Chapter 5

SKELETONS AND SEGMENTATION OF SHAPES

Jayant Shah
Department of Mathematics
Northeastern University
shah@neu.edu

Abstract The emphasis of the two approaches described in this chapter is on determining shape skeletons in the presence of noise. The parts of the skeleton which are most robust in the presence of noise correspond to the parts of the shape where its width varies slowly. In contrast, singular points of the skeleton such as junctions tend to be sensitive to noise. The strategy therefore is to concentrate on determining the robust parts and fill in the details afterwards as needed. The first approach takes advantage of the geometry of the distance function while the second analyzes the local symmetries of the level sets of a smooth analog of the distance function.

1. INTRODUCTION

Shapes are commonly described in terms of their parts. For example, at a superficial level, a cat may be described in terms of its head, legs, tail and body. A finer description would include how the parts are connected or related to each other. As a further refinement, one might describe individual parts themselves in terms of subparts. The cat's leg may be described in terms of the foot with its claws, the ankle, the knee and the portions of the leg in between. In effect, what we have is a hierarchical description of the shape. To this qualitative picture, one may then add quantitative information such as the length, thickness and curvature of the leg or its parts. Statistical measures may be associated with this quantitative data to account for variation in shapes within individual categories.

Shape skeletons provide an intuitive method for generating such descriptions. In practice, determination of shape skeletons on a discrete

grid has turned out to be a difficult problem. An information-preserving shape skeleton tends to contain a huge number of branches, most of which turn out to be related to noise or insignificant details. It would be worthwhile to avoid spending too much computational effort on determining unimportant details. It is easy to determine the skeleton of the parts of the shape with slowly varying width, namely, ribbons of approximately uniform width in the 2D case and such sheets in the 3D case. These are the shape parts which are usually the most important features of the shape. It is the accurate determination of the branches of the skeleton corresponding to noise and secondary shape details which are computationally difficult. In this chapter, we present two approaches for determining shape skeletons in which the emphasis is on determining quickly and accurately parts of the shape with slowly varying width and then fill in the missing details as needed. The amount of information lost depends on the choice of these details. Such a strategy is analogous to the one used by Zhu (1999) for 2D shapes; he finds the axes of the ribbon-like parts of the shape by minimizing a functional defined in terms of chords and extends these axes in an optimal way to form junctions.

The first approach follows the traditional definition of shape skeletons. For 2-dimensional shapes, it is the locus of the centers of maximal circles inscribed inside the shape. For 3D shapes, circles in the definition are replaced by spheres. The pruning strategy is based on the "object angle" which is defined at every point of the shape. The success of this approach rests on the fact that special properties of the distance function allow us to use large neighborhoods for accurate determination of the object angles, at least at the nonsingular points of the skeleton. The second approach is essentially a method for finding ridges in the graph of a smooth function defined over an n -dimensional domain and may be of interest for applications besides determination of shape skeletons. A smooth function is interpolated inside the shape boundary and a smoothed version of the shape skeleton is recovered from the locus of local symmetries of the level sets of the interpolated function. The smooth interpolation attenuates noisy branches of the skeleton. A major advantage of this approach is that the locus of local symmetries is easy to define in all dimensions and is easy compute as zero-crossings of certain differential quantities. It is stratified with strata ranging in codimension from 1 to $n-1$ for an n-dimensional shape. The difficulty arises when we try to extract the actual skeleton from the locus of local symmetries. Another drawback is that, since the computation uses differential quantities of order up to three, its accuracy deteriorates in narrow parts of the shape. Both approaches determine the branches

of the skeleton corresponding to parts with slowly varying widths quite accurately.

2. APPROACH 1

In the following, a shape is simply a connected open bounded set $D \subset \mathbf{R}^n$; its boundary will be denoted by ∂D. We denote the topological shape skeleton by K which is assumed to have the regularity properties usually assumed in practice. For example, its codimension is at least 1. The term "sphere" will mean a circle in the case of 2D shapes, a 2D sphere in the case of 3D shapes and a hypersphere in the case of shapes of dimension ≥ 3. Similarly, the term "plane" is used to mean a line, a plane or a hyperplane, depending on the context.

2.1. GRAY SKELETON

Construction of the skeleton by drawing maximal spheres at all points within the shape is clearly impractical. An alternative is to use the distance function, ρ:

$$\rho(P) = \max_{X \in \partial D} dist(X, P) \tag{1}$$

It may be computed very quickly by solving the eikonal equation $\|\nabla\rho\| = 1$ using, for example, the fast marching algorithm of Sethian (1996). The topological shape skeleton, K, is the discontinuity locus of the gradient of the distance function ρ. Associated with each point P on the skeleton is the radius of the maximal sphere centered at P. The shape boundary may be recovered as the envelope of the maximal spheres centered on K. Alternatively, K is the locus of the singular points of the level sets of ρ. The level sets of ρ are the fronts of the "grassfire", the usual analogy used to visualize the construction of the shape skeleton. The shape is imagined to be filled with dry grass and a fire is started at the shape boundary. The time of arrival of the grassfire front at a point equals the distance ρ of that point from the shape boundary. The shape skeleton is the locus of points where fronts from two or more directions meet. When the colliding fronts come from opposite directions, the point of their collision can be determined fairly accurately. As the angle between the front normals decreases, their collision becomes numerically more and more difficult to detect. It is even more difficult to locate the singular points of the skeleton where more than two fronts come together.

The distance function ρ has many special properties. One of these is that in the complement of $\partial D \cup K$, the gradient lines of ρ are straight lines and $||\nabla\rho|| = 1$. If P is a point not on $\partial D \cup K$, the gradient line of ρ passing through P connects P to a unique point Q on the shape

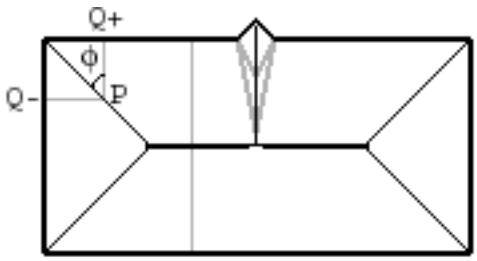

Figure 5.1. Geometry of the distance function

boundary nearest to P. The vector $\overrightarrow{QP}$ is the radius of the maximal sphere at P. It is normal to the fronts advancing to P. If P is a point of K where there are exactly two boundary points, Q^+ and Q^-, nearest to P, the maximal sphere centered at P touches the shape boundary only at Q^+ and Q^-. Let φ denote one half of the angle between the vectors $\overrightarrow{Q^+P}$ and $\overrightarrow{Q^-P}$; it is called the object angle and takes values between 0 and $\frac{\pi}{2}$ (Fig. 5.1). If the shape boundary is smooth at Q^+, then the vector $\overrightarrow{Q^+P}$ is orthogonal to the shape boundary and φ is the angle between the chord $\overrightarrow{Q^+Q^-}$ and the tangent plane at Q^+. Giblin and Kimia (2000) have shown that the bisector of the angle between $\overrightarrow{Q^+P}$ and $\overrightarrow{Q^-P}$ is tangent to the shape skeleton at P. Since $\|\nabla^{\pm}\rho\| = 1$ where $\nabla^{\pm}\rho$ denotes the gradients in the directions $\overrightarrow{Q^{\pm}P}$, we have

$$\cos\varphi = \frac{d\rho}{ds} \tag{2}$$

where ds denotes the infinitesimal arc-length along the skeleton in the direction of increasing ρ. It is the projection of the gradients $\nabla^{\pm}\rho$ onto the plane tangent to the skeleton at P. Thus, the larger the value of φ, the slower the rate of change in the width of the shape. We also have

$$\sin\varphi = \frac{1}{2}\left\|\nabla^{+}\rho - \nabla^{-}\rho\right\| \tag{3}$$

which is one half the jump in $\nabla\rho$ across the shape skeleton. The larger the object angle φ, the larger the jump in $\nabla\rho$, and the easier it is to detect the corresponding portion of the shape skeleton.

In the complement of $\partial D \cup K$, $\nabla\rho$ is continuously differentiable and we set $\varphi = 0$ there.

We define the gray skeleton, Γ, of the shape by setting its value equal to $\sin\varphi$. The gray skeleton is now defined everywhere except on the set

J of points of K where the maximal sphere touches the shape boundary at more than two points. We assume that J has codimension ≥ 2. Our strategy is to determine K/J from the gray skeleton in the complement of $\partial D \cup J$ and then extend it to parts of J which lie in the closure of K/J. This strategy still leaves out special cases such as discs, circular cylinders and balls where the closure of K/J does not contain all of J. Therefore, it would be useful to extend the definition of the gray skeleton to all of K. An elegant method to define such an extension due to Dimitrov et al. (2003) is based on the notion of average flux. These authors use the divergence theorem and calculate the average flux as a limit of a surface integral over spherical neighborhoods as the neighborhood size shrinks to zero. A possible alternative is to integrate the laplacian of ρ in the sense of distributions. Only the singular part of the laplacian (including its poles) contributes to the average flux and since the singular part depends only on the object angles, it may be computed accurately using large neighborhoods. In this chapter, since Γ is used only for the purpose of extracting K, it is sufficient to use a simple approximation to define Γ at points in J. At a point P in J, we pick two points, Q^+ and Q^-, among the set of boundary points nearest to P such that the object angle φ determined by the vectors $\overrightarrow{Q^+P}$ and $\overrightarrow{Q^-P}$ is the largest possible and set $\Gamma(P) = \sin\varphi$. Note that the set J may be quite large because there may be numerous intersecting branches produced by noise. The suggested method amounts to either ignoring or lumping together contribution of less important branches to.

