Medial Axis, Symmetry Axis, Skeletonization

Literature: Harry Blum, Biological Shape and Visual Science, J. theor. Biol. (1973) 205-287

Harry Blum and Roger Nagel, Shape description using weighted symmetric axis features Pattern Recognition, Vol. 10, 167-180

Motivation: Characterization of object structures

so far: edges, contours: local properties, need to be linked to global structures by edge-linking or grouping

Hough: global grouping, but restricted to predefined shapes, no flexibility

Topology: too general

Geometry: too restricted: Euclidean, projective

Goal:
- full description of object shape
- lossless (reconstruction possible)
- access to geometric features
- flexibility

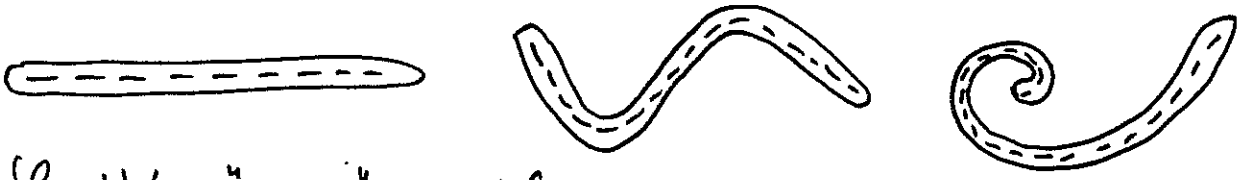

flexible "worm" model

medial axis: same length, no junctions,
radius constant
only difference: course of axis
=> same "animal", different instantiations

MAT: Medial Axis Transform

- higher flexibility
- more natural description
- hierarchical description
- figure – subfigure relationship, graph description

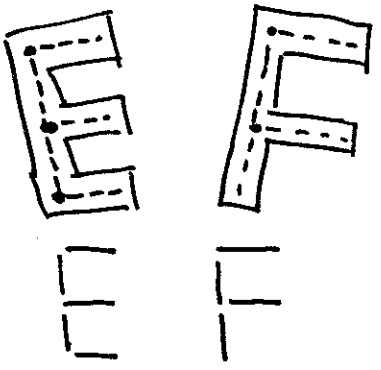

Construction of MAT

Prairie-fire analogy: Propagation of fire fronts, starting at the boundary and propogating with constant speed.

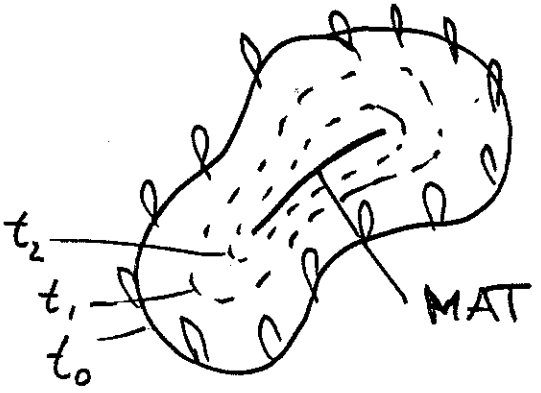

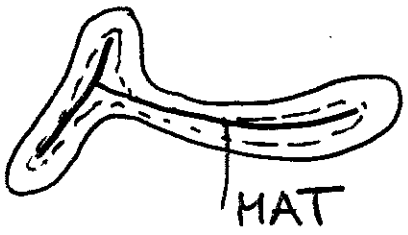

At locations where fire extinguishes: Location of axis. time & radius

Time level set: Distance Transforms

Location where fire extinguishes: skeleton

see Figs 218 223

Geometrical description:

Skeleton: Locations of center points of all largest disks fully inside the object

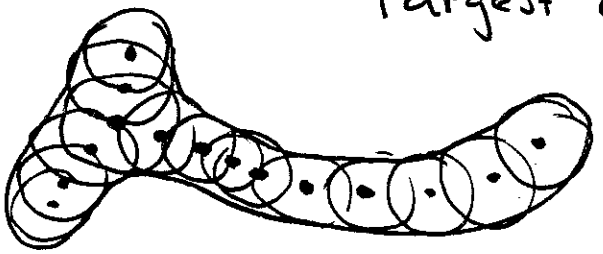

(smaller disks fully covered by larger disks are removed!)

Reconstruction:

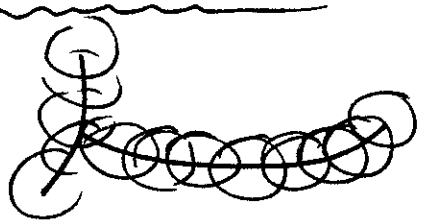

Superposition of all disks defined by the skeleton.

