Skeletons of 3D Shapes*

Jayant Shah
Mathematics Department, Northeastern University, Boston MA
email: shah@neu.edu

Abstract

A new method for determining skeletons of 3D shapes is described. It is a combination of the approach based on the "grass-fire" technique and Zhu's approach based on first finding portions of the shape where its width is approximately constant. The method specifically does not require presmoothing of the shape and is robust in the presence of noise. In an appendix, a method based on variational calculus is formulated for determining pruned, smoothed shape skeletons by minimizing a functional.

1 Introduction

Although a large number of papers have appeared on ways to determine shape skeletons, until recently [3, 7, 8, 11, 18], these papers have dealt with 2D shapes [4, 6, 10, 12, 15, 16, 17, 19]. Extension of these techniques to 3D shapes is nontrivial since skeletons for 3D shapes consist of intersecting 2-dimensional surfaces [7]. This paper presents a generalization of the method developed in [17] to 3D shapes. The emphasis is on developing a computationally robust method for determining shape skeletons in the presence of noise and the inevitable numerical inaccuracies inherent in computation on a discrete grid.

The usual definition of the skeleton of a 3D shape is that it is the locus of the centers of maximal spheres contained in the shape. If the radii of these spheres are recorded at the corresponding points on the skeleton, the shape can be fully recovered as an envelope of the spheres centered on the skeleton with radii recorded on the skeleton. Construction of the skeleton by drawing maximal spheres at all points within the shape is clearly impractical and determining the tangency between a sphere and a noisy shape boundary is problematic. An alternative is the "grassfire" technique in which 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. Alternatively, it is the discontinuity locus of the gradient of the distance function ρ. When the colliding fronts

*This work was supported by NIH Grant I-R01-NS34189-08.

come from opposite directions, the point of their collision can be determined fairly accurately. As the angle between the front normals decreases, it becomes numerically more and more difficult to detect their collision. It is even more difficult to locate the singular points of the skeleton where more than two fronts come together. Another difficulty is that the skeleton preserves the noise present in the shape boundary in the form of numerous extraneous branches which must be pruned. It is also necessary to identify and prune branches corresponding to shape features which are deemed irrelevant for the task at hand.

An alternative to the approach described above is the one proposed by Zhu [19] for 2D shapes. A 2D shape is viewed as a collection of ribbons which are glued together. A ribbon is a part of the shape where the shape width is approximately constant, so that, after some smoothing, the opposite boundaries of the ribbon are nearly parallel. Chords may then be found which are approximately orthogonal to the smoothed boundary. Zhu minimizes a functional designed to determine optimal chords which are as closely normal to the smoothed boundary as possible. The functional also incorporates terms such as a penalty for parts of the shape in which proper chords cannot be found and a penalty for having many short ribbon segments instead of a few long ones. The relative importance of these desirable properties depends on the values of the parameters in the functional. The centers of the chords define the medial axes of the ribbons which are then extended to form junctions in an optimal way. Note that these junctions are not necessarily the points where the maximal discs touch the shape boundary at two or more points. Instead, their location is governed by criteria such as a minimal number of junctions and angles subtended by the branches at the junctions. Parts of the shape which are neither ribbon-like nor associated with junctions are ignored (pruned). It is not clear how to determine a good initial approximation to start the minimization process, how to tune the numerous parameters involved in the functional and how to generalize the method to 3D shapes.

The approach proposed in this paper is a combination of the two approaches described above and consists of three steps. The "gray skeleton" of a shape is defined in §2 by associating each point inside and outside the shape with a numerical value which is a monotonic function of the angle at which the fronts intersect. One half of this angle is what is called the object angle in the literature [3,5]. We define this angle to be zero at points away from the skeleton. The first important step is an accurate calculation of the object angle at every point based on the observation that normals to the fronts may be determined by using fairly large neighborhoods. An alternative method for calculating the object angle based on the divergence theorem is given in [5]. Corresponding to Zhu's chords, we now have a set of boundary points associated with each point P on the gray skeleton where the front normals at P intersect the shape boundary.

Gray skeletons contain numerous points associated with noise and its pruning constitutes the second step which is described in §3. Parts of the shape with slowly varying width are determined by thresholding the gray skeleton. Then, branches of the thresholded skeleton are extended along the gray-skeleton in such a way that they join up without picking up extraneous branches caused by noise.