To estimate φ, we need to determine the gradient directions $\overrightarrow{Q^+P}$ and $\overrightarrow{Q^-P}$ which are the directions in which the directional derivative of ρ is maximum and equals 1. We compute the directional derivatives of ρ in all *inward* radial directions at every point P in a given shape and determine the directions in which it is nearly equal to 1. On a discrete grid there are only finitely many radial directions available depending upon the size of the neighborhood N which determines the numerical accuracy of φ. The key point is that since the gradient lines are straight lines and $||\nabla\rho|| = 1$, we may use fairly large neighborhoods. If there are exactly two rays from P along which the directional derivative of ρ approximately equals 1, then the point must be on the smooth part of the skeleton and the angle between the two rays may be used to calculate the value of φ at that point. The presence of more than two gradient directions at a point indicates a point in J and the value of φ there is determined as indicated above. If there is a single gradient direction at a point, we set $\varphi = 0$. The algorithm in detail is as follows:

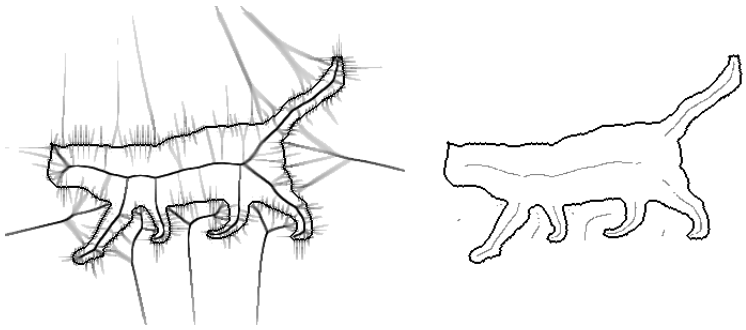

Figure 5.2. Left: Gray skeleton. Right: result after thresholding

Step 1: Scan the boundary ∂N of a neighborhood N of P . If X is a point belonging to ∂N, the derivative of ρ in the direction $\overrightarrow{XP}$ is

$$D_X(\rho) = \frac{\rho(P) - \rho(X)}{\left|\overrightarrow{XP}\right|} \tag{4}$$

Step 2: Determine the local maxima of $D_X(\rho)$ along ∂N. Among these maxima, choose the ones which are approximately equal to 1 within the tolerance determined by the size of N.

Step 3: If there are two or more local maxima chosen in *Step 2*, calculate the object angle between rays corresponding to all possible pairs of maxima and choose the largest value.

The gray skeleton of the shape of a cat is shown in Fig. 5.2 (left). The darker the shade, the higher the value of φ. The neighborhood at each point was chosen adaptively. At most points, the size of the neighborhood used was 15×15 pixels. If a 15×15 neighborhood at a point did not fit inside the shape, then its size was gradually decreased until it did. If even a 3×3 neighborhood did not fit, then the point was declared ineligible to be on the skeleton. As expected, the skeleton includes numerous extraneous branches which need to be pruned. Notice that the smaller the value of φ, the fuzzier the gray skeleton, reflecting the fact that the smaller the angle between the front normals, the harder it is to locate precisely the point where the fronts intersect.

2.2. PRUNING

Fig. 5.2 (right) shows the result obtained by thresholding out all points with $\varphi < 75^o$. This is a fairly high threshold, suitable for finding ribbon-like portions of the shape and the result is a set of disconnected skeleton branches. The corresponding ribbons may be identified

by drawing chords or maximal circles from the points with $\varphi \geq 75^o$. If the threshold is set too low, the skeleton will include branches belonging to noise. The following algorithm gets around this difficulty by extending the branches of the thresholded skeleton along the gray skeleton into thicker parts of the shape, possibly forming junctions, without picking up extraneous branches belonging to noise. An alternate pruning method based on homotopically shrinking the shape boundary onto the shape skeleton is described by Bouix and Siddiqi (2000).

Step 1: Choose two thresholds $\overline{\theta}$ and $\underline{\theta}$ for angle φ with $\overline{\theta} > \underline{\theta}$.

Step 2: Threshold the gray skeleton by $\overline{\theta}$. This eliminates irrelevant branches so that what remains corresponds to significant parts of the shape.

Step 3: Extend the branches of the thresholded skeleton in the direction of increasing ρ provided that φ remains greater than $\underline{\theta}$ throughout. The effect is to extend the skeleton branches from narrow necks towards the more blobby parts of the shape. As the branch is extended, it may encounter junctions with noise or protrusion related branches, but *these branches are not followed since they ascend from the junctions towards the shape boundary.* If the value of $\overline{\theta}$ is set too high, some numerical difficulties may be encountered during this step due to inaccuracies in the values of ρ and the discreteness of the grid. If the value of $\overline{\theta}$ is too high, ρ decreases very slowly along the skeleton branch; for example, if $\overline{\theta} = 75^o$, $\frac{d\rho}{ds} = \cos 75^o = 0.258$. Consequently, along the actual directions of descent permitted by the grid, the rate of decrease in ρ may be smaller still. Therefore, in searching for the lower values of ρ, the size of the neighborhood should be adjusted in order to ensure that it includes a direction of decreasing ρ.

Fig. 5.3 shows several examples of skeletons obtained by this procedure with $\overline{\theta} = 75^o$ and 60^o and $\underline{\theta} = 10^o$.

An example of a 3D shape is shown in Fig. 5.4. It is a multiheaded figure created from MRI slices of a head. A visualization of its complete skeleton, including all the noise is shown in Fig. 5.5: On the left is the complete skeleton by itself while on the right, the same with the original figure superimposed. What is seen is the external skeleton whose main feature is the four sheets emanating from the four concavities. The internal skeleton is obscured by the external skeleton. Fig. 5.6 shows the result of pruning. To see internal details, one has to view the skeleton either section by section or layer by layer. Figs. 5.7 and 5.8 depict several cross-sections of the gray skeleton (left column), the pruned skeleton with $\overline{\theta} = 70^o$ and $\underline{\theta} = 45^o$(middle column), and the pruned skeleton with $\overline{\theta} = 70^o$ and $\underline{\theta} = 30^o$ (right column). The points on the gray skeleton within 2 voxels of the shape boundary were removed before thresholding

as this portion was judged to be too noisy to be relevant. Fig. 5.7 depicts 5 successive vertical sections of the heads from front to back. Fig. 5.8 depicts 7 successive horizontal sections proceeding from the top of the triple heads to the "neck". Notice the effectiveness of pruning. With $\overline{\theta} = 70°$and $\underline{\theta} = 45°$, the pruned skeleton consists of 3 connected components, one inside the shape and two outside. (Extremely short components were removed) With $\overline{\theta} = 70°$ and $\underline{\theta} = 30°$, the pruned skeleton has a single connected component inside the shape and a single component outside.

As long as $\underline{\theta}$ is sufficiently low, the resulting skeleton is a connected set since in theory, the skeleton is connected if $\underline{\theta} = 0^o$. Since what is noise or an extraneous feature must depend on the context, the values of thresholds $\overline{\theta}$ and $\underline{\theta}$ must also depend on the context. The procedure is not too sensitive to the choice of $\overline{\theta}$ and $\underline{\theta}$ except around certain critical values. Topological and geometric changes in the skeleton of a shape at critical values of $\overline{\theta}$ and $\underline{\theta}$ reveal critical features of the shape. For instance, in the case of the rectangle in Fig. 5.1, if $\underline{\theta} > 45°$, the skeleton does not include the diagonal branches; the remaining skeleton, namely the main axis, represents the shape as a pure ribbon. If $\overline{\theta}$ and $\underline{\theta}$ are set sufficiently low, the skeleton will include the diagonal branches as well as the branch emanating from the protrusion.