Disk: center & radius

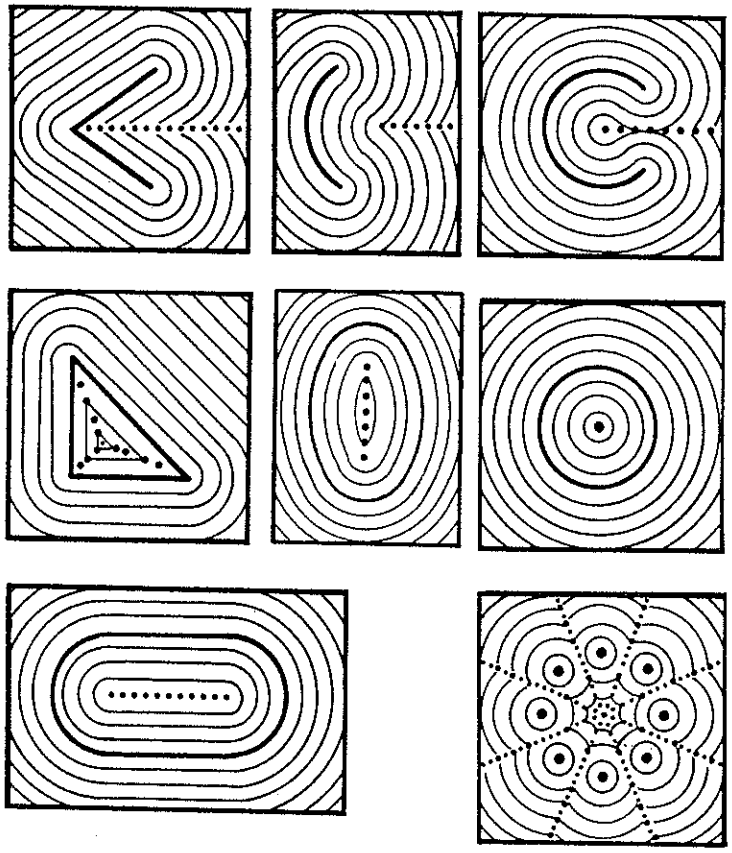

FIG. 8. A succession of grassfire wavefronts for some simple inputs. At the top, the grassfire is started along open contours. The sym-ax (shown dotted) occurs on the inside of the angle only, starting at the center of curvature and starting at a pinch in the space—ending at the center. The center panel shows some closed contours combining the above features. The sym-ax disappears at the largest inscribed circle. Note that the boundaries are convex and have no outside sym-ax. The bottom panel shows the sym-ax for a parallel boundary and for a set of points on a circle. In the parallel oval, the grassfire disappears all at once. The points on a circle give an example whereby the object is in the ground and discrete points can be treated as equivalent to a contour in generating an object.

concept analogous to "phase velocity" in wave physics. The circular arc shows the sudden appearance of the sym-ax when the wavefront reaches a contour's center of curvature. These open line inputs are extended to form closed convex object boundaries. No sym-ax appears outside the object. The internal wavefront disappears at the largest sym-dist point, the center of the largest inscribed circle. The circle shows an isolated sym-point, and the parallel oval, a sym-ax with constant sym-dist. The last input is a scene of dots arranged in a circle. It shows that the sym-ax may branch in the ground as well as in the object. An interesting figure completion property also appears. The two point scene of Fig. 5 showed the sym-ax issuing along the bisector of the two points; starting with infinite velocity, and decreasing

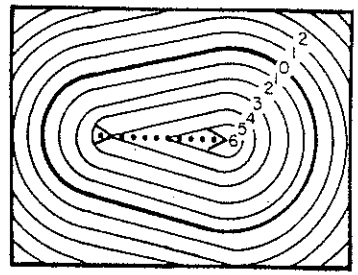

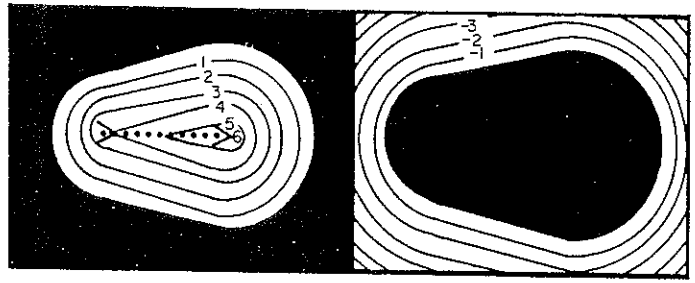

FIG. 9. Grassfire and inverse grassfire for boundary, object and ground. Wavefront times are shown for boundary. When the object is known, grassfire can be separated. Object times (or distances) are positive; ground, negative. The inverse grassfire is generated by exciting the internal sym-ax in reverse times. It generates the wavefronts in reverse order. At time zero, the original object is generated. Note that new inverse grassfire points are always excited outside the wavefront generated by previous points since the sym-ax velocity is faster.