The method is illustrated by means of an example in §4. An alternative approach to pruning the gray skeleton by contracting the shape boundary towards the gray skeleton homotopically is described in [16].

Due to noise, the boundary of the skeleton obtained by the method described above may have a tattered appearance. The skeleton may also have gaps and short isolated branches. Regularization of the skeleton is the third step. In §5, we propose an approach based on minimization of a functional. The 2D version of this functional has been implemented in [17]; however, its 3D version involves regularization of the skeleton boundary which is a space curve and has yet not been implemented.

Zhu's functional includes a term which penalizes wiggles in the skeleton. In the appendix, we present an approach based on variational calculus for smoothing (and pruning) shape skeletons by minimizing a functional which is analogous to functionals used for segmenting gray-scale images.

2 Gray Skeleton

In this section, we define the gray skeleton Γ of a shape from which the topological skeleton K is extracted. A shape is simply an open bounded set D; its boundary will be denoted by ∂D. The topological shape skeleton K is assumed to have the regularity properties usually assumed in practice. For example, its dimension is at most 2. We start with the distance function

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

which may be computed quickly, for example, by Sethian's fast marching algorithm [1] which solves the eikonal equation $\|\nabla\rho\| = 1$ by the "upwind" finite difference scheme. We estimate the angle between the normals to the intersecting fronts exploiting the property that the gradient lines of ρ are straight lines except where they cross the shape skeleton [14]. The gradient line of ρ passing through a point P off the shape skeleton connects P to the point Q on the shape 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^-. Exactly two fronts intersect at P and the angle between their velocity vectors is equal to the angle between the vectors $\overrightarrow{Q^+P}$ and $\overrightarrow{Q^-P}$ (Fig. 1). Let φ denote one half of the angle between the vectors $\overrightarrow{Q^+P}$ and $\overrightarrow{Q^-P}$; it takes values between 0 and $\frac{\pi}{2}$. 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^+.

It is shown in [7] 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

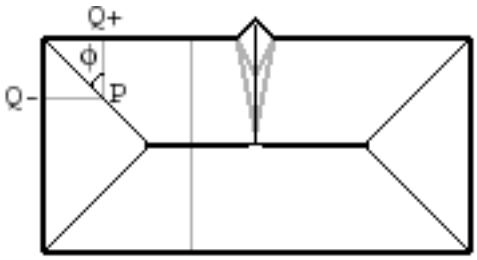

Figure 1: Geometry of the distance function

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 the bisector. 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.

Let $\varphi = 0$ in the complement of $\partial D \cup K$ where $\nabla\rho$ is continuously differentiable. We define the gray skeleton, Γ, of the shape by letting its value equal to $\sin\varphi$. The gray skeleton is defined everywhere except on the set J of points of K where the maximal sphere touches the shape boundary at more than two points. (J is assumed to have dimension ≤ 1. Our strategy is to determine K/J from the gray skeleton in the complement of $\partial D \cup J$ and then extend it over parts of J which lie in the closure of K/J . This strategy still leaves out special cases such as 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, Damon and Siddiqi [5] is based on the notion of average flux. These authors use the divergence theorem and calculate the average flux as a limit of the surface integral of $\nabla\rho$ over spherical neighborhoods as the neighborhood size shrinks to zero. A possible alternative method for calculating the average flux 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 paper, the purpose of calculating Γ is to extract K from it and it is sufficient to use a simple approximation of Γ

at a 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 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.Since the gradient lines are straight lines and $||\nabla\rho|| = 1$, we may use fairly large neighborhoods. The algorithm in detail is as follows:

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

$$D_R(\rho) = \frac{\rho(P) - \rho(R)}{\left|\overrightarrow{RP}\right|} \tag{4}$$

Step 2: Determine the local maxima of $D_R(\rho)$ 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 is a single gradient direction at a point, we set $\varphi = 0$. If there are two or more maxima chosen in *Step 2*, calculate the object angle between rays corresponding to all possible pairs maxima and choose the largest value.

3 Pruning

A straightforward approach to pruning is by thresholding. A high threshold results in a set of disconnected skeleton branches correspond to the shape parts with slowly varying width. If the threshold is set too low, the skeleton will include branches due to noise. The following algorithm gets around this difficulty by extending the branches of the thresholded skeleton along the gray skeleton into the thicker parts of the shape, possibly forming junctions, without picking up extraneous branches belonging to noise.