2.3. SEGMENTATION FUNCTIONAL

The skeleton obtained by the pruning method described in the last section may still have small gaps, especially if the lower threshold $\underline{\theta}$ is fairly high. It may also contain short extraneous branches. Even if $\underline{\theta}$ is sufficiently low so that the skeleton is connected, its boundary in the case of 3D shapes may have a tattered appearance. An alternative is to determine the skeleton by minimizing a functional. In the case of 2D shapes, the functional is defined as

$$E(K_o) = \beta\,|n_{K_o}| + \int_{K_o} (c - \alpha) \tag{5}$$

where K_o is the pruned skeleton, to be determined as a subset of the complete skeleton K defined as the set of points where $\sin\varphi > 0$, n_{K_o} is the number of connected components of K_o, $c = 1 - \sin\varphi$ is the cost function, α is a cost threshold and β is the penalty assigned to each component of K_o. In the case of 3D shapes, the functional is defined as

$$E(K_o) = \beta\,|\partial K_o| + \int_{K_o} (c - \alpha) \tag{6}$$

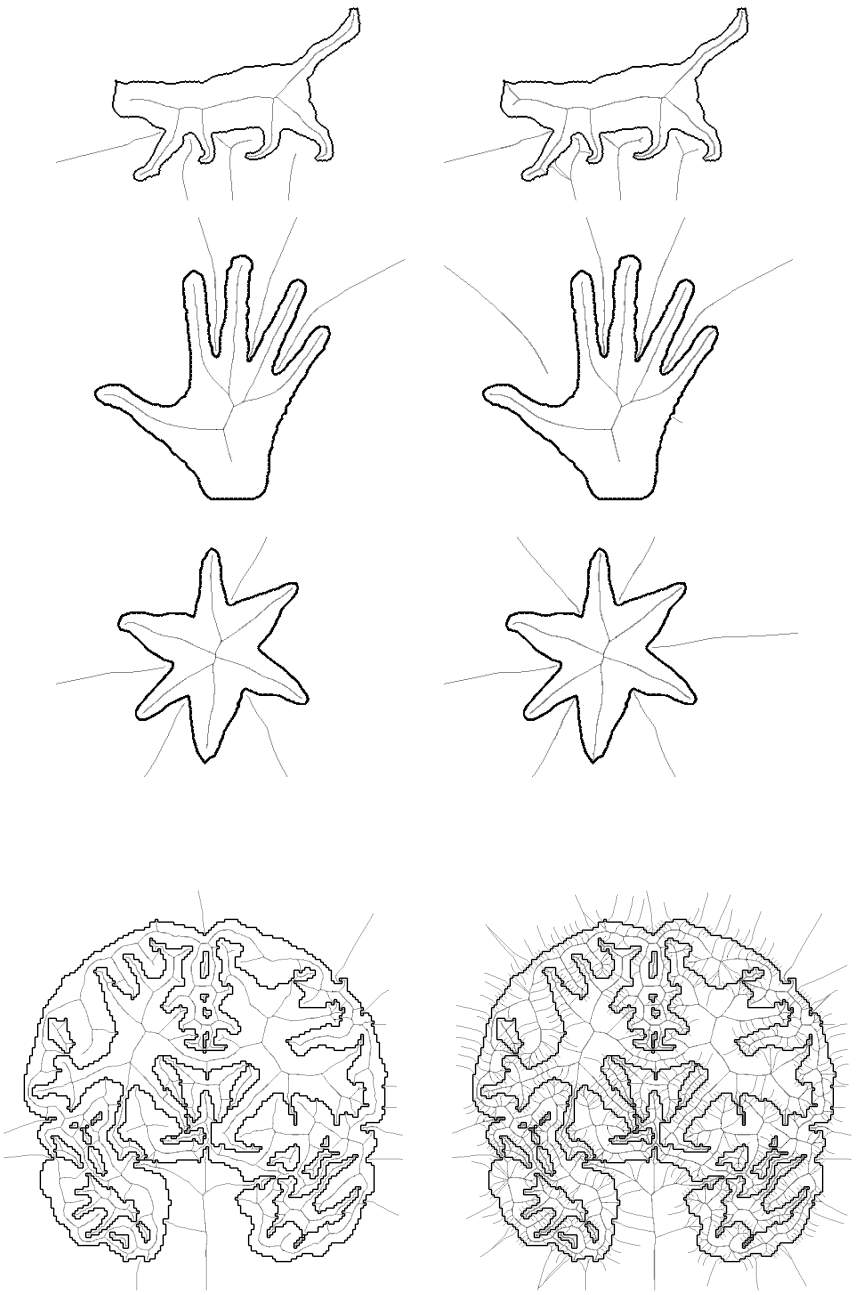

Figure 5.3. Examples of Skeletons. Left: $\overline{\theta} = 75^o$ Right: $\overline{\theta} = 60^o$ (right)

x

Figure 5.4. A 3D image

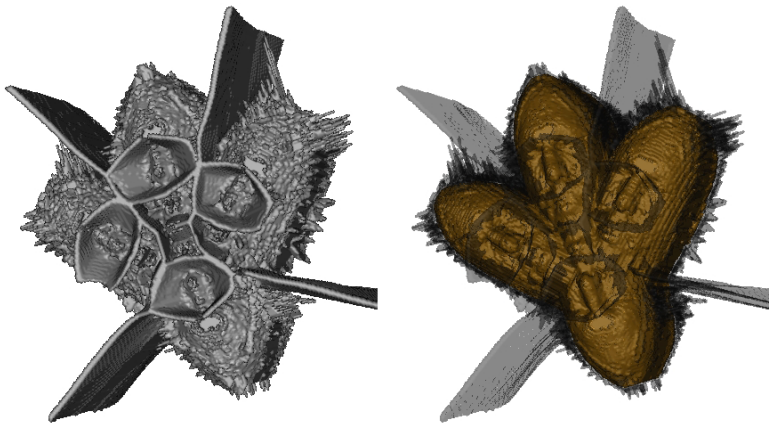

Figure 5.5. Complete Skeleton

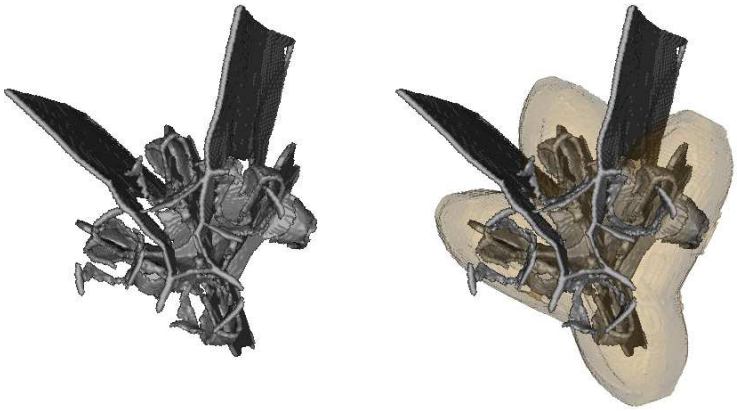

Figure 5.6. Pruned Skeleton

where $|\partial K_o|$ is the length of the boundary of K_o. If $\beta = 0$, E is minimized by setting K_o equal to the set of points of K where $c \leq \alpha$. Positive values of β permit (a) pruning of short components of K_o even if all the points on a short component have $c \leq \alpha$ and (b) inclusion of points in K_o with $c > \alpha$ if it results in the elimination of a short gap in K_o or shortening the boundary of K_o. At present, it is not clear how to minimize the functional in the 3D case. In the case of 2D shapes, the minimizing K_o satisfies two necessary conditions which may be used to reduce $E(K_o)$ with initial K_o as the set of points of K where $c \leq \alpha$. To derive these conditions, we rewrite $E(K_o)$ in the 2D case. If X is a connected component of K_o, let L_X be its the length and $\overline{c}_X$ its average cost:

$$\overline{c}_X = \frac{\int_X c ds}{L_X} \tag{7}$$

Then

$$E = \alpha \sum_{i=1}^{n_K} L_{X_i} \left[\frac{\overline{c}_{X_i}}{\alpha} - \left(1 - \frac{\gamma}{L_{X_i}} \right) \right] \tag{8}$$

where $\gamma = \beta/\alpha$. If X be a component of K_o and E is minimum, then

$$\frac{\overline{c}_X}{\alpha} \leq 1 - \frac{\gamma}{L_X} \tag{9}$$

otherwise, E can be reduced by omitting X. For example, if X has length $\leq \gamma$ and average cost $\overline{c}_X > 0$, then E can be reduced by eliminating X from K_o. That is, γ is a lower bound on the lengths of the components of K_o.