ment by complicating the grassfire, coding distance in a way other than absolute time. In the unsynchronized case, the inputs at a point may arrive at different times and flowthrough must be permitted. Each point is assumed to contribute an independently increasing circularly symmetric function. Each point of the space must respond to the first value of the function from each point of the excitation, even though it arrives at a later time. Distance may be coded by amplitude change, rise-time change, variance spread (in a noisy velocity grassfire), velocity change, waveshape change, frequency change, or a host of other distance dependent processes in a plane. An extremely interesting system can be set up using a two velocity propagation in one or two planes. A point of excitation would lead to two waves, the time difference between which would be a code of distance. The two plane version is particularly suited to doing the direct and inverse transform, since the first excitation would be inserted in the slow plane and the second in the fast—both arriving in synchrony at the object boundary. Fortunately, in all of these schemes each point must know only that it has received two or more excitations at minimal distance. It need only remember the minimal distance that has passed it. An awareness of the variety of implementations possible for the grassfire is essential in a search for the process in biological systems and the synthesis in artificial systems. However, it is only diverting tutorially with respect to shape properties. Therefore, the development continues with the simple grassfire. We should remember, however, that it can take a much wider variety of forms than the simple one we are using to explore the geometry.

The inverse grassfire excitation may be thought of as a wavefront generator as well as a disc initiator. Consider the previous figure again. The sym-ax at the non-end points has been generated in a continuous way, the sym-point always moving faster than the fire velocity. When this process is inverted, new points are started outside the expansion of earlier ignited points since the sym-ax velocity exceeds this fire velocity. The rate and direction is precisely that needed to generate an identical and oppositely moving wavefront to the one that generated the sym-ax. Thus, such smoothly moving points generate the same pannormals at the point as those that created the point in the first place. We may look at the sym-ax as an antenna surface receiving a wave at the sym-dist times. Again, this reciprocity between time inversion and the receiving–transmitting relationship of antennas is well known. The main innovation introduced here is the non-linear propagating space of the grassfire and the definition of the receiving surface by the input itself rather than by a predetermined boundary condition. The receiver antenna is now generated by the signal. In addition, sine waves and associated phase angles have been avoided by the use of impulse signals.

Another view of the transform can be obtained by substituting a three-dimensional static process for the two-dimensional dynamic grassfire (see Fig. 11). If a space–time plot of the grassfire is made with the input plane as the x, y surface, and time as an orthogonal z coordinate, a surface will be generated. Each plane at time t will contain the firefront at that time. Of course, for time we may substitute a parallel distance coordinate and get a purely geometric surface.† Since each parallel curve is the union of discs of

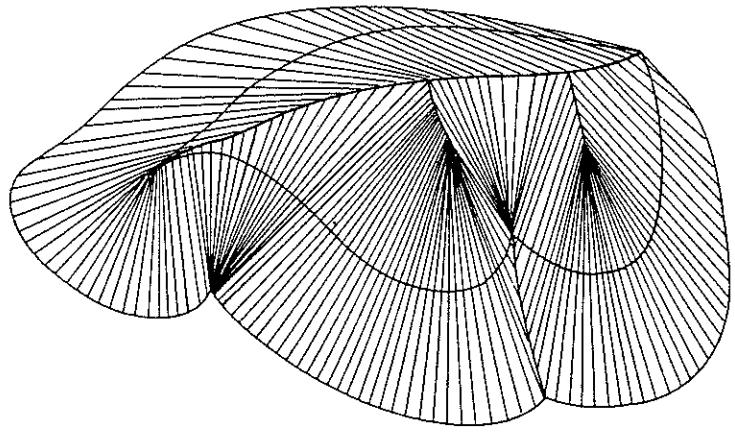

FIG. 11. Static three-dimensional representation of a grassfire. The vertical coordinate represents time, or distance between parallels. The slope of this surface, which is the reciprocal of the grassfire velocity, is everywhere one except at the sym-ax, where it lies between zero and one. The surface is the boundary of the union of right angle cones with apices on the object boundary.

increasing radius, the surface generated is the boundary of the union of right cones whose apex sits on the input excitation and whose axis is perpendicular to the *object plane* (the x, y plane, in which the input excitation occurs). Figure 12 shows the surface for a humanish boundary generated in precisely such a way with a conical cutter on an engraving machine. The surface is smooth and has a maximum directional slope of one everywhere except at the sym-ax points, where it lies between zero (for parallel curves) and one. When the "outside" surface is made negative and reflected across the input

† Such a surface appears in analytic function theory and is called a "distance surface" there (see, for example, Goodman, 1964). It appears as a pursuit strategy surface in the Theory of Games and is called a "sandpile function" there (Isaacs, 1965). Neither of these fields develop the geometry we do here.