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 towards the more blobby or thicker 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*

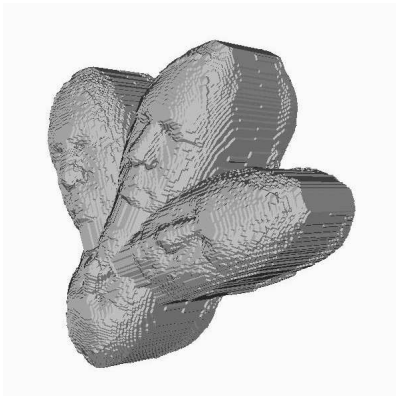

Figure 2: A 3D image

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 to be adjusted in order to ensure that it includes a direction of decreasing ρ.

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. 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.

4 An Illustrative Example

The test shape is a multiheaded figure created from MRI slices of the human head shown in Fig. 2.[1] Figs. 3 and 4 depict several cross-sections of the gray skeleton (left column), the pruned skeleton with $\overline{\theta} = 70°$and $\underline{\theta} = 45°$, (middle

[1]I thank Professor Jackie Shen of University of Minnesota for providing the MRI slices and Professor Kaleem Siddiqi of McGill University for the 3D visualization depicted in Fig, 2.

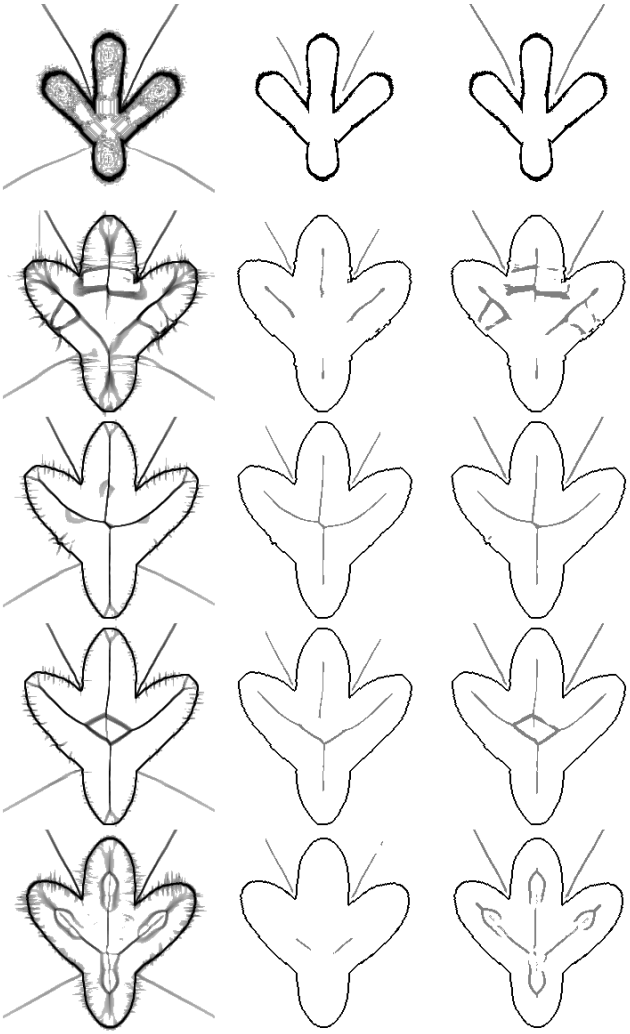

Figure 3: Gray skeleton and pruned Skeletons: Vertical sections

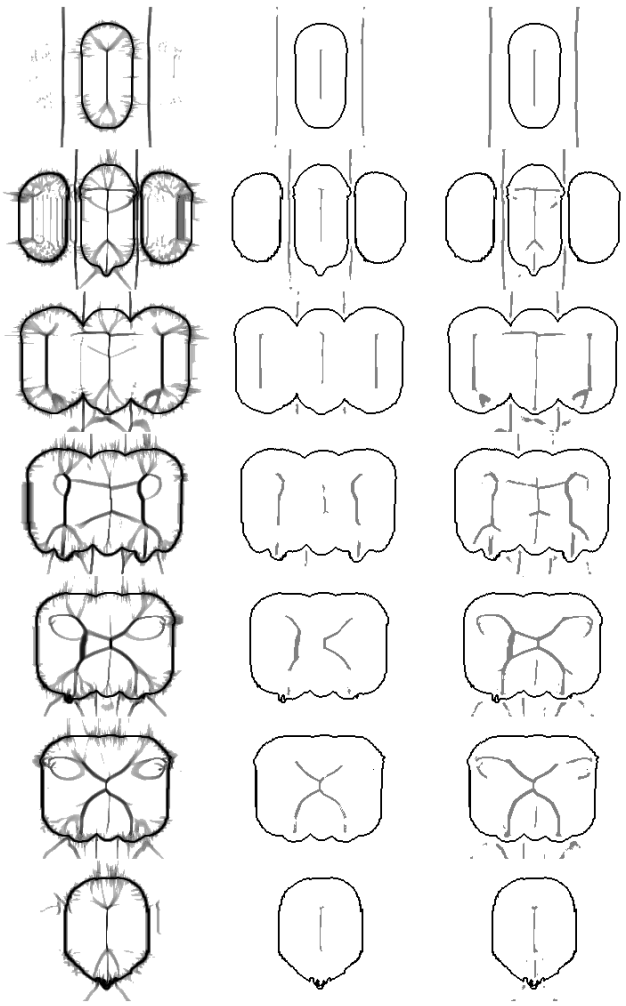

Figure 4: Gray skeleton and pruned skeletons: Horizontal sections

column), and the pruned skeleton with $\overline{\theta} = 70°$ and $\underline{\theta} = 30°$ (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. 3 depicts 5 successive vertical sections going from front to back. Fig. 4 depicts 7 successive horizontal sections proceeding from the top of the triple head 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, see §5.) With $\overline{\theta} = 70°$ and $\underline{\theta} = 30°$, the pruned skeleton has a single connected component inside the shape and a single component outside. Notice that as the shape get less cylinder-like (small φ), the gray skeleton gets fuzzier reflecting the difficulty in locating the skeleton accurately.