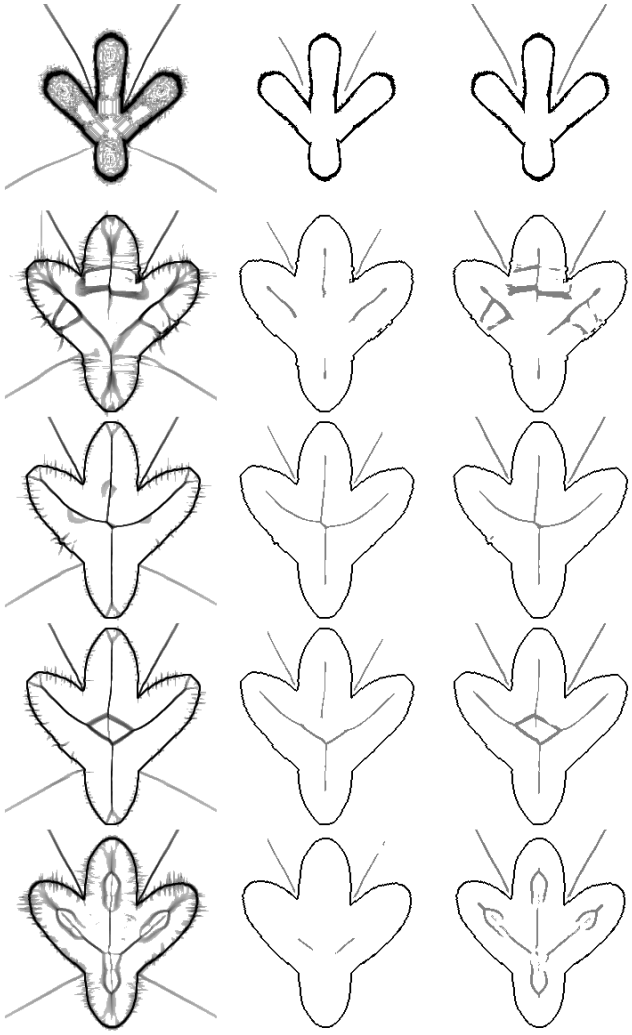

Figure 5.7. Gray skeleton and pruned Skeletons: Vertical sections

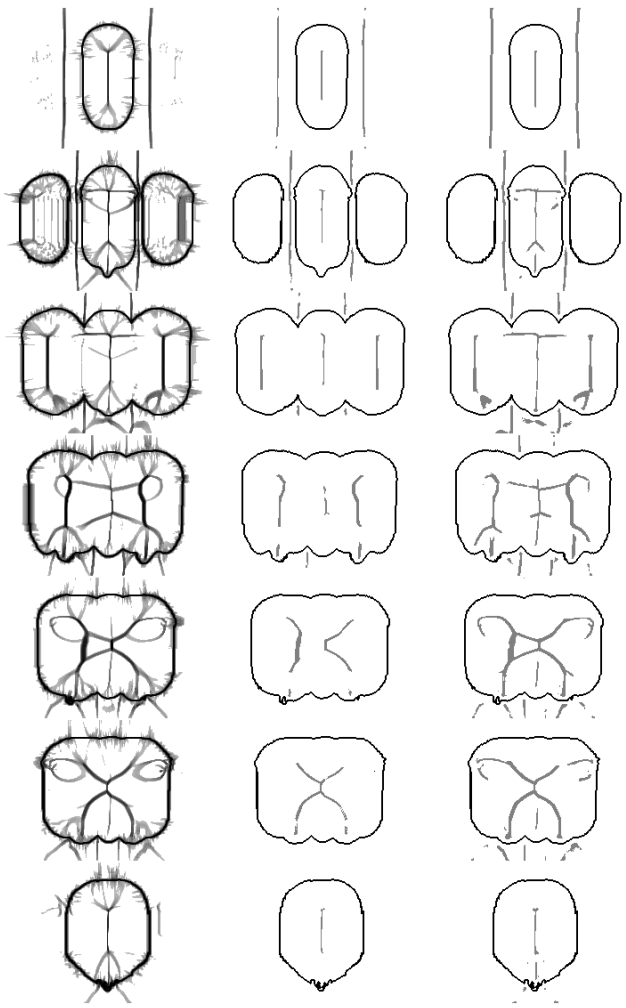

Figure 5.8. Gray skeleton and pruned skeletons: Horizontal sections

If X is a path in $K \backslash K_o$ bridging a gap between two components of K_o, let L_X be its the length and $\overline{c}_X$ its average cost. If X is adjoined to K_o, then the increase in E is given by the formula

$$\delta E = \int_X c ds - \alpha L_x - \beta = \alpha L_X \left[\frac{\overline{c}_X}{\alpha} - \left(1 + \frac{\gamma}{L_X} \right) \right] \tag{10}$$

since bridging the gap by X reduces the number of components of K_o by one. Hence if E is minimum, we must have

$$\frac{\overline{c}_X}{\alpha} \geq 1 + \frac{\gamma}{L_X} \tag{11}$$

The necessary conditions (9) and (11) provide a strategy for eliminating short branches and short gaps:

Step 1: Choose a threshold α. Let initial K_o be equal to the set of points on K at which $c \leq \alpha$.

Step 2: Eliminate all gaps in K_o that can be bridged by paths which violate the necessary condition (11).

Step 3: Remove all components of the modified K_o which violate the necessary condition (9).

Illustrative examples of segmentation obtained by this procedure with $c = \cos\varphi$ and $\gamma = 2$ pixels are shown in Fig. 5.9. On the left is the case where $\alpha = \cos 75^o$ and on the right, $\alpha = \cos 60^o$. Protrusions of the shape may be identified by drawing chords connecting the pairs of boundary points nearest to the endpoints of the protrusion branch (Fig. 5.10). The shape may be further segmented by drawing chords through the saddle points.

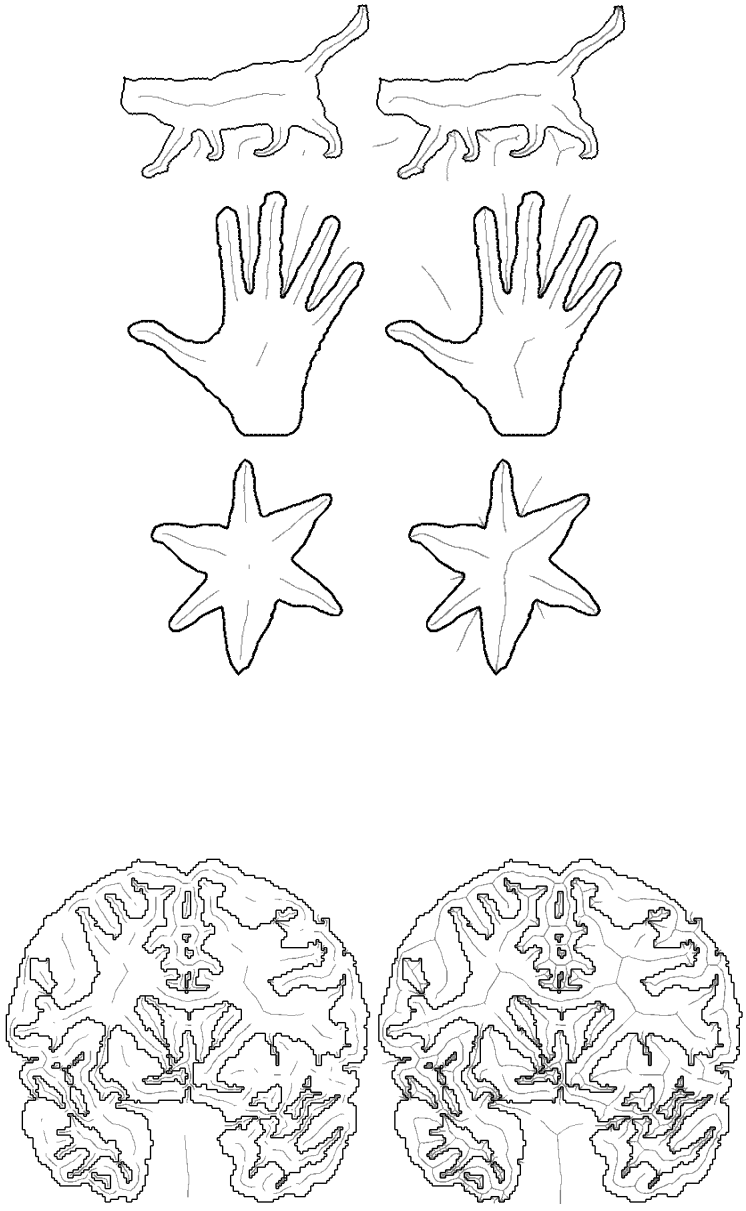

Figure 5.9. Examples of axes of ribbon-like parts found using the segmentation functional. Left: $\overline{\theta} = 75^o$, Right: $\overline{\theta} = 60^o$