Natural language for shape description based on MAT (H. Blum)

see Figs. p.160, 240, 263, 260

Example: Qualitative & quantitative characterization

	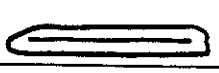		
# segments	1	1	1
point types	ENN...NNE	..	..
length	L	L	L
shape	worm	worm	worm
axis	straight	right-/left-circular	spiral left
axis curvature	∅	neg/pos	pos
axis curvature change	∅	∅	pos

see pages 12 and 13,14 for elements of shape language

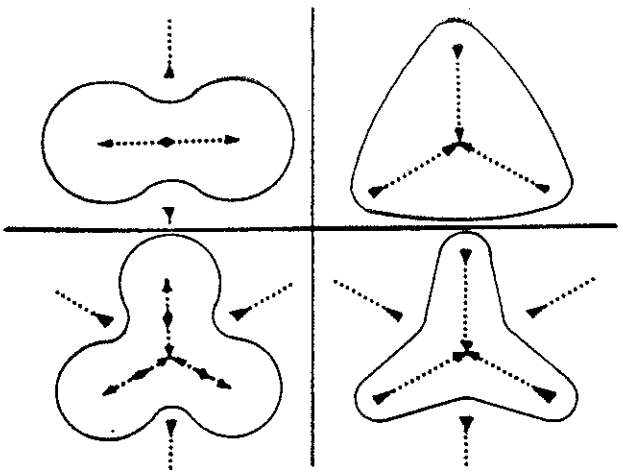

Fig. 5. Representation of simple shape properties by directed graph structure of the MAF.

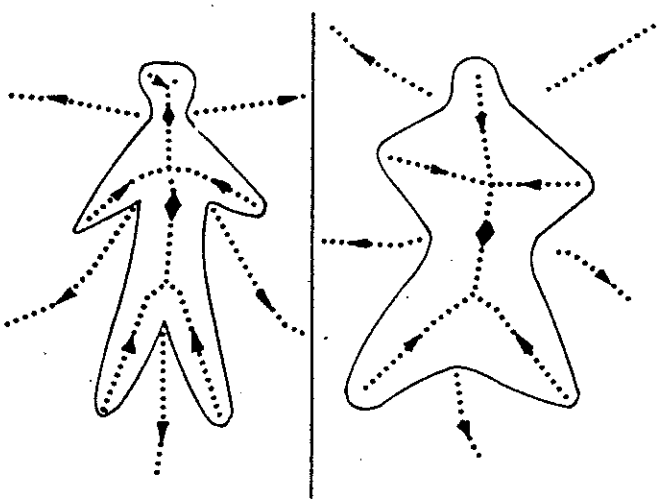

Fig. 6. Two anthropomorphs and their MAFs.

it to be in defining natural shape properties on the new primitive descriptors. It could be argued that the properties I have discussed in the last paragraph could be gotten by such simple contour properties as counting the inflection points and the total curvature of the shape.

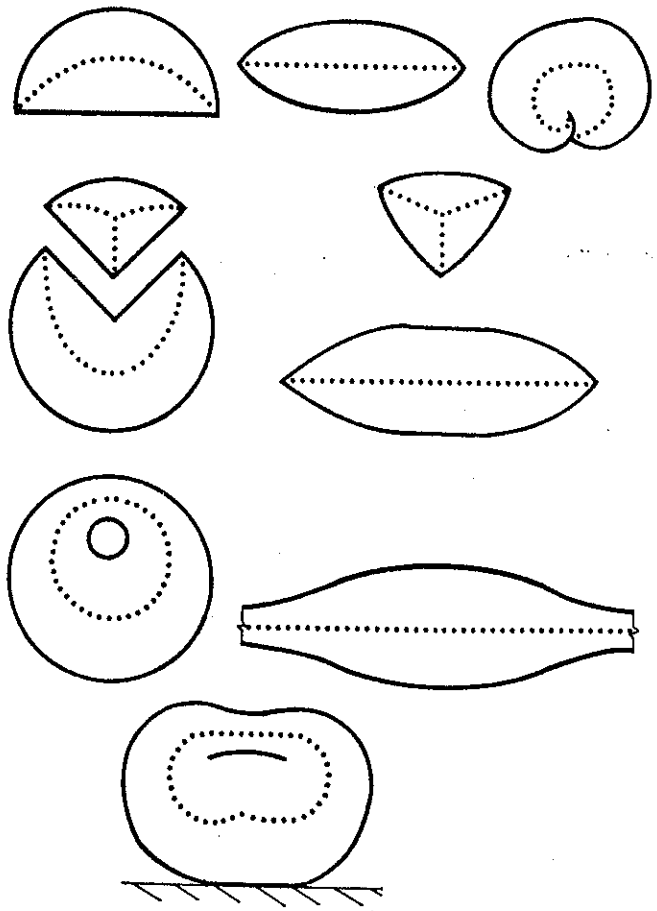

FIG. 21. Flexure of a number of objects. A careful examination of these figures will help considerably in understanding the process. The eccentric annulus, which must be pasted to form a cigar band, is a good example of the unexpected results.