5 Regularization of the Skeleton Boundary

As mentioned in the Introduction, the boundary of the skeleton obtained by the method described above may have a tattered appearance due to noise in the data. It may also have gaps and and short isolated branches. To address these defects, we define a regularizing functional

$$E(K) = \int_K (c - \alpha) + \beta |\partial K| \tag{5}$$

where K is a subset of the gray skeleton Γ representing the pruned skeleton, $c = \cos\varphi$ is the cost function defined on the gray skeleton, α is a cost threshold, ∂K is the boundary of K and $|\partial K|$ is its length. If $\beta = 0$, E is minimized by setting K equal to the set of points of the gray skeleton G where $c \leq \alpha$. If $\beta > 0$, we may locally shorten ∂K even if it means extending K along G to include voxels where $c > \alpha$. Thus, minimization of E requires the techniques of curve evolution on a surface. Details of this approach are being worked out.

6 Appendix

We present an approach based on variational calculus to determine a pruned, smoothed skeleton directly by minimizing a functional. We exploit a basic identity that the distance function satisfies [9]. Let

$$u = \nabla(\frac{1}{2}\rho^2) \tag{6}$$

where ρ is the distance function. Since $\|\nabla\rho\| = 1$ and $u = \rho\nabla\rho$, $\|u\| = \rho$. It follows that u satisfies the identity

$$u = \nabla(\frac{1}{2}u \cdot u) \tag{7}$$

We have the following converse due to Gomes and Faugeras [9]:

Proposition: Suppose $u : R^n \to R^n$ satisfies Identity (6) almost everywhere and u is continuous at all points of the set $M = u^{-1}(0)$. Then, $u = \nabla(\frac{1}{2}\rho^2)$ where ρ is the distance from M.