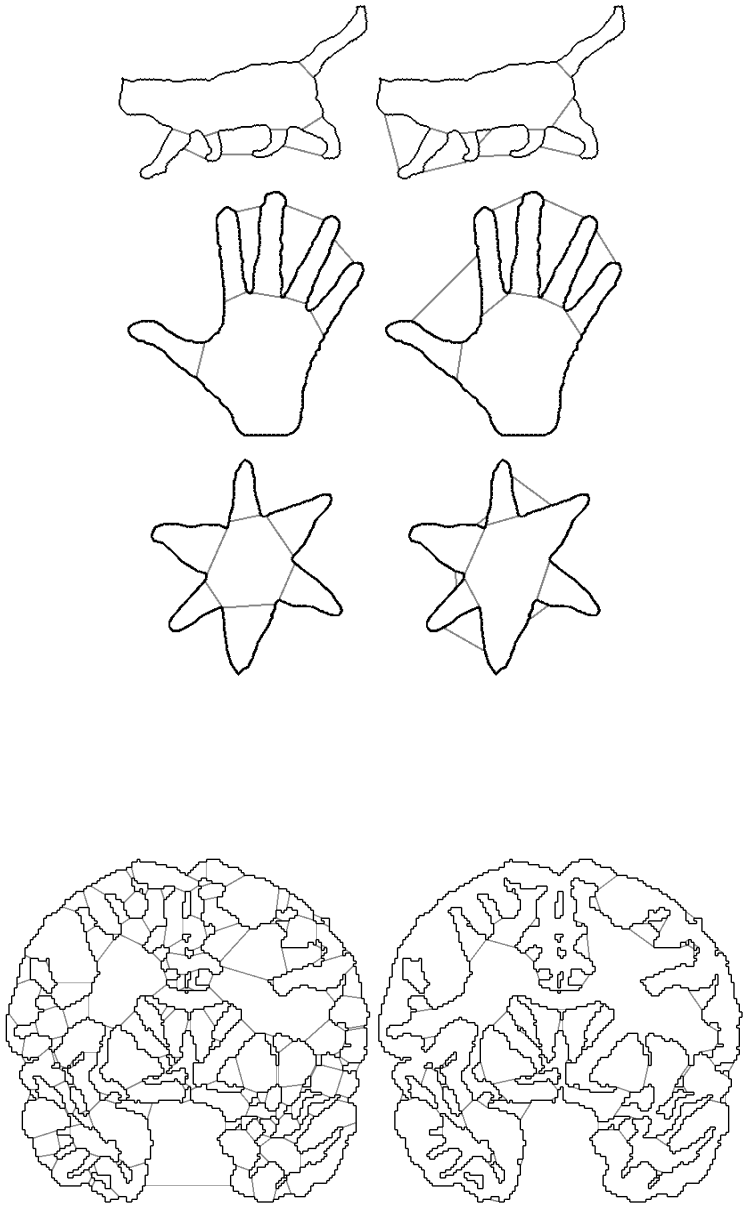

Figure 5.10. Protrusions (limbs) by the first method (§5.2)
Left: $\overline{\theta} = 75^o$, Right: $\overline{\theta} = 60^o$

3. APPROACH 2

The level sets of the distance function may be thought of as evolution of the shape boundary moving with a constant speed equal to one. The skeleton branches are traced out by the singular locus of the evolving front. This suggests an alternate approach based on smoothing the level sets of the distance function. Consider the case of a square shape. The level curves of the distance function are squares as well. Suppose we smooth the corners of the level curves. Then we have a well-defined flow field inside the square, defined by the normals to the level curves. The trajectories near the diagonals of the square have to travel much longer distance to reach the center than the trajectories starting at the midpoints of the sides of the square. Therefore, the speed of the evolving curve has to be faster near the corners than near the principal axes of the square. The faster speed near the corners corresponds to the fact that the absolute gradient of the distance function is less than one along the skeleton. Therefore, we expect the smoothed skeleton to be the locus of points where the gradient of the smoothed distance function is minimum along the level curves. The number of such points on a level curve and hence the degree of noise suppression will depend on the extent to which the level curve is smoothed. Since shallow corners of the level curves will smooth out before the sharp ones, noise related branches will disappear quickly while the branches associated with large object angles will persist.

An obvious way to smooth the level curves would be to smooth the distance function. Another possibility is introduce smoothing as the boundary curve evolves. The dynamics of the flow defined by the smoothed level curves in the example outlined above suggests that we consider evolution of the boundary curve in which we augment the constant speed by a component proportional to the curvature. The difficulty is that the evolving curve may backtrack and hence the later level curves may cross the earlier ones. The time of arrival may thus fail to define a single-valued function over the shape domain. As an alternative, we choose a model based on an idea developed by Modica and Mortola (1977) for deriving an approximation of the perimeter of a set by minimizing an elliptic functional. The level curves of the minimizing function, v, in the 2D case behave like the usual model of curve evolution. In the following, we develop an approach based on the analysis of the local symmetries of the level sets of the function v. The analysis below does not depend on the choice of a particular v. An advantage of our choice of v is that it provides a link between the problem of determining shape skeletons and the problem of image segmentation [Tari et al. (1997)].

As level cruves evolve and become progressively more circular, branches of the skeleton terminate. Consequently, the smoothed skeleton determined by this method is usually not connected. Nonetheless, it is sufficient for the purpose of locating the main parts of the shape and determining the hierarchical relationship among them. This is sufficient for applications like object recognition as demonstrated recently by Aslan and Tari (2005).

This approach was developed jointly with Sibel Tari.

3.1. MODICA-MORTOLA APPROXIMATION OF THE BOUNDARY

Define the functional

$$E_\sigma(v) = \int_D \left[\sigma \left\|\nabla v\right\|^2 + \frac{(v-\chi_{\partial D})^2}{\sigma}\right] \tag{12}$$

$$= \int_D \left[\sigma \left\|\nabla v\right\|^2 + \frac{v^2}{\sigma}\right] \tag{13}$$

with the boundary condition $v = 1$ along ∂D. The characteristic function $\chi_{\partial D} = 1$ along ∂D and $= 0$ elsewhere.

We have

$$\lim_{\sigma \to 0} \inf E_\sigma(v) = volume(\Gamma) \tag{14}$$

The minimizer of E_σ satisfies the elliptic differential equation

$$\nabla^2 v = \frac{v}{\sigma^2} \tag{15}$$

with boundary conditions $v = 1$ along the boundary of ∂D. The parameter σ plays the role of a nominal smoothing radius. This is a linear differential equation and its numerical solution may be obtained by standard methods such as the finite element method.

3.2. 2-DIMENSIONAL SHAPES

3.2.1 Analogy with curve evolution. When σ is small compared to the local width of the shape and the local radius of curvature of ∂D, the level curves of v locally capture the smoothing of ∂D by curve evolution. As shown by Mumford and Shah (1989) in Appendix (3), when σ is small,

$$v(x,y) = \sigma\left(1 + \frac{\sigma\kappa(x,y)}{2}\right)\frac{\partial v}{\partial \eta}(x,y) + O(\sigma^3) \tag{16}$$

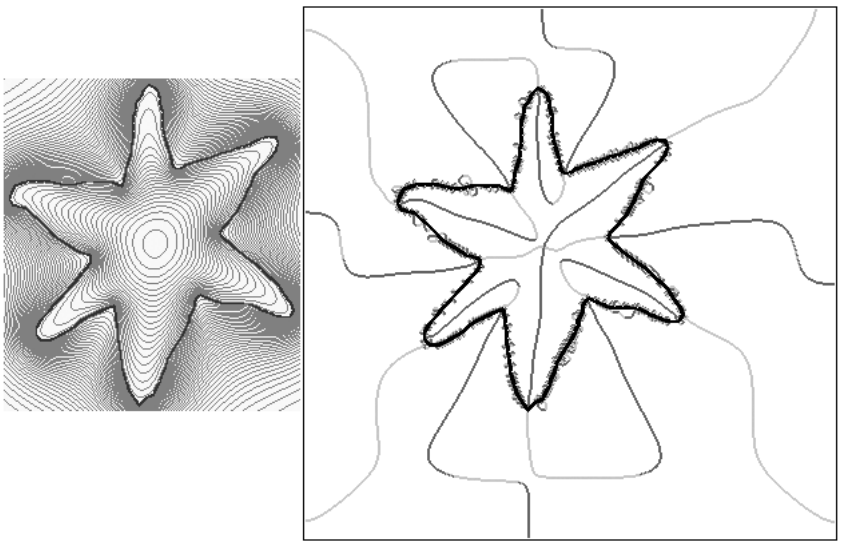

Figure 5.11. Left: Level Curves of v Right: S_1^+

where $\kappa(x, y)$ is the curvature of the level curve passing through the point (x, y) and η is the direction of the gradient. $\kappa(x, y) = v_{\xi\xi}/\|\nabla v\|$ where $v_{\xi\xi}$ is the second derivative of v in the direction of the unit vector ξ tangent to the level curve. To see that near the shape boundary, level curves of v mimic the curve evolution moving with a speed consisting of a constant component and a component proportional to curvature, imagine moving from a level curve to a level curve of v along the normals. A small change of δv in the level requires movement

$$\delta r \approx \frac{\sigma}{v}(1 + \frac{\sigma\kappa}{2})\delta v \tag{17}$$

where r denotes the arc length along the gradient lines of v. Define time τ such that $\frac{d\tau}{\sigma^2} = \frac{dv}{2v}$. Then

$$\frac{dr}{d\tau} \approx \frac{2}{\sigma} + \kappa \tag{18}$$

3.2.2 Local Symmetries and Shape Skeletons. We now analyze the structure of the level curves of v. Fig. 5.11 (left) shows the level curves of v inside a starlike shape. As expected, the level curves are more separated from each other along the apparent medial axes of the protruding arms than they are in the neighborhood of indentations. The tips of protrusions are in some sense furthest away from the apparent

center of the shape. The distance between two adjacent level curves, $v = v_0$ and $v = v_0 + dv$ is $\frac{dv}{\|\nabla v\|}$ which is maximum when $\|\nabla v\|$ is minimum along the level curve. Therefore, we define the locus of local symmetries, S_1 as the set of points where