intercepts of all perpendiculars to the line and splitting them symmetrically about the line. This process, as well as other types of symmetrization, circularizes (maintains area but reduces perimeter). *Sym-ax symmetrization*, or simply *symmetrization*, where no ambiguity exists, shall consist of making the asymmetric or axis curvature zero everywhere, that is straightening the sym-ax. This procedure is always doable for a *linear object*, one whose

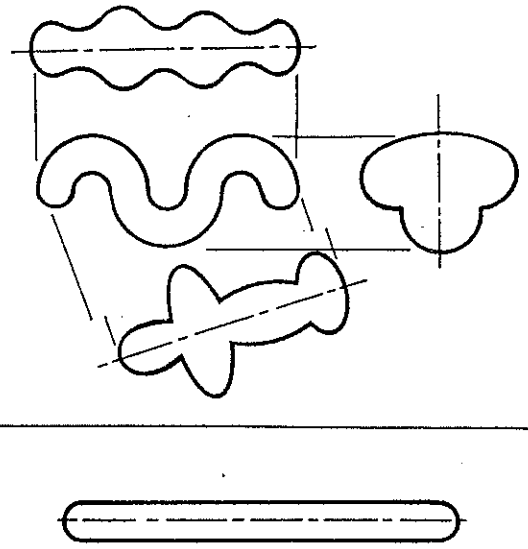

FIG. 22. Steiner symmetrization and sym-ax symmetrization. Conventional symmetrization consists of splitting chords through an object about a line (Steiner) or a point. Steiner symmetrization about three lines is shown. Its result is not intuitively satisfying for the above. Conventional symmetrization circularizes—maintains area but reduces perimeter. Sym-ax symmetrization maintains both.

sym-ax is not closed and has no branches. We showed some examples earlier in Fig. 21, including non-linear objects, but do not pursue the topic further. Symmetrization and other flexures for a number of "plane geometry" objects have been included in Fig. 21. Because flexing does not always give the expected results for these, they are illuminating examples. It reminds us that the transform is not designed for these kinds of objects. They are the special cases.

We now define *conflexity*, the property of an object whose symmetrization is convex. Thus, the kidney bean of Fig. 21 is conflex, while the tadpole of that figure is not. It is apparent that an object possesses this property only if its width curvature is everywhere non-positive—it does not curve away from its sym-ax.

8. Symmetric Disc Coordinates

It is not the intent of this shape mathematics to become involved with the calculus and analysis. But they are so insightful that one cannot bypass them.

B-morphs may be combined. For example, the isosceles triangle in perspective can be described by the vertical sym-ax to define the vertical and by the base line to define the horizontal. Note that there is an error in this vertical that is an increasing function of the perspective angle of the figure, suggesting some testable visual experiments. When the trapezoid of the figure is

