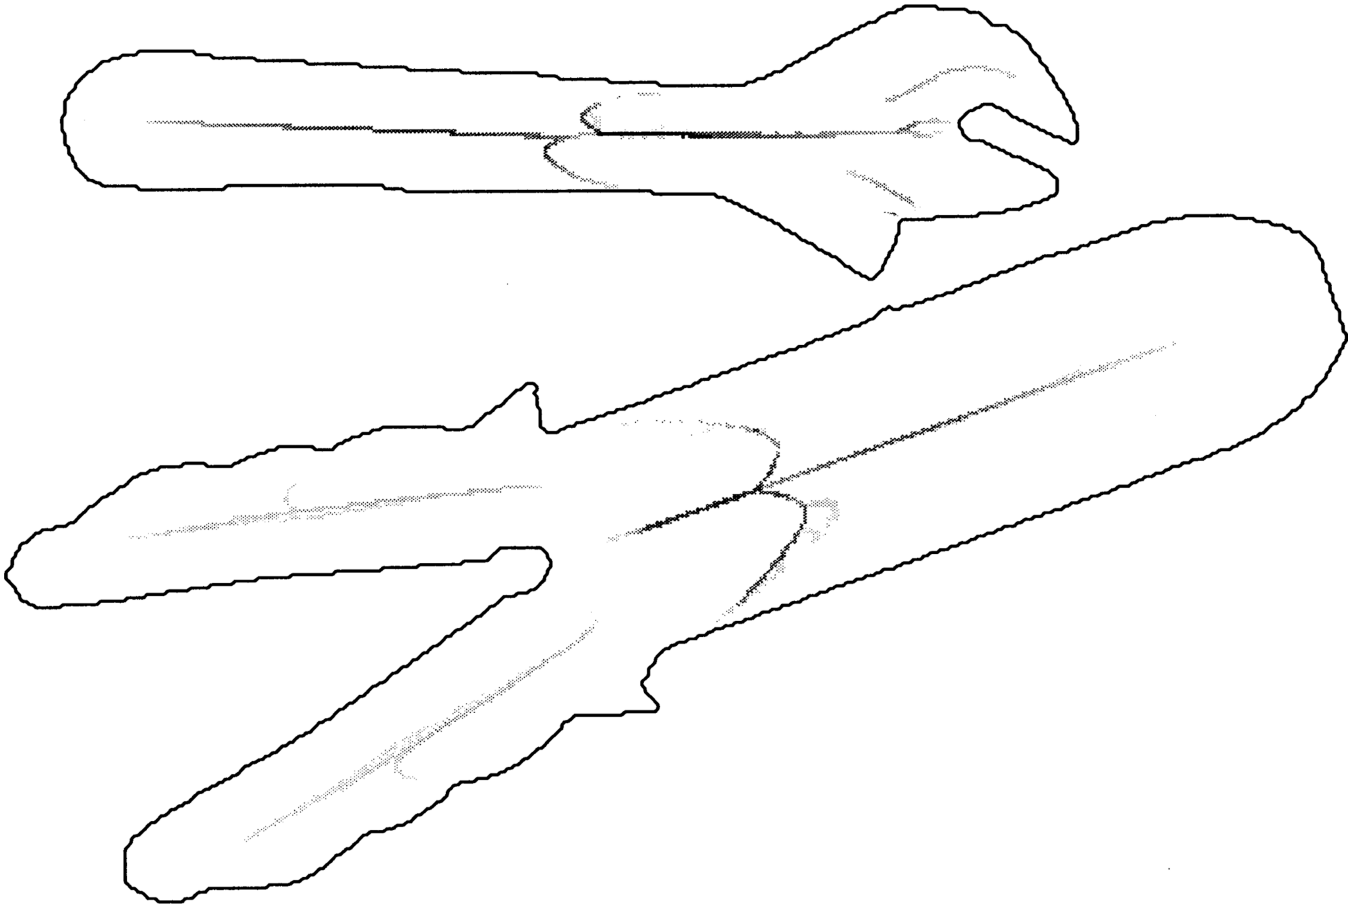

Figure 17.
Gray level affine medial axis of the tool shapes.