$$\frac{d\,\|\nabla v\|}{ds} = 0 \tag{19}$$

where s is the arc-length along the level curves of v. Since $\frac{d\|\nabla v\|}{ds} = \eta \cdot H\xi$ where H is the Hessian of v, η is the unit vector in the direction of ∇v and ξ is the unit vector tangent to the level curve, $\frac{d\|\nabla v\|}{ds} = 0$ means that the Hessian of v along the axes of local symmetry is diagonalized when expressed in terms the local frame, η and ξ. In particular, the gradient vector ∇v is an eigenvector of the Hessian at points on the axes of local symmetry. Geometrically, at a point P on an axis of local symmetry, the level curve is asymptotically a conic section whose one of the principal axes coincides with the gradient vector. As explained above, along the middle of protrusions, the distance between adjacent level curves is the greatest, that is, $\|\nabla v\|$ is minimum along the level curve. Let S_1^+ denote the set of zero-crossings of $\frac{d\|\nabla v\|}{ds}$ where $\frac{d^2\|\nabla v\|}{ds^2}$ is positive and let S_1^- denote the set of zero-crossings of $\frac{d\|\nabla v\|}{ds}$ where $\frac{d^2\|\nabla v\|}{ds^2}$ is negative. Let $S_1^0 = S_1 \backslash \left(S_1^+ \cup S_1^-\right)$. The expressions for the derivatives $\frac{d\|\nabla v\|}{ds}$ and $\frac{d^2\|\nabla v\|}{ds^2}$ in terms of global coordinates x and y are

$$\frac{d\,\|\nabla v\|}{ds} = v_{\eta\xi} \text{ and } \frac{d^2\,\|\nabla v\|}{ds^2} = v_{\eta\xi\xi} + \frac{v_{\xi\xi}(v_{\xi\xi} - v_{\eta\eta})}{\|\nabla v\|} \tag{20}$$

where

$$v_{\eta\xi} = \frac{\{(v_y^2 - v_x^2)v_{xy} - v_x v_y (v_{yy} - v_{xx})\}}{\|\nabla v\|^2} \tag{21}$$

$$v_{\xi\xi} = \frac{\left\{v_y^2 v_{xx} - 2v_x v_y v_{xy} + v_x^2 v_{yy}\right\}}{\|\nabla v\|^2} \tag{22}$$

$$v_{\eta\eta} = \frac{\left\{v_x^2 v_{xx} + 2v_x v_y v_{xy} + v_y^2 v_{yy}\right\}}{\|\nabla v\|^2} \tag{23}$$

$$v_{\eta\xi\xi} = \frac{\left\{v_x v_y^2 v_{xxx} + v_y(v_y^2 - 2v_x^2)v_{xxy} + v_x(v_x^2 - 2v_y^2)v_{xyy} + v_x^2 v_y v_{yyy}\right\}}{\|\nabla v\|^3} \tag{24}$$

The quantity $v_{\eta\xi}/\|\nabla v\|$ is the curvature of the gradient lines. Therefore, another geometric characterization of the axes of local symmetries is that they are the locus of points where the gradient lines of v are straight like

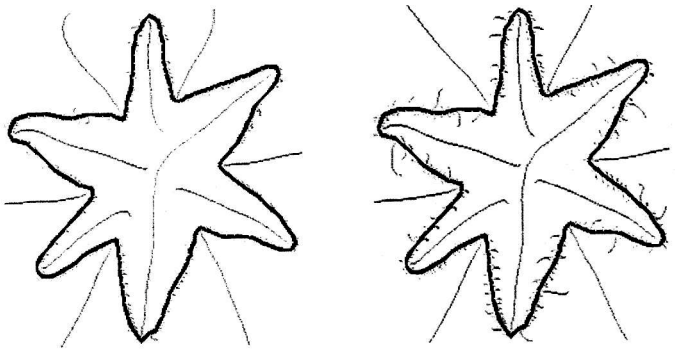

Figure 5.12. Left: S_1^+ with $\rho = 128$ Right: S_1^+ with $\rho = 8$

the gradient lines of the distance function. The loci S_1^+ and S_1^- may be distinguished from each other without calculating $\frac{d^2\|\nabla v\|}{ds^2}$ by observing that as we follow S_1^+ in the direction of decreasing v, $\frac{d\|\nabla v\|}{ds}$ is positive to its right and negative to its left while the pattern is reversed along S_1^-.

Fig. 5.11 (right) illustrates the locus S_1^+ (dark) and S_1^- (light) in the case of the star-like figure. It suggests the following dynamics as the level curves of v of the starlike figure become more and more circular. The locus S_1^+ tracks evolution of its protrusions while S_1^- tracks the evolution of its indentations. During the evolution, a protrusion may merge with an indentation, joining a branch of S_1^+ with a branch of S_1^- and thus terminating both the branches. The branches near the shape boundary created by the noise in the boundary are short. The smaller the protrusion, the shorter the branch. The medial axes of two of the arms join up to define the longest ribbon that can be extracted from the shape. The construction described above depends on the choice of the smoothing parameter σ. Axes shown in Fig. 5.11 were obtained with $\sigma = 32$ pixels. (The size of the frame around the shape was 400×400 pixels.) Fig. 5.12 depicts the symmetry locus S_1^+ determined using $\sigma = 8$ and $\sigma = 128$. The larger values of σ cause faster termination of the protrusion axes.

Since a level curve of $\frac{d\|\nabla v\|}{ds}$ is tangent to a level curve of v at a point if and only if $\frac{d^2\|\nabla v\|}{ds^2} = 0$, the gradient of v is nonzero along $S_1^+ \cup S_1^-$. We orient each branch of $S_1^+ \cup S_1^-$ in the direction of *decreasing* v. The set S_1^0 is the union of two sets, S_0 and J: S_0 is defined by the equation $\|\nabla v\| = 0$ and J is defined by the equations $\frac{d\|\nabla v\|}{ds} = \frac{d^2\|\nabla v\|}{ds^2} = 0$. At

points in J, the symmetry axes are tangent to the level curves of v and should not be part of the shape skeleton.

Ideally, we would like S_1^+ to be orthogonal to the level curves of v as in the case of symmetric ribbons with a straight axis of symmetry. In this case, v is symmetric about the geometric axis and S_1^+ will coincide with the geometric axis; S_1^+ will be orthogonal to the level curves of v. As S_1^+ and S_1^- approach J, their behavior deviates further and further from this ideal. A measure of this deviation is the angle, ψ, between the vectors ∇v and $\nabla(\frac{d\|\nabla v\|}{ds})$ which may be used to threshold S_1^+ and S_1^-.

Computation of $\nabla\left(\frac{d\|\nabla v\|}{ds}\right)$ involves third order derivatives of v and may present difficulties where the shape feature is too small or the shape is too narrow. The approach of S_1^+ and S_1^- to J may be quite gradual which makes the numerical determination of J difficult. The join between a branch of S_1^+ and a branch of S_1^- may appear as an alternating series of points or sections of S_1^+ and S_1^- with small associated numerical values of $\frac{d^2\|\nabla v\|}{ds^2}$.

The set S_0 may be further subdivided into the set S_0^+ of elliptic points where the determinant of the Hessian of v is positive, the set S_0^- of hyperbolic points where it is negative and the set S_0^0 of parabolic points where it is zero. At an elliptic point, v has a local minimum and has the Taylor expansion of the form $a_{00} + a_{20}x^2 + a_{02}y^2 +$ *higher order terms*. By applying the definition of S_1^+ and S_1^- to this local expression, it is easy to see that at an elliptic point, there are two branches of S_1^+ in the direction of the maximum second derivative and two branches of S_1^- in the direction of the minimum second derivative. All four branches are directed towards the elliptic point.

At a hyperbolic point, v has the Taylor expansion of the form $a_{00} + a_{11}xy +$ *higher order terms* and calculations show that at a hyperbolic point, there are four branches of S_1 all of which belong to S_1^+. Hyperbolic points are of course saddlepoints in that two of these branches are directed away from the saddle point and two are directed towards it. It is the two branches of the latter type which make the extraction of the shape skeleton difficult. Let S_{1h}^+ denote the set of branches of S_1^+ which terminate at a saddle point. Since saddle points occur in the neck-like parts of the shape, branches of S_{1h}^+ do not start from the shape boundary, but from a point in J. An example of this can be seen in Fig. 5.11. The midpoints of the square frame are hyperbolic points. Four of the six branches of S_1^+ outside the star which start from the six concave corners change into short branches of S_1^- which in turn change into branches of S_1^+ terminating at the hyperbolic points.