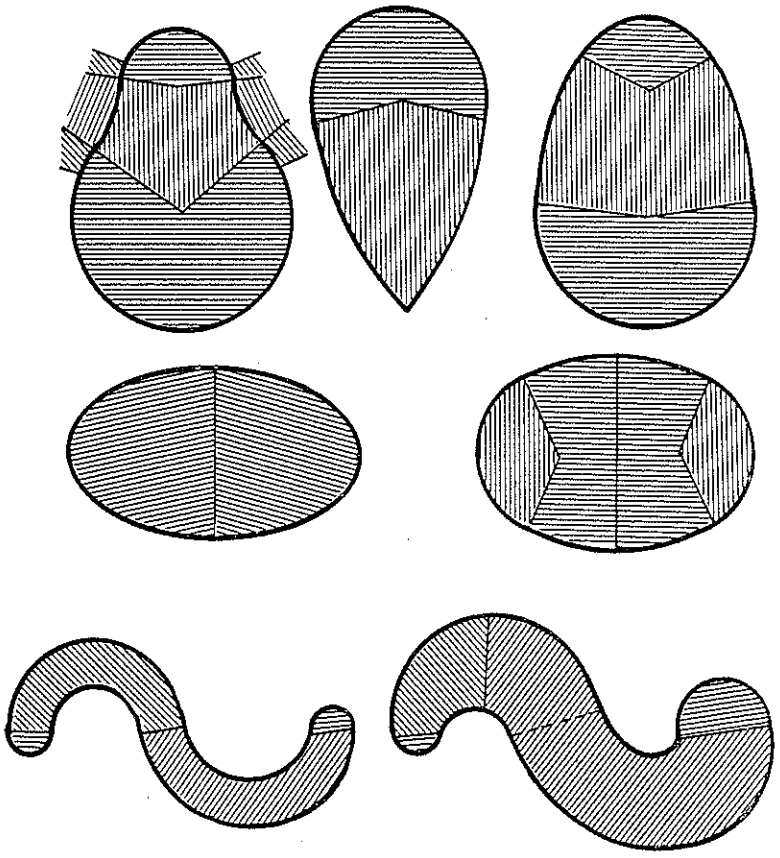

FIG. 35. Description of some simple curved objects by A-morphs. Where a finite contact point exists, a finite length of arc is described by a node A-morph. Note in the "pear" that the central section can be described by both object and ground A-morphs. The curvature property is simpler in the ground, since it exists as a location. The representation in the object requires precise velocity and acceleration information. The bottom shows two objects with the same sym-ax. Note that as A-morphs are arbitrarily defined here, the width and axes A-morphs need not terminate at the same places.

considered to be a rectangle in perspective, the central wedge angle is a good cue for the perspective plane.

We go on to curvilinear objects in Fig. 35. Shown there are an egg, a pear and a pointed object—all single symmetric arc A-morphs with end node A-morphs. The ellipse shown in Fig. 24 consists of two cups of finite entry width and zero contact entry angle, joined by a 2,0 chain, with no disc contact there. As discussed earlier, this defines a whole range of curves which need not be mathematical ellipses at all. The "three-sided" objects of Fig. 10 have a common feature in a three-bulb. The difference between the tapered and bulbous ends are represented by a number of A-morphs: a pinch, the presence or absence of other bulbs, and the presence of flares in the "arms". Figure 36 shows a dog, an early scheme for obtaining the computer derived sym-ax and its computer reconstruction, using simple node locations and straight linear wedges connecting them (Philbrick, 1968). The tail has been straightened out by this simplistic procedure, but it already gives a good reconstruction. The objects of Fig. 37 make the important point that the equivalence of exact boundary and sym-ax descriptions, no longer holds when objects are caricatured by A-morphs and B-morphs.

The A-morphs introduced thus far extract only the width data at distinguished points and the consistency of curvature changes between these points. This is obviously inadequate for many important shape distinctions. For example, all triangles have the same A-morphs. It is necessary to add some quantitative descriptors or at least some relations among them. Fortunately the zeros and extrema of many quantities exist at distinguished points. At the points, we may extract sym-width, sym-angle, object angle, object orientation, contact angle(s), symmetric and asymmetric curvature. Over intervals between these points, we may wish to extract distance and sym-length. The simplest way of using them is to specify the desired quantities and some tolerance limits on them. Unfortunately, that does not even allow for the identification of similar (size changed) objects. To accomplish this it is necessary to introduce global properties of the object. Normalization may be done by setting the object width to unity. We then divide lengths and widths, and multiply curvature by the object width. Other object parameters may also be used. A more compact scheme would simply order these parameters by size. Such a scheme has been proposed for intensity and color estimation models in human vision (Land, 1964). A less global method would represent the increase, decrease or sameness of the parameter between adjacent A-morphs. Some of this information is inherent already in the width and curvature A-morphs. It could be extended to include A-morph sym-length, for example. The advantage of such an "adjacent A-morph" scheme is in its minimizing the level at which parameter comparisons must be made.