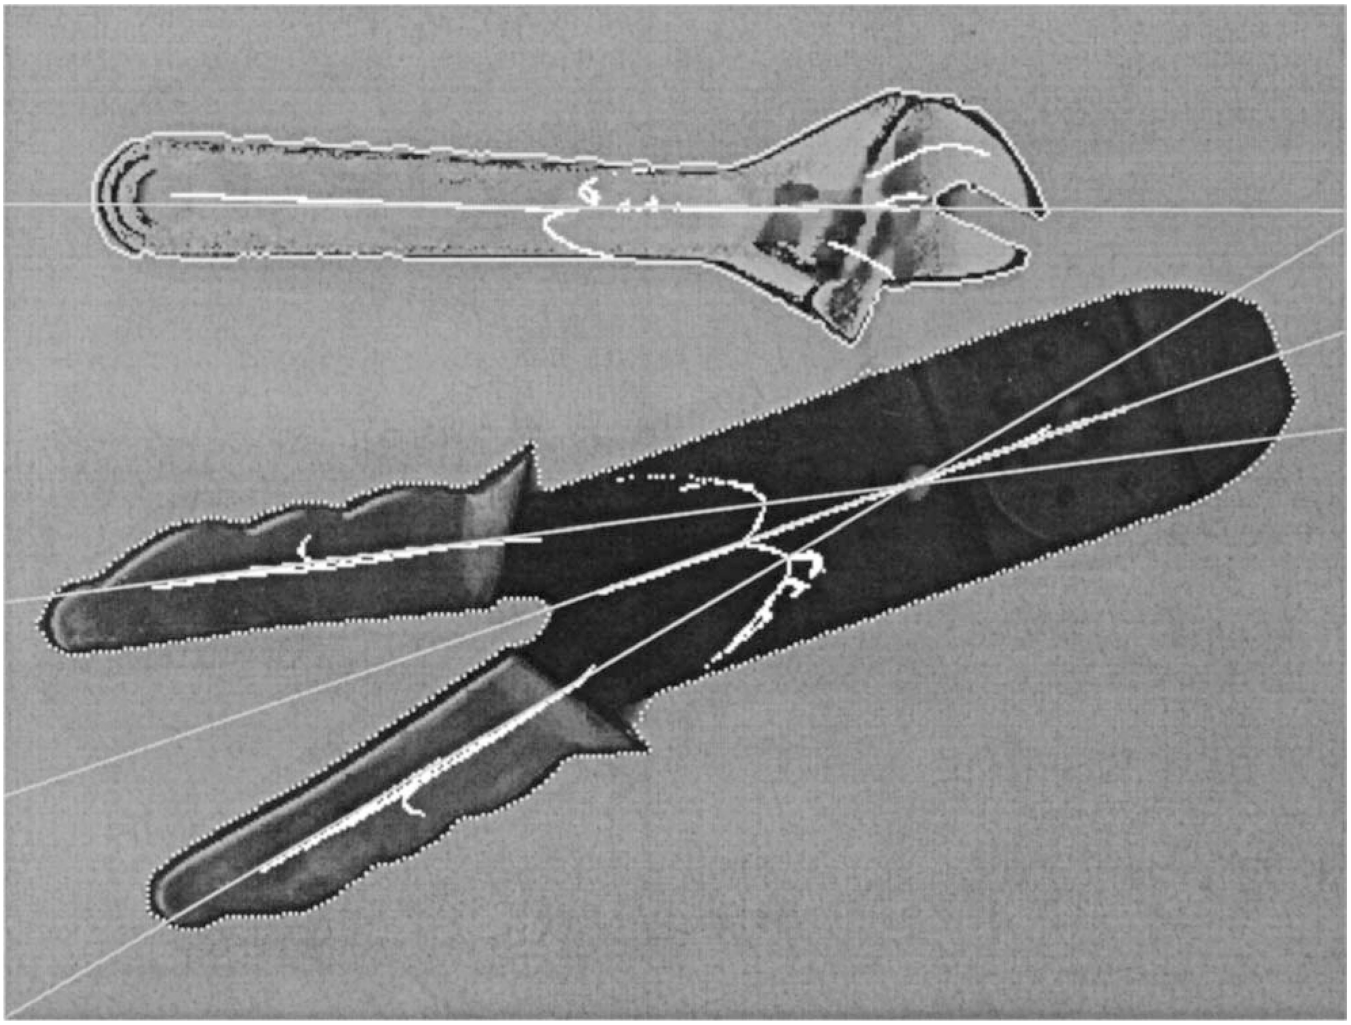

Figure 18.
Thresholded affine medial axis of the tool shapes, overlaid with lines detected by a Hough transform of the thresholded affine medial axis.