Taking into account that the identity fails where u is discontinuous, namely, on the shape skeleton, we define the functional

$$\int_{R-K} \alpha \left\| \frac{1}{2}\nabla(u \cdot u) - u \right\|^2 + |K| \tag{8}$$

subject to the condition that $u = 0$ on the shape boundary. Here R is an open subset of R^n containing the shape, K is the discontinuity locus of u and $|K|$ is its volume (length, area, etc). The problem with this functional is that $\left\|\frac{1}{2}\nabla(u \cdot u) - u\right\|^2 = O(\rho^2)$ and hence, the penalty for violationg Identity (6) depends on its distance from ρ. To remedy this, we consider the normalized functional

$$E(u, K) = \int_{R-K} \alpha \left\| \nabla \|u\| - \frac{u}{\|u\|} \right\|^2 + |K| \tag{9}$$

Since minimization of E balances the cost of the modifying or removing a segment of K against the cost of violating Identity (6), the result is a pruned and smoothed skeleton K. The functional does not attempt to regularize the boundary of K; inclusion of such a term would make the functional considerably more difficult to implement. To apply the method of gradient descent, we use the Ambrosio-Tortorelli appoximation of E:

$$E_\lambda(u, v) = \int_R \left\{ \alpha \left\| \nabla \|u\| - \frac{u}{\|u\|} \right\|^2 (1 - v)^2 + \frac{\lambda}{2} \|\nabla v\|^2 + \frac{v^2}{2\lambda} \right\} \tag{10}$$

The corresponding gradient descent equations are:

$$\frac{\partial u}{\partial t} = \beta u + (1 - v)\nabla \|u\| \tag{11}$$

$$\frac{\partial v}{\partial t} = \nabla \cdot \nabla v - \frac{v}{\lambda^2} + \frac{2\alpha}{\lambda}(1 - v) \left\| \nabla \|u\| - \frac{u}{\|u\|} \right\|^2 \tag{12}$$

where

$$\beta = (1 - v) \left\{ \nabla \cdot \nabla \|u\| - \frac{u}{\|u\|} \right\} - 2 \left(\nabla \|u\| - \frac{u}{\|u\|} \right) \cdot \nabla v \tag{13}$$

We illustrate this approach for determining shape skeletons by means of an example of a 2D shape shown in Fig. 4. The top row shows the gray skeleton and the skeleton obtained by pruning the gray skeleton. The bottom row depicts the function v obtained using Eqs. 10 and 11 with two different values of α. Just as in the case of segmentation functionals, two objections may be raised against this approach: First, we have much less control over pruning and smoothing of the skeleton than when we use the method of gray skeletons; second, function v is spread out and we need a ridge finding method to locate the skeleton precisely by following the ridges of v. Nonetheless, it is of interest to formulate the problems of determining shape skeletons and of segmenting images within the same variational framework.