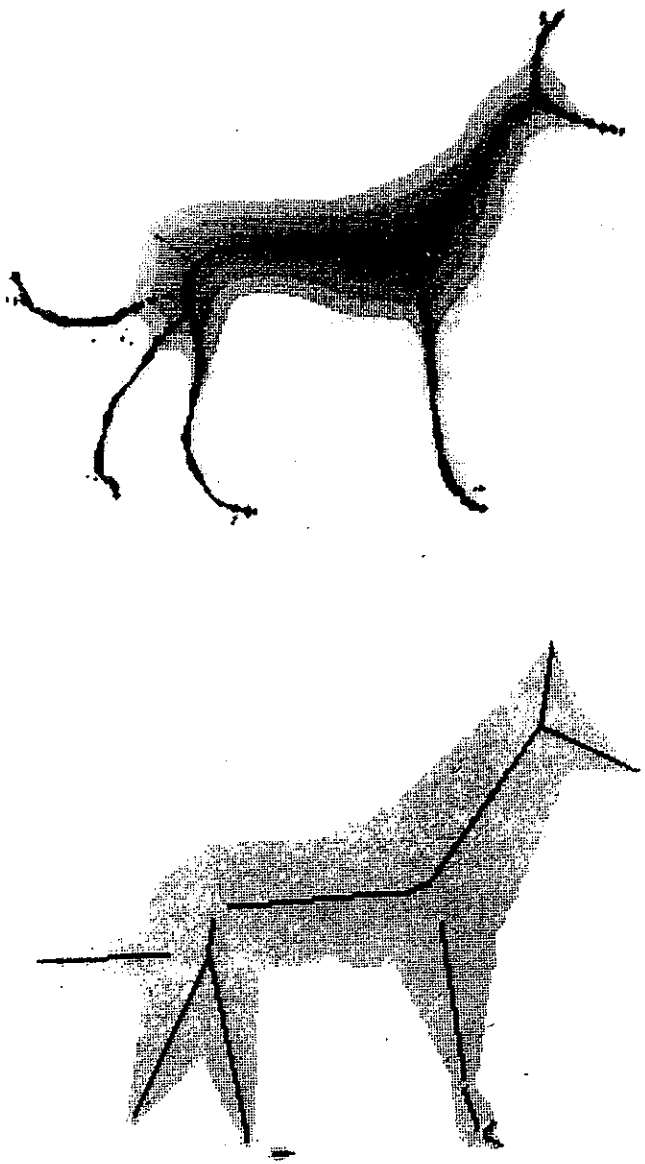

FIG. 36. Early automatic extraction of sym-function by computer. Beneath is shown a reconstruction using node A-morphs and wedges as connections (Philbrick, 1966).

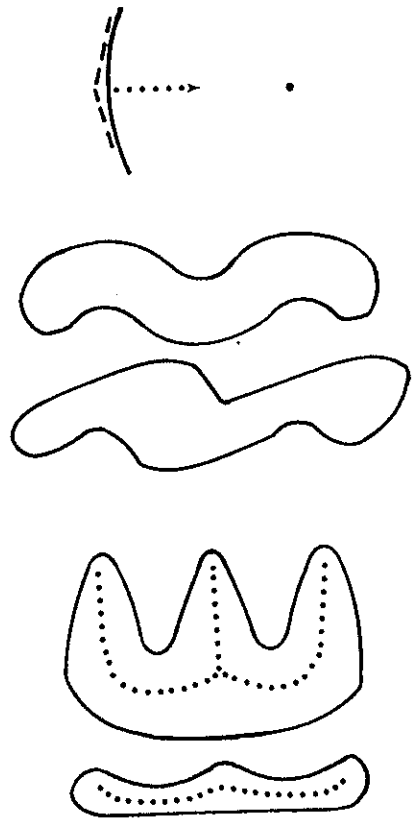

FIG. 37. Some un-nice geometric properties. Top shows a boundary curve and a sym-ax point for it. Whereas a smooth perturbation of the boundary will generally lead to a smooth perturbation of the sym-ax, the insertion of a corner, no matter how small, will generate a new sym-ax from that point discontinuously. To avoid this discontinuous change, the velocity of the sym-ax must be considered part of its description. The middle and bottom figures show the *lack* of invariance of the sym-ax description with affine transformation. The middle shows the skewing of a worm. The bottom shows the compression of a "fingered" object.

Other pairwise statements seem extremely useful also. One may wish to make a statement about the curvature of a tadpole boundary, one that is most simply imaged in a ground A-morph, in relation to the tadpole size. This would define the maximum boundary curving of the normalized tadpole. One may wish instead to make such a maximum bending condition

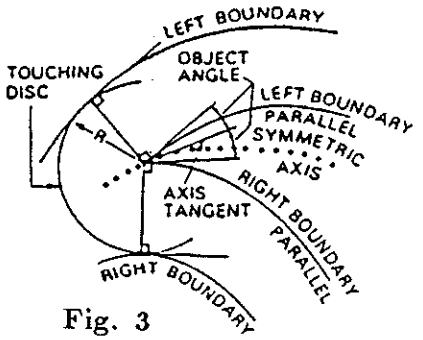

Fig. 3

- object width : 2 × radius
- object opening : ∠ angle of tangents at touching points
- axis curvature
- object curvature : change of object angle along axis

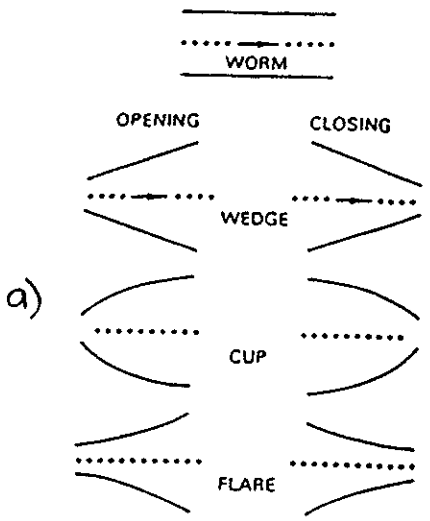

a)

b)

Object curvature (k_β)	Object angle (β): Neg	Zero	Pos
Pos	Opening flare		Closing cup
Zero	Opening wedge	Worm	Closing wedge
Neg	Opening cup		Closing flare

c)

from \ to	Worm	Open wedge	Close wedge	Open cup	Close cup	Open flare	Close flare
Worm	·	·	·	·	×	×	·
Opening wedge	·	·	·	×	·	×	·
Closing wedge	·	·	·	·	×	·	×
Opening cup	•	×	·	·	•	×	·
Closing cup	·	·	×	·	·	·	×
Opening flare	·	×	·	×	·	·	·
Closing flare	·	·	×	·	×	•	×

× indicates transition always possible.
• indicates transition at width extreme.