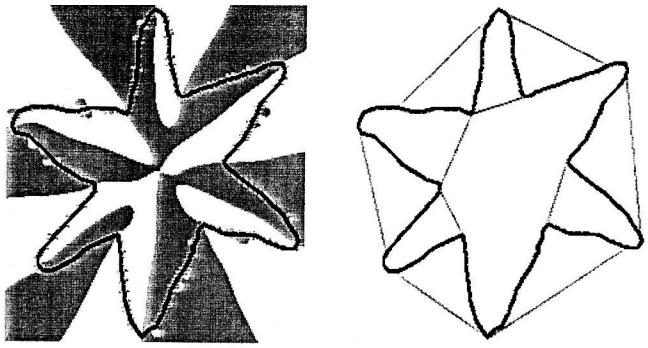

Figure 5.13. Left: $\frac{d\|\nabla v\|}{ds} > 0$ dark, < 0 light Right: Segmentation

It is numerically difficult distinguish a true parabolic point from a near-degenerate elliptic or a hyperbolic point. The medial axes of ribbons of approximately constant width are like a line of parabolic points and numerically they may exhibit a series of elliptic and hyperbolic points with associated short branches of S_1.

We call $S_1^+ \backslash S_{1h}^+$ the skeleton of the shape. A branch of the shape skeleton ending at an elliptic point is called a *main axis* while a branch ending in J is called a *protrusion axis*. Branches of $S_1^+ \backslash S_{1h}^+$ may be pruned by thresholding the angle ψ. As noted before, the shape skeleton is not connected.

3.2.3 Segmentation. The objective is to segment protrusions and necks by means of their medial axes. The basic idea is to find the two nearest points on the shape boundary from the terminal point of the protrusion axis, one on each side of the axis and connect these two points to segment the protrusion. To restrict the search to a suitable neighborhood of the protrusion axis, we use the fact that S_1 segments the shape and inside each connected component of $D \backslash S_1$, $\frac{d\|\nabla v\|}{ds}$ is either positive or negative. Segmentation of D in the case of the star-figure is shown in Fig. 5.13 (left). Each protrusion axis neighbors exactly two of these components. Terminate the protrusion axis either where it meets J^+ or where the angle ψ falls below a fixed threshold. From the terminal point of the axis, search for the nearest boundary point in the interior of each of the two components of $D \backslash S_1$ adjoining the axis. Segment the protrusion by connecting the two boundary points. Admittedly, the boundary points found in this way depend on where the terminal point of the protrusion axis is which in turn depends on the choice of σ. If the ends of the protrusion are marked by a sharp change in the local width of

the shape, the boundary points are insensitive to σ. An example where this is not true is the case of a parallelogram where the medial axes starting from the obtuse corners terminate in J^+. The obtuse corners are interpreted as protrusions. In this case, the smaller the value of σ, the longer the branch starting from the obtuse corners and the further away the segmenting chords from the obtuse corners.

Fig. 5.13 (right) shows the segmentation of the star figure. The shape is segmented from inside as well as from outside. As noted before, in its attempt to extract the longest possible ribbon, the algorithm disregards the approximate symmetry of the star and includes in the main ribbon two of what would normally be perceived as protrusions.

It is possible to segment shapes across necks by means of the associated saddlepoints. This is a more delicate construction since the medial axes of long narrow necks act like lines of parabolic points. The problem is how to avoid spurious saddlepoints arising from the numerical breakup of parabolic lines. A characteristic of such a situation is that there are small components of $D\backslash S_1$ adjoining the saddlepoint which do not extend to the shape boundary. Therefore, a hyperbolic point is called a *true saddlepoint* if it adjoins at least three components of $D\backslash S_1$ which touch the shape boundary. Find the two nearest boundary points, one on each side of the medial axis going through the saddle point, restricting the search for them to these components of $D\backslash S_1$. Segment the shape by a chord connecting the two boundary points.

Additional examples of shape segmentation are shown in Fig. 5.14. Each shape was scaled to make sure that it was nowhere less than 4 pixels wide. The approximate sizes of the frames around the shapes were: the star 208×240 pixels, the cat 208×300 pixels, the hand 280×290 pixels. All the quantities needed in the algorithm are computed using 3×3 neighborhoods except the sign of $d^2 \left\|\nabla v\right\| / ds^2$ which required 4×3 neighborhoods. In Fig. 5.14, the top row on the left shows the S_1^+ loci while the bottom row shows the corresponding segmented shapes. The example of brain segmentation shown on the right illustrates the case of a complex of shapes involving non-simply connected shapes and triple junctions. (The shape boundary is outlined by thick jagged lines while the segmentation lines are straight and thinner.)

The segmentation has the structure of a directed graph. Its set of vertices consists of the true saddlepoints and one vertex for each of the shape segments. If a segment X is a protrusion with its medial axis terminating at a point in J which is contained in another segment Y, then YX is an edge in the graph. (The direction of an edge is always in the direction of increasing v.) Each segmentation line corresponding to a saddlepoint is the common boundary between two shape segments.

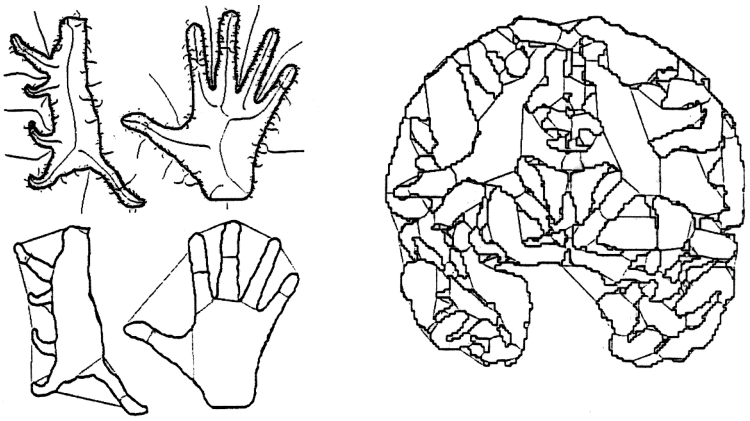

Figure 5.14. Examples of shape segmentation

The vertices corresponding to these two segments are connected to the vertex corresponding to the saddlepoint by edges directed towards the saddlepoint.

3.2.4 Grayscale Images. The function v is used by Ambrosio and Tortorelli (1992) to approximate the Mumford-Shah segmentation functional and provides a way to determine shape skeletons directly from grayscale images. The Mumford-Shah functional is

$$E_{MS}(u,B) = \alpha \int_{R\backslash B} \|\nabla u\|^2 + \int_R \beta(u-g)^2 + |B| \tag{25}$$

where R is a connected, bounded open subset of $\mathbf{R}^2$, g is the feature intensity, B is a curve segmenting R, u is the smoothed image $\subset R\backslash B$, $|B|$ is the length of B and α, β are weights. Its elliptic approximation is given by Ambrosio and Tortorelli as

$$E_{AT}(u,v) = \int_R \left\{ \alpha(1-v)^2 \|\nabla u\|^2 + \beta(u-g)^2 + \frac{\sigma}{2}\|\nabla v\|^2 + \frac{v^2}{2\sigma} \right\} \tag{26}$$

where v is the Modica-Mortola approximation of B as described above. The Ambrosio-Tortorelli approximation is usually implemented by gradient descent. We can then analyze the level curves of v by the method described above. A difficulty in applying this approach to determining shape skeletons directly from the grayscale images is that the value of the function v on the shape boundary depends on the local contrast; it is nearly equal to 1 when the contrast is high and may be much less than

1 or even near 0 if the contrast is low. Therefore, the method readily determines branches of the shape skeleton corresponding to parts of the shape with high contrast. Parts of the shape boundary where v is near zero behave like gaps in the shape boundary. See Tari et al. (1997) for more details.

3.3. SHAPES IN DIMENSION$\geq$ 3

3.3.1 Loci of partial symmetries. Extension of the definition of the axes of local symmetries in the case of 2-dimensional shapes to higher dimensions is not difficult. The link is provided by the following proposition.

Proposition: $||\nabla v||$ is stationary along a level hypersurface at a point P if and only if ∇v is an eigenvector of the Hessian H of v.

Proof: $||\nabla v||$ is stationary along a level hypersurface at a point P if and only if the derivative of $||\nabla v||$ in any direction tangent to the level hypersurface at P vanishes. That means that at P, $\nabla||\nabla v||$ cannot have a component tangent to the level hypersurface at P. In other words, the directions of $\nabla||\nabla v||$ and ∇v must coincide at P. That is, $\nabla||\nabla v||$ must be a multiple of ∇v. Since

$$\nabla||\nabla v|| = \frac{H\nabla v}{||\nabla v||} \tag{27}$$

it follows that the necessary and sufficient condition for P to be a stationary point is that

$$H\nabla v = c\nabla v \quad \textit{for some constant} \quad c \tag{28}$$

■

If ∇v is the vector $\{\alpha_1, \alpha_2, \cdots, \alpha_n\}$ and $H\nabla v$ is the vector $\{\omega_1, \omega_2, \cdots, \omega_n\}$, then we have $n-1$ equations

$$\frac{\omega_1}{\alpha_1} = \frac{\omega_2}{\alpha_2} = \cdots = \frac{\omega_n}{\alpha_n} \tag{29}$$

to determine the 1–dimensional locus of the stationary points. The Hessian at a stationary point is diagonalized if we choose the direction of the gradient vector as one of the local coordinates and choose the other coordinates appropriately in the hyperplane tangent to the level hypersurface. Thus the convexity of v is locally symmetric with respect to the vector ∇v.