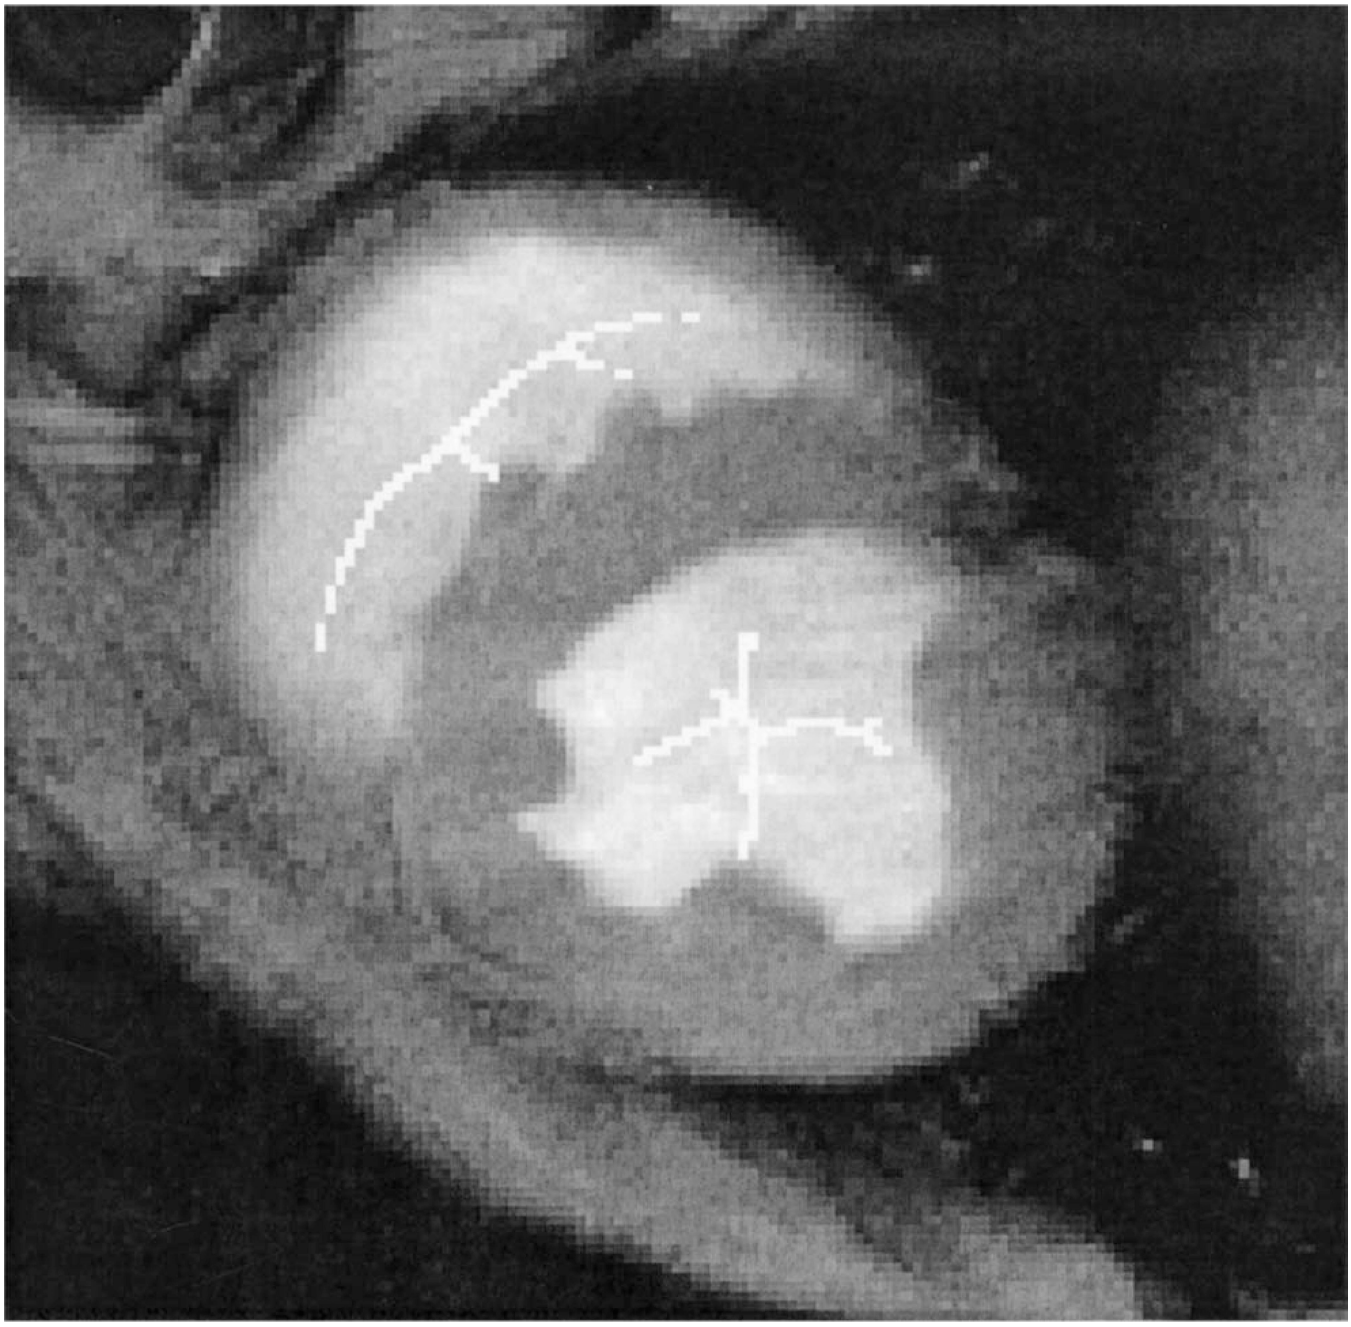

Figure 19.
Thresholded affine medial axis of the heart.