Elements of shape language

a) elementary shape types

b) mathematical characterization

c) transition rules

a)

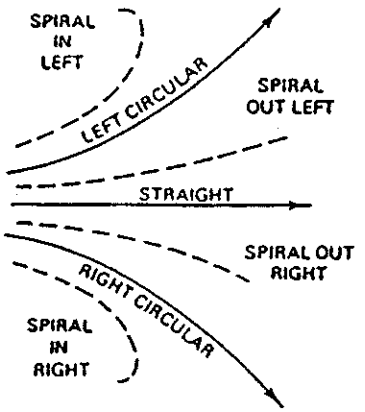

b)

Axis curvature change (k_s)	Axis curvature (k_a): Pos	Zero	Neg
Pos	Spiral in left		Spiral out right
Zero	Circular left	Straight	Circular right
Neg	Spiral out left		Spiral in right

Elements of shape language

a) elementary types of axis changes

b) mathematical characterizatio

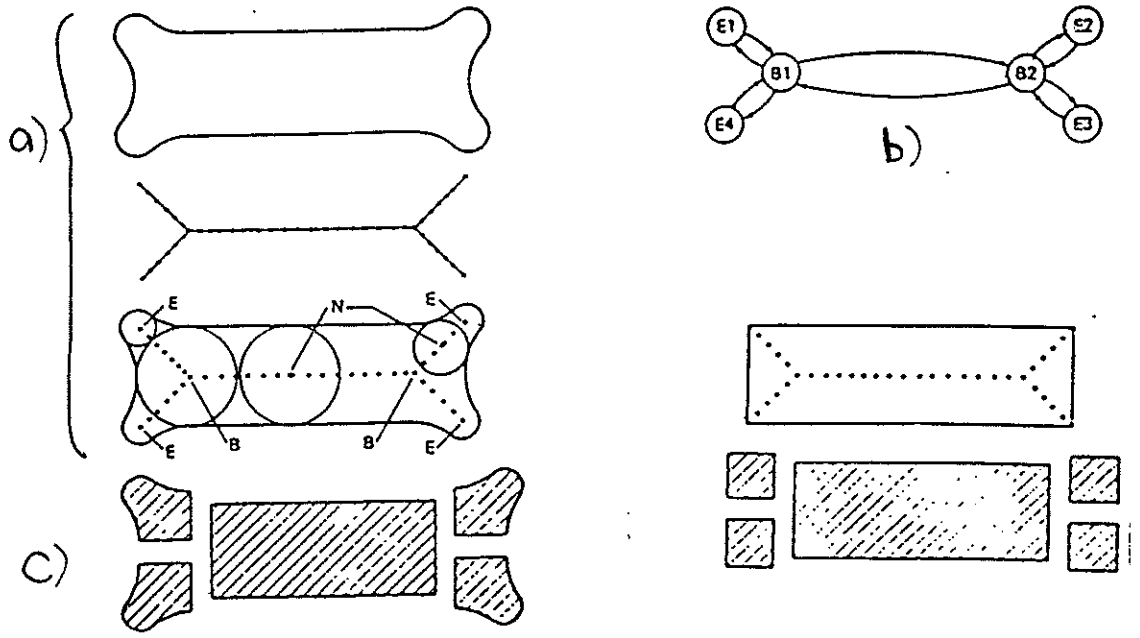

a) Elementary Shape Decomposition
b) relational graph structure
c) non-intuitive decomposition

Blum: Types of skeleton points

		# touching points
E	endpoint	1
N	normal point	2
B	knot point	≥ 3

04/15/2006

(15)

Implementations:

binary figures:

a) distance transform → ridge extraction

b) disk inclusion test : centers of set of largest disks completely inside the object and not covered by another disk

c) Voronoi Diagram of contour points → simplification, pruning

d) propagation of contour ⇒ shocks

$$\frac{\partial C}{\partial t} = \alpha \vec{N}$$

e) mathematical morphology

grey-level images:

f) wave equation on level sets

g) mathematical morphology (grey-level)

h) "cores" of multi-scale representation (Pizer, Morse, Eberly, Furst etc.)