To obtain more information about the shape, we look for partial symmetries. That is, we require the level hypersurface to be symmetric only about some linear space containing ∇v. We look for the smallest linear

subspace L, invariant under the action of H on $\mathbf{R}^n$ and which contains ∇v. Let Σ be the matrix whose columns are

$$\nabla v, H\nabla v, H^2\nabla v, \cdots, H^{n-1}\nabla v \tag{30}$$

Then L is spanned by the columns of Σ and its dimension equals the rank of Σ. If ∇v is an eigenvector of H, the rank of Σ is one. Therefore we define the partial symmetry locus S_k of dimension k as the locus of points where Σ has rank $\leq k$. We have a sequence of nested loci of successively increasing degree of symmetry:

$$D = S_n \supset S_{n-1} \supset \cdots \supset S_1 \supset S_0 \tag{31}$$

where S_0 is the locus of points where ∇v vanishes, making the rank of Σ equal to zero. S_{n-1} is the given by the vanishing of the determinant of Σ. At a point in S_k, v is locally symmetric with respect to L.

In the case of 3–dimensional shapes, the 1–dimensional locus S_1 is given by the equations

$$\frac{v_x v_{xx} + v_y v_{xy} + v_z v_{xz}}{v_x} = \frac{v_x v_{xy} + v_y v_{yy} + v_z v_{yz}}{v_y} = \frac{v_x v_{xz} + v_y v_{yz} + v_z v_{zz}}{v_z} \tag{32}$$

which may be rewritten as

$$v_x v_y \left(v_{xx} - v_{yy}\right) + v_{xy}\left(v_y^2 - v_x^2\right) + v_z\left(v_y v_{xz} - v_x v_{yz}\right) = 0 \tag{33}$$

$$v_y v_z \left(v_{yy} - v_{zz}\right) + v_{yz}\left(v_z^2 - v_y^2\right) + v_x\left(v_z v_{yx} - v_y v_{zx}\right) = 0 \tag{34}$$

$$v_z v_x \left(v_{zz} - v_{xx}\right) + v_{zx}\left(v_x^2 - v_z^2\right) + v_y\left(v_x v_{zy} - v_z v_{xy}\right) = 0 \tag{35}$$

As an example, in the case of a square prism, the locus S_2 consists of 3 orthogonal planes parallel to the prism faces and passing through its center, and 12 essentially planar sheets emanating diagonally from the 12 edges of the prism. The locus S_1 consists of the three principal axes of the prism passing through the centers of its faces, 12 lines starting from the midpoints of the 12 edges which are the intersections of the 3 principal planes and the 12 diagonal sheets, and 8 lines starting diagonally from the 8 corners. S_0 consists of a single point, the centroid [Tari and Shah (1998)].

3.3.2 Shape Skeletons. It is not clear what the generalization of the definition of the shape skeleton of a 2D shape should be in higher dimensions. Below, we outline some alternatives.

In the 2D case, the level curves of v were seen to approximate the evolution of the shape boundary moving with speed equal to the curvature of the level curve plus a constant. Therefore, the larger the curvature,

the smaller the gradient of v. In the starlike figure, the curvature is large inside the arms and their axes are part of the skeleton. Near the concave corners (indentation), the curvature is negative and the gradient of v is small. The axis of symmetry starting at a concave corner is eliminated. At a saddle point, the curvature of the level curves along the branch of S directed towards the saddle point has negative curvature and the corresponding axes of local symmetry eliminated as well. Since along an axis of local symmetry, the Hessian is diagonalized when expressed in terms of the local frame, η and ξ, its eigenvalues are given by $\lambda_1 = v_{\eta\eta}$ and $\lambda_2 = v_{\xi\xi}$ and the curvature of the level curve $\kappa = \lambda_2 / \|\nabla v\|$. This suggests a way to use the eigenvalues of the Hessian H to classify the axes of local symmetry in higher dimensions. Let L be the linear subspace spanned by the columns of Σ and $L^\perp$ be the linear subspace orthogonal to L. Then, analogous to the classification of S_0 in the 2D case, points on the symmetry locus may be classified according to the various combinations of the signs of the eigenvalues of H corresponding to the eigenvectors contained in $L^\perp$ [Tari and Shah (2000)]. A coarser measure may be defined by analogy with curve evolution in higher dimensions where the speed of evolution has a component proportional to the mean curvature of the level sets. The mean curvature is given by the formula

$$\mu = \nabla \cdot \frac{\nabla v}{\|\nabla v\|} = \frac{1}{\|\nabla v\|} \left[\nabla^2 v - \frac{\nabla v}{\|\nabla v\|} \cdot H \frac{\nabla v}{\|\nabla v\|} \right] \tag{36}$$

We adapt this formula as follow. At points in $S_k \backslash S_{k-1}$, the decomposition $L \oplus L^\perp$ may be used to put H in a block diagonal form consisting of a $k \times k$ block, Λ_k, and a $(n-k) \times (n-k)$ block, Ξ_{m-k}. The expression in the bracket in the formula is just the trace of Ξ_{m-1}. Therefore, at a point in $S_k \backslash S_{k-1}$, We define

$$\mu_k = \frac{trace(\Xi_{m-k})}{\|\nabla v\|} \tag{37}$$

and use μ_k to classify the loci of local symmetry. In the 3D case, at a point in $S_2 \backslash S_1$, $L^\perp$ is one-dimensional and spanned by the unit vector t:

$$t = \frac{\nabla v \times H \nabla v}{\|\nabla v \times H \nabla v\|} \tag{38}$$

Therefore,

$$\mu_2 = \frac{t \cdot Ht}{\|\nabla v\|} \tag{39}$$

The symmetry loci have to be pruned as in the case of 2D shapes. The most important part of the skeleton is its $(n-1)$-dimensional part.

Let θ denote the determinant of Σ. The $(n-1)$-dimensional locus S_{n-1} is the level set $\theta = 0$. As in the 2D case, let ψ be the angle between the $\nabla\theta$ and ∇v and use ψ to threshold S_{n-1}.

Bibliography

Ambrosio, L. and V. M. Tortorelli: 1992, 'On the Approximation of Functionals depending on Jumps by Quadratic, Elliptic Functionals'. *Boll. Un. Mat. Ital.* [xxv]

Aslan, C. and S. Tari: 2005, 'An Axis Based Representation for Recognition'. In: *Proceedings of the IEEE International Conference on Computer Vision*. [xviii]

Bouix, S. and K. Siddiqi: 2000, 'Divergence-Based Medial Surfaces'. In: *Proceedings of the European Conference on Computer Vision*. Dublin, Ireland, pp. 603–618. [vii]

Dimitrov, P., J. N. Damon, and K. Siddiqi: 2003, 'Flux Invariants for Shape'. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. [v]

Giblin, P. J. and B. B. Kimia: 2000, 'A formal Classification of 3D Medial Axis Points and Their Local Geometry'. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Vol. 1. pp. 566–573. [iv]

Modica, L. and S. Mortola: 1977, 'Il limite nella Γ-convergence di una familiglia di funzionali ellittici'. *Boll. Un. Mat. Ital.* **14-A**, 526–529. [xvii]

Mumford, D. and J. Shah: 1989, 'Optimal Approximations by Piecewise Smooth Functions and Associated Variational Problems'. *Comm. Pure Appl. Math.* **42**(5), 577–684. [xviii]

Sethian, J. A.: 1996, 'A Fast Marching Level Set Method for Monotonically Advancing Fronts'. *Proceedings of the National Academy of Sciences, USA* **93**, 1591–1595. [iii]

Shah, J.: 2001, 'Segmentation of Shapes'. In: *Workshop on Scale-Space and Morphology*. [-]

Shah, J.: 2005a, 'Grayscale Skeletons and Segmentation of Shapes'. *Computer Vision and Image Understanding* **99**(1), 96–109. [-]

Shah, J.: 2005b, 'Skeletons of 3D Shapes'. In: *The Fifth International Conference on Scale-Space and PDE Methods in Computer Vision*. [-]

Tari, S. and J. Shah: 1998, 'Local Symmetries of Shapes in Arbitrary Dimension'. In: *Proceedings of the IEEE International Conference on Computer Vision*. Bombay, India. [xxvii]

Tari, S. and J. Shah: 2000, 'Nested Local Symmetry Set'. *Computer Vision and Image Understanding* **79**(2), 267–280. [xxviii]

Tari, S., J. Shah, and H. Pien: 1997, 'Extraction of Shape Skeletons from Grayscale Images'. *Computer Vision and Image Understanding* **66**, 133–146. [xvii, xxvi]

Zhu, S. C.: 1999, 'Stochastic Jump-Diffusion Process for Computing Medial Axes in Markov Random Fields'. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **21**(11). [ii]