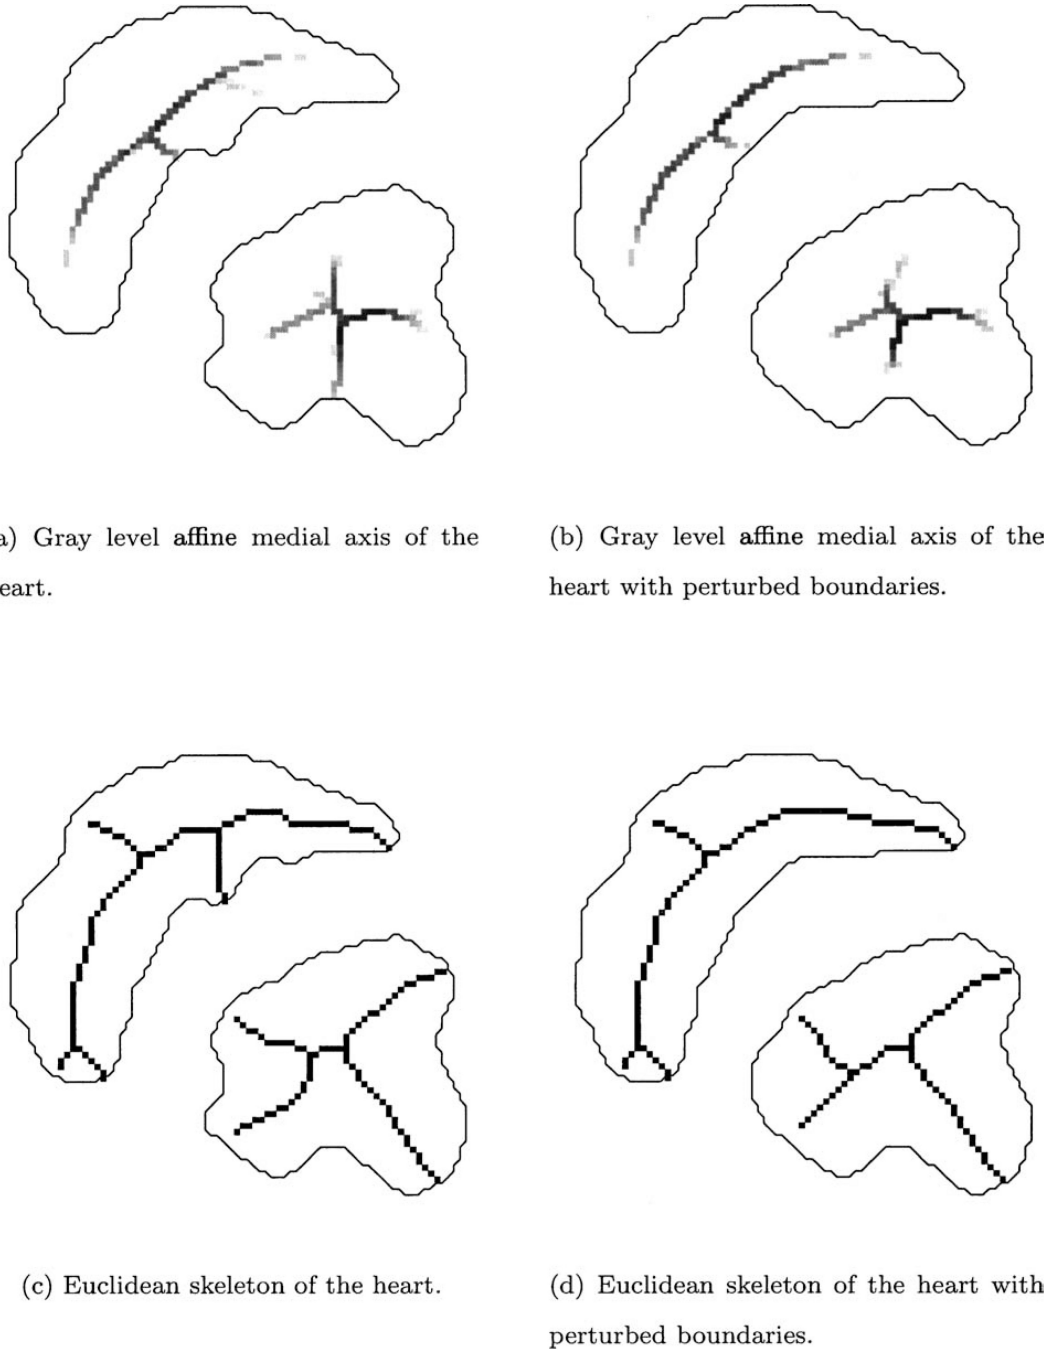

(a) Gray level affine medial axis of the heart.

(b) Gray level affine medial axis of the heart with perturbed boundaries.

(c) Euclidean skeleton of the heart.

(d) Euclidean skeleton of the heart with perturbed boundaries.

Figure 20.
Gray level affine medial axes and Euclidean skeletons of the heart.