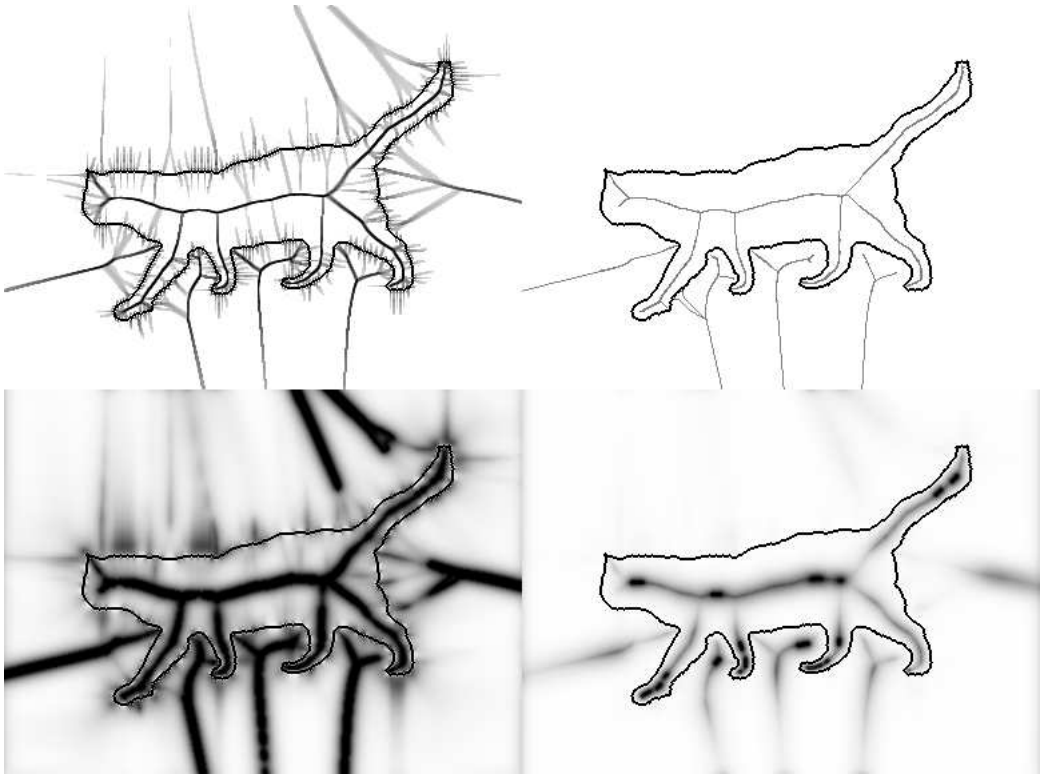

Figure 5: Top row: Gray skeleton and pruned skeleton
Bottom row: Function v with two different values of α

7 References

1. D. Adalsteinsson and J.A. Sethian: "The fast construction of extension velocities in level set methods", J. of Computational Physics, 1999, pp. 2-22.

2. L. Ambrosio and V.M. Tortorelli: "On the Approximation of Functionals depending on Jumps by Quadratic, Elliptic Functionals", Boll. Un. Mat. Ital. (1992).

3. N. Amenta, S. Choi and R. Kolluri: "The Power Crust", Proc. of the sixth ACM Symposium on Solid Modeling, pp.249-260, 2001.

4. S. Bouix and K. Siddiqi: "Divergence-based medial surfaces", Sixth European Conference on Computer Vision, Dublin, Ireland, June, 2000.

5. P. Dimitrov, J.N. Damon and K. Siddiqi: "Flux Invariants for Shape", Conference on Computer Vision and Pattern Recognition (CVPR'03), 2003.

6. P. Giblin and B.B. Kimia: "On local form and transitions of symmetry sets and medial axes and shocks in 2D", International Conference on Computer Vision, 1999, pp.385-391.

7. P. Giblin and B.B. Kimia: "A formal classification of 3D medial axis points and their local geometry', IEEE Trans. Pattern Analysis and Machine Intelligence, 2003.

8. J.D. Furst, S.M. Pizer, D.M. Eberly: "Marching Cores: A Method for Extracting Cores from 3D Medical Images", Proc. of the Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA), pp. 124-130, June, 1996.

9. J. Gomes and O. Faugeras: "Using the Vector Distance Functions to Evolve Manifolds of Arbitrary Codimension", Third International Conference on Scale Space and Morphology in Computer Vision, 2001.

10. B.B. Kimia, A.R. Tannenbaum, and S.W. Zucker: "Shapes, shocks and deformations, I", International Journal of Computer Vision, 15(3), 1995.

11. F. Leymarie and B.B. Kimia: "Symmetry-based representation of 3D data", International Conference on Image Processing, October, 2001, pp.581-584.

12. T.-L. Liu, D. Geiger and R.V. Kohn: "Representation and self-similarity of shapes", International Conference on Computer Vision, Mumbai, India, 1998.

13. G. Malandain and S. Fernandez-Vidal: "Euclidean Skeletons", Image and Vision Computing, v.16, pp.317-327, 1998.

14. C. Mantegazza and A.C. Mennucci: "Hamilton-Jacobi equations and distance functions on Riemannian manifolds", Appl. Math. Opt. v.47, no.1, 2003, pp.1-25.

15. J. Siddiqi, A. Shokoufandeh, S. Dickinson and S. Zucker: "Shock graphs and shape matching", International Journal of Computer Vision, 35(1), September, 1999, pp. 13-32.

16. K. Siddiqi, S. Bouix, A. Tannenbaum and S. Zucker: "The hamilton-jacobi skeleton", International Conference on Computer Vision, 1999, pp. 828-834.

17. J. Shah: "Gray Skeletons and Segmentation of Shapes", to appear in CVIU.

18. S. Tari and J. Shah: "Local symmetries of shapes in arbitrary dimension", Sixth International Conference on Computer Vision, 1998.

19. S.C. Zhu: "Stochastic jump-diffusion process for computing medial axes in markov random fields", IEEE Trans. Pattern Analysis and Machine Intelligence, 21(11), 1999.