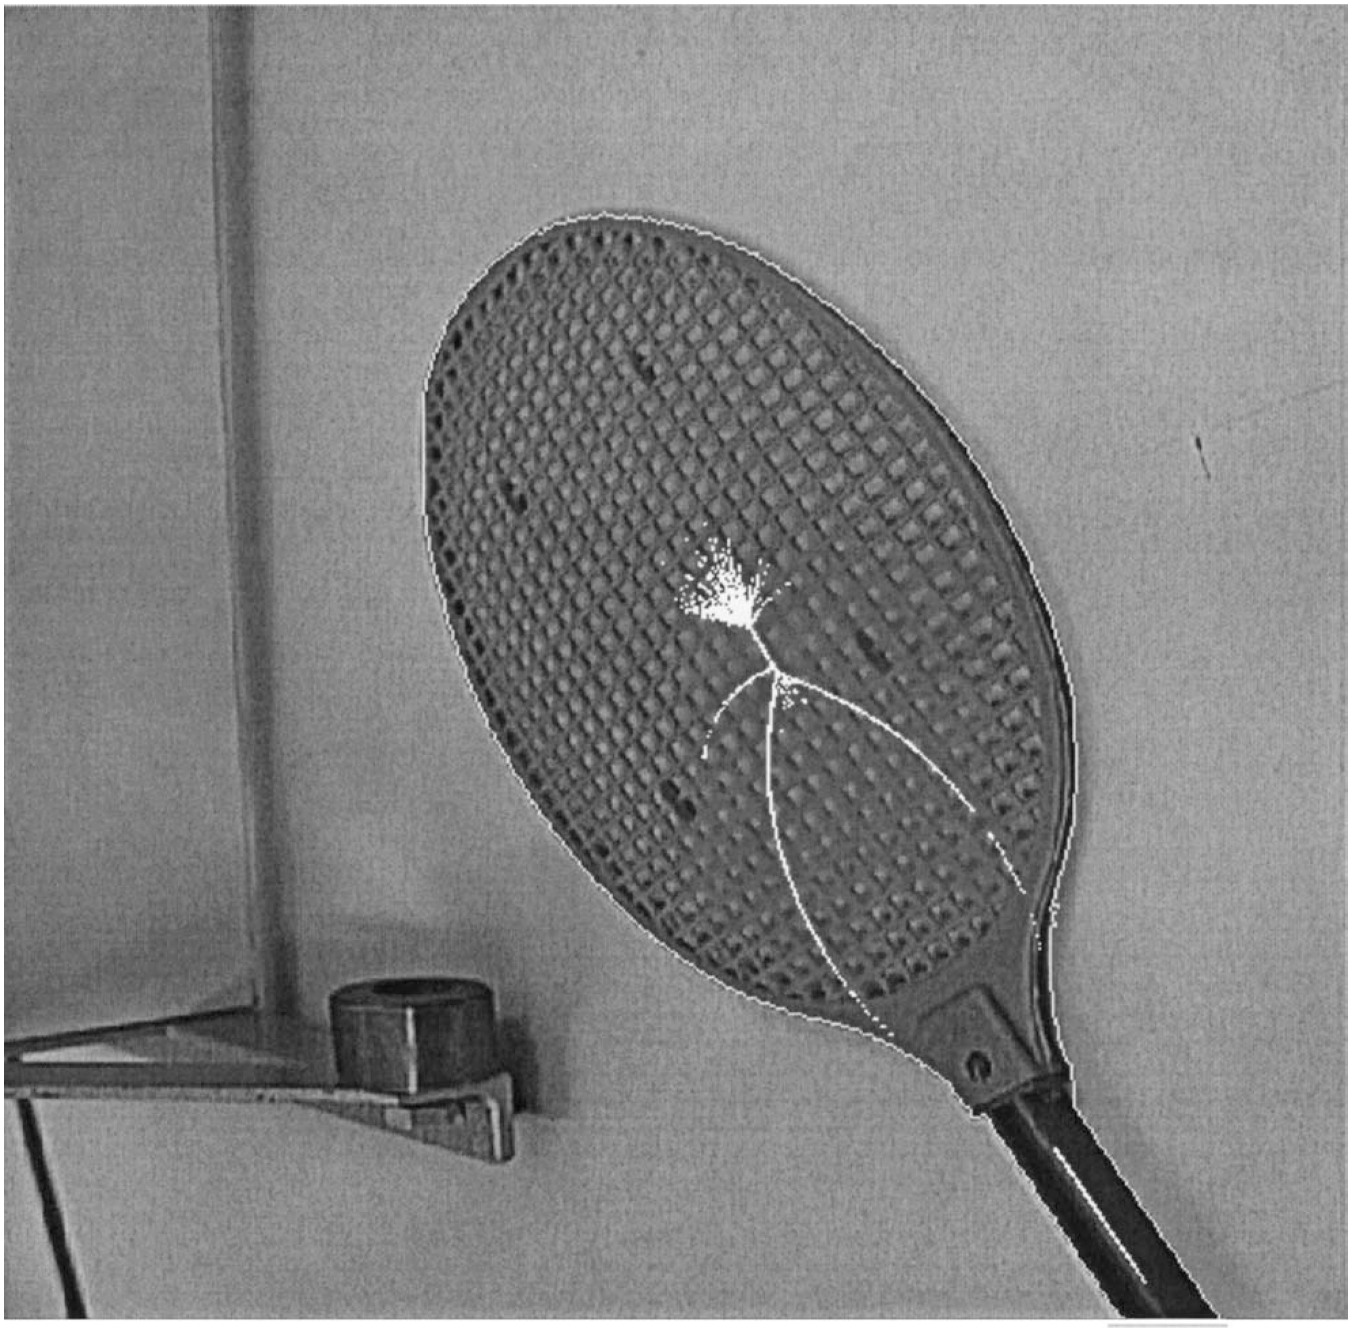

Figure 21.
Thresholded affine medial axis of the tennis racket.

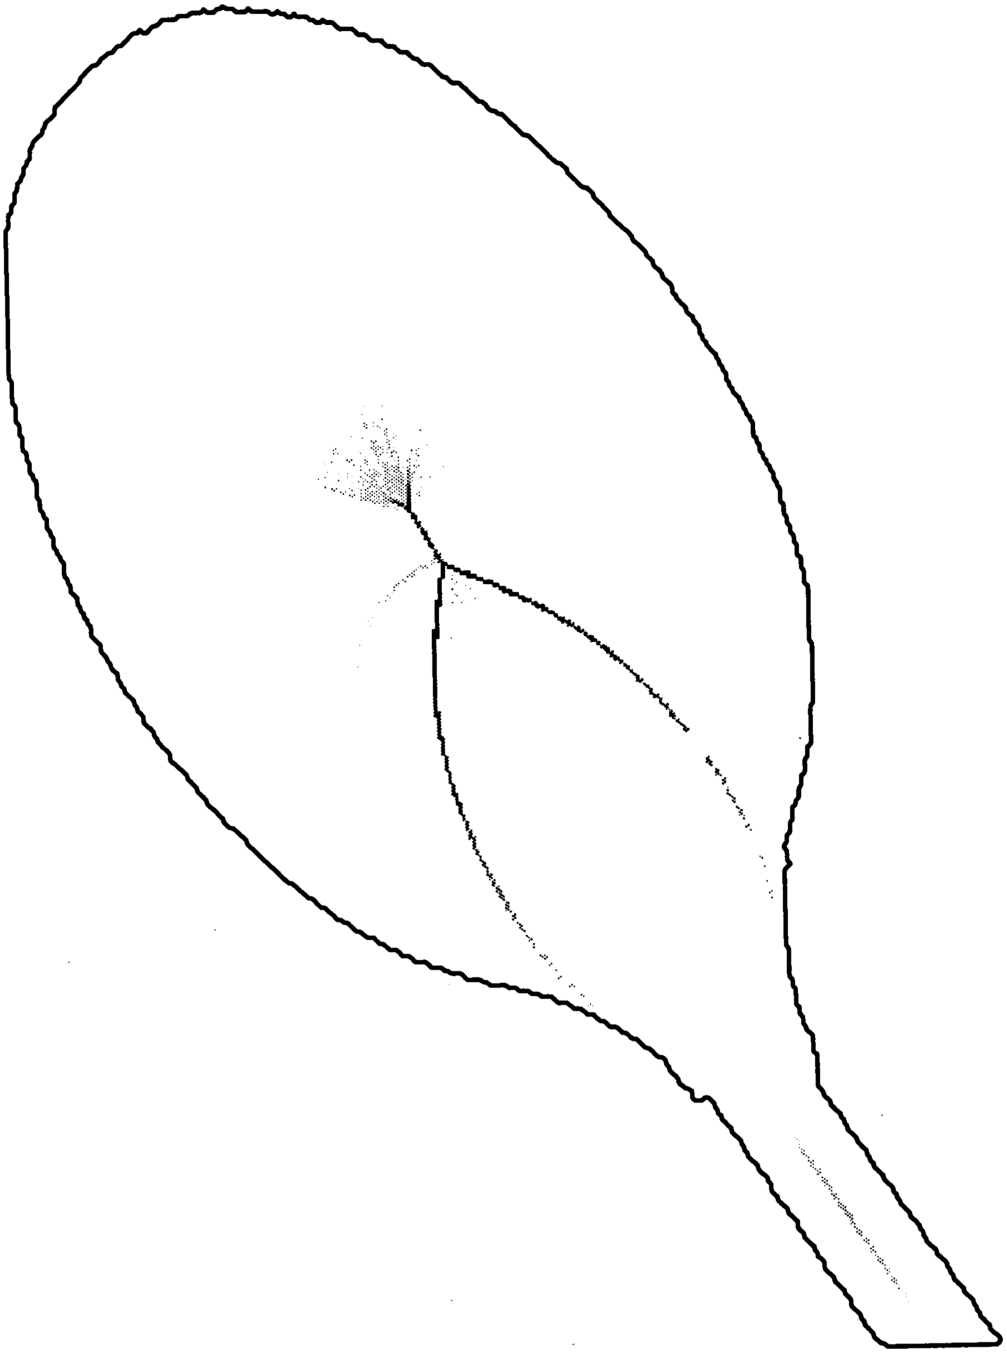

Figure 22.
Gray level affine medial axis of the tennis